COHESITY

Cohesity User Guide

Version 7.4.1

July 27, 2026

© 2026 Cohesity, Inc. All rights reserved.

Cohesity, the Cohesity logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, the Helios logo, DataGovern, SiteContinuity, DataHawk, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an "AS IS" basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.

No part of this documentation or any related software may be reproduced, stored, transmitted, or otherwise distributed in any form or by any means (electronic or otherwise) for any purpose other than the purchaser's personal use without the prior written consent of Cohesity, Inc. You may not use, modify, perform or display this documentation or any related software for any purpose except as expressly set forth in a separate written agreement executed by Cohesity, Inc., and any other use (including without limitation for the reverse engineering of such software or creating compatible software or derivative works) is prohibited, except to the extent such restrictions are prohibited by applicable law.

Published on July 27, 2026

Contents

Chapter 1: Cohesity 7.4.1 13
Getting Started 13
Explore Our Videos 13
Chapter 2: Supportability 15
Supported Software 16
Supported Workflows and External Targets 72
Supported Multitenancy Workflows 72
Supported Multitenancy Workflows for Hybrid Extender 73
Supported Multitenancy Workflows for Hybrid Extender Realms 73
Supported Cloud Spin and Other Workflows 121
Supported Platforms for Heterogeneous Cohesity cluster 121
Agent Compatibility Within a Cohesity cluster 125
Cohesity Cloud Edition Supportability 130
Replication Compatibility 152
Supportability for Qualifications 155
Chapter 3: Data Protection Key Concepts 157
Cohesity Cluster Architecture 157
Protection Overview 161
About Scale Out 173
About Cloud Archival 175
About Cloud Tier 178
About Tape Archival 179
About Fault Tolerance 183
Secure Software Development Practices 193
Cohesity Security Features 200
MegaFile 226
Data Retention for Audits 229
Dynamic Disk Support 232
Chapter 4: Cluster Administration 235
Initial Cluster Setup 236
Configure Settings 236
Customize Preferences 268
Cohesity Cluster Services 271
Software Update 282
Power Off and Power On a Cohesity Cluster 301
Search the Cluster and Take Action 315

Manage Storage Domains ... 317
Manage Nodes ... 342
Use Cohesity Data Cloud (Self-managed) in DoD Mode ... 386
About Node-to-Node Encryption ... 395
Quorum Protected Operations ... 398
Antivirus Software to Protect Cluster Data ... 398
Cohesity Licensing ... 403
Manage Apps ... 414
Networking Workflows ... 416
Manage Firewall Rules and Ports ... 487
Access Management ... 585
Web Server Certificates ... 700
Support Toolbox Overview ... 701
Available Service Health Checks ... 705
Groot uptime check ... 710
Available Third-party Health Checks ... 719
Available Backup Health Check ... 720
Available Vendor Health Checks ... 720
Available Software Health Checks ... 721
Available Networking Health Checks ... 728
Available Archival Health Checks ... 737
Available Archival Health Checks ... 752
Available Performance Health Checks ... 753
Available Upgrade Health Checks ... 754
Available Vendor Health Checks ... 764
Environment Setup ... 774
Access Cohesity Host OS ... 803
Hardware Replacement ... 804

Chapter 5: FortKnox Self-Managed ... 805

Key Concepts ... 805
Supported Workloads ... 806
Get Started ... 807
Convert a Cluster to Vault Cluster ... 811
Pair Primary Cluster to Vault cluster ... 813
Create Protection Policy for Vault Cluster from Primary Cluster ... 816
Register a Source on Primary Cluster ... 820
Create a Protection Group on Primary Cluster ... 820
Vaulting Status ... 824
Recover from Vault Cluster ... 826
Clean Room Recovery using FortKnox Self-Managed ... 829
Vaulting Scenario for Organizations ... 834

Considerations ... 835
Best Practices ... 837
Frequently Asked Questions ... 840

Chapter 6: Cohesity Data Cloud for NetBackup ... 843

Deployment Modes ... 843
Deployment Overview ... 844
Topologies and use cases ... 844
Plan and Prepare for Data Cloud with NetBackup ... 847
Limitations ... 848
Set Up the Cluster using the UI ... 849
Create Cluster using CLI ... 852
Get Started with Data Cloud for NetBackup ... 853
NetBackup Media Servers ... 854
Manage Storage Domains ... 856
Manage Storage Servers ... 860
Manage Logical Storage Units ... 861
Add Nodes to the Cluster ... 863

Chapter 7: Organizations (Multitenancy) ... 865

Service Providers (SPs) ... 866
Large Enterprises ... 866
Organization User Types ... 867
Organizations Subfeatures ... 868
Prerequisites and Considerations for Organizations ... 870
Privileges for Service Provider Admin and Organization Admin ... 871
Best Practices for Organizations ... 872
Enable Organizations (Multitenancy) on a Cluster ... 873
Add an Organization (Tenant) ... 874
Infrastructure for Organizations ... 876
Access Management for Organizations ... 908
Users and Groups for Organizations ... 916
Data Protection for Organizations ... 917
Impersonate an Organization ... 920
Deactivate an Organization ... 920
Activate an Organization ... 921
Delete an Organization ... 921
CloudArchive and CloudTier for Organizations ... 922
Replication for Organizations ... 923
Reports and Chargeback for Organizations ... 923
Supported Multitenancy Workflows ... 927
Tenant-Level View Allowlist ... 1213
Key Management Service (KMS) in Multitenancy ... 1214

Chapter 8: Manage Protection Policy 1217
Create or Edit a Standard Policy 1217
Manage Policies 1231
Customize Indexing 1233
Configure Pre & Post Scripts 1235
DataLock 1244
Extended Retention 1247
Chapter 9: SmartFiles (File and Object Services) 1263
SmartFiles (File and Object Services) Overview 1264
Ideal Use Cases for SmartFiles 1267
Fundamentals of SmartFiles 1269
Cohesity Views 1281
File Shares (SMB) 1343
Network File System (NFS) 1372
Objects (S3) 1393
SmartFiles Dashboard 1450
OpenStack Swift Storage Protocol 1467
NAS Tiering 1482
File Services Audit Logs 1505
Antivirus Software to Protect Cluster Data 1514
Access Management 1518
Chapter 10: Cloud Platforms 1549
AWS 1549
Troubleshooting 1617
Execute the CloudFormation Template 1630
Amazon S3 Entity Hierarchy 1714
What are Materialized Views 1775
Why Refresh is Needed After Recovery 1775
Materialized View Refresh Options 1775
Best Practices 1778
What are Materialized Views 1826
Why Refresh is Needed After Recovery 1826
Materialized View Refresh Options 1826
Best Practices 1829
Requirements 1930
Considerations 1930
Register SQL Server on Linux OS 1931
Azure 2093
GCP 2341
Chapter 11: S3-Compatible Storage 2525

Plan and Prepare for S3-Compatible Storage Protection 2525
Register and Manage S3-Compatible Storage 2527
Protect S3-Compatible Storage 2529
Recover S3-Compatible Storage Buckets 2537

Chapter 12: Virtual Machines 2541

About Encrypted VM Backup 2542
Plan and Prepare for VM Protection 2543
Register or Edit a Hypervisor Source 2552
Register a Storage Snapshot Provider 2569
Add or Edit a Protection Group for Virtual Servers 2577
Manage the VMware VM Backup Configuration 2597
Creating Application-Consistent Snapshots 2602
Cancel Protection Run at Object Level 2605
Migrate VMware VMs Using Cohesity 2606
Protect Migrated VMs or vCD 2610
Cohesity CDP for VMware 2613
Failover Protection Group for VMs 2626
Viewing VM Snapshot Types 2636
Truncate Logs in Microsoft Exchange Servers on VMs 2637
Configuring Custom Quiesce Scripts for Linux Virtual Machines 2638
Recover Virtual Machines 2639
Azure VMware Solution (AVS) 2706
VMware vCloud Director (vCD) 2710
OpenStack 2774

Chapter 13: Databases 2785

Amazon RDS and Amazon Aurora 2786
Cassandra 2795
Couchbase Using Connector Agent 2843
MongoDB 2857
MongoDB via Ops Manager 2878
Pinecone 2896
Microsoft SQL Server 2906
Oracle Database 3022
CockroachDB 3146
Epic Database Protection using Cohesity Physical Agents 3164
IBM DB2 3168
MySQL EE 3195
Percona MySQL on Linux 3213
SAP HANA 3232
SAP Sybase IQ 3264
SAP MaxDB on Linux 3286

SAP MaxDB on Windows ... 3309
SAP Sybase ASE on Linux ... 3331
SAP Sybase ASE on Microsoft Windows ... 3352
SAP on Oracle for Linux ... 3372
SAP on Oracle for Microsoft Windows ... 3399
PostgreSQL ... 3419

Chapter 14: NAS ... 3457

Plan and Prepare for NAS Protection ... 3457
Register or Edit NAS ... 3460
Protect NAS ... 3466
Create a Protection Group for NAS Volumes ... 3467
Manage the NAS Backup Configuration ... 3482
NAS File DataLock ... 3488
NetApp ONTAP ... 3490
File DataLock Modes ... 3530
Things to remember for File DataLock ... 3531
Add Inclusion ... 3532
Add Exclusion ... 3533
Considerations ... 3535
Nutanix Files ... 3568
Dell EMC Isilon ... 3593
Recover NAS Source ... 3619
Replicate NAS Data ... 3640
Archive NAS Data ... 3640

Chapter 15: Microsoft 365 ... 3643

Microsoft 365 Data Backup and Restore: Process Overview ... 3644
GCC High: Microsoft 365 Data Backup and Recovery ... 3645
Using Microsoft Graph APIs ... 3645
Prerequisites ... 3646
Register Microsoft 365 Sources ... 3664
Exchange Online Mailboxes ... 3684
Exchange Online Public Folders ... 3721
OneDrive for Business ... 3738
SharePoint Online ... 3760
Groups ... 3784
Teams ... 3803
Replicate Microsoft 365 Data ... 3834

Chapter 16: Physical Servers ... 3837

Plan and Prepare for Physical Server Protection ... 3838
Register or Edit a Physical Server ... 3844

Network Bandwidth and CPU Throttling for Physical Server 3859
Add or Edit a Protection Group for Physical Servers 3861
Manage the Physical Server Backup Configuration 3881
Start On-demand Multiple Protection Runs for Physical Server 3885
Source-Side Deduplication and Cache Optimization for Physical Server 3886
Protect Windows Failover Cluster 3887
Protect File Server Role deployed on Windows Server Failover Cluster 3888
Physical Server Recovery 3891

Chapter 17: MS Applications 3927

Active Directory 3927
MS Exchange Server 3944

Chapter 18: SAN Volume 3963

IBM Storage FlashSystem 3963
Pure Storage 3971
HPE Alletra Storage MP 3981

Chapter 19: Hadoop 3985

Core components 3985
Hadoop Features 3986
Plan and Prepare for Hadoop Protection 3986
Register a Hadoop Source 3992
Backup Hadoop 3994
Recover Hadoop 4003

Chapter 20: About the Remote Adapter 4007

Recommendations for Dedupe-focused Backups with RMAN 4010
Add or Edit a Remote Adapter Protection Group 4011
About SAP HANA Database Backups 4015
SAP on Oracle 4064

Chapter 21: Kubernetes 4115

Plan and Prepare for Protecting Kubernetes Clusters 4115
Register and Manage Kubernetes Cluster 4120
Upgrade Velero and Datamover 4129
Backup Kubernetes Cluster Resources, PVCs, and Namespaces 4132
Recover Kubernetes 4149
Red Hat OpenShift Virtualization 4166

Chapter 22: Private Cloud Platforms 4233

VMware Cloud Foundation (VCF) Automation 4233

Chapter 23: About Bare Machine Recovery 4245

CBMR and CoBMR Comparison ... 4246
CBMR ... 4246
CoBMR ... 4251

Chapter 24: Manage External Targets ... 4261
CloudArchive Direct ... 4262
Archive Object Lock ... 4270
Archive Data Over AWS PrivateLink ... 4273
Archive Data Over Azure Private Endpoints ... 4274
Register an External Target ... 4279
AWS Security Token Service Support ... 4362
Kerberos Authentication for NAS External Target ... 4364
QStar Tape Configuration ... 4371
Unregister an External Target ... 4371

Chapter 25: Disaster Recovery ... 4373
About Replication ... 4373
Replication and Remote Access Setup ... 4391
About CloudRetrieve ... 4403
About Failover Ready ... 4408
Disaster Recovery Using CloudEdition ... 4409

Chapter 26: Manage Sources ... 4411
Manage Existing Sources Using Actions Menu Options ... 4412
Maintenance Mode ... 4412
Register Applications Running on Virtual Servers ... 4413
Manage Registered Applications Running on Virtual Servers ... 4413
Manage Cohesity Installed Agents Running on Virtual Servers ... 4413
Manage Existing Physical Servers Using the Actions Menu ... 4414
Health Check for SQL Sources ... 4414

Chapter 27: Manage Protection Groups ... 4419
Cancel Replication and Archive Copy Tasks of a Protection Run ... 4421
Delete Protection Groups, Snapshots and Objects ... 4422
Protection Group Run Issues ... 4427
Add or Remove a Legal Hold to a Snapshot ... 4427
Failover a Protection Group ... 4429

Chapter 28: Manage Recovery ... 4441
Use Wildcard * in Search for Recovery ... 4441
Cancel a Recovery ... 4442

Chapter 29: Cohesity Agents ... 4443
What are Cohesity Agents? ... 4443

What Do Cohesity Agents Do? 4443
Cohesity Agents Deployment in Different Scenarios 4444
Install and Manage Cohesity Agents 4445
Cohesity Java Agent Upgrade 4445
Install and Manage the Agent on Windows Servers 4447
Install and Manage the Agent on Linux Servers 4457
Install and Manage the Agent on Hyper-V Linux VMs 4481
Install and Manage the Cohesity Java Agent on AIX Servers 4483
Install and Manage the Agent on Solaris Servers 4487
Install and Manage the Agent on HP-UX Servers 4489
Deploy Cohesity Agent from the Cohesity UI 4492
Troubleshoot Cohesity Agent 4511

Chapter 30: NoSQL and Hadoop Service 4517

Supported Software 4517
Features 4518
Prerequisites 4519
Process Overview 4519
Install NoSQL and Hadoop Service 4520
Run the NoSQL and Hadoop Service 4522
Verify the Status of NoSQL and Hadoop Service 4523
Using NoSQL and Hadoop Data sources 4523
Post Installation Procedures 4523
Managing NoSQL and Hadoop Services 4539
Additional Configuration for NoSQL and Hadoop Services 4546
Monitor the NoSQL and Hadoop Service Health 4552
FAQs on NoSQL and Hadoop Services 4560

Chapter 31: Test & Dev (Cloning) 4563

Test & Dev Overview 4564
Clone Views 4565
Clone VMs 4567
Clone Databases 4583
Test & Dev Using CloudSpin 4589
Clone vCloud Director (vCD) VMs 4589

Chapter 32: Monitoring 4591

Performance 4592
Monitor Performance of Workloads 4595
Configure and Test SNMP 4602
Storage Statistics 4608
Health Dashboard 4621
Reports 4820

Audit Logs 4868
Advanced Diagnostics 4893
Chapter 33: Virtual Edition 4909
Cohesity Virtual Edition 4910
Initial Virtual Edition Cluster Setup 4911
Configure vSphere Availability for Virtual Edition for VMware 4911
Extend the Capacities of a Virtual Edition for VMware 4912
Chapter 34: Cohesity Cloud Edition 4923
Cohesity Cloud Edition on AWS 4923
Cohesity Cloud Edition on Azure 4938
Cohesity Cloud Edition on GCP 4963
Cohesity Cloud Edition on OCI 4978
Chapter 35: Tools 4981
Using the Cohesity CLI 4981
Cluster Health Check 5137
Support Toolbox 5140
Using the Secure Shell 5140
Cohesity REST API 5143
Chapter 36: Supplemental Information 5147
Open Source in Cohesity Products 5147
Explore Our Videos 5148
Helios Documentation 5149
Cohesity Hardware Data Sheets and Compliance Guide 5149
Hardware Replacement 5149
Technical Guides 5149
Cohesity Support 5149
Glossary 5150
Documentation Archive 5154
Documentation Feedback 5155

Chapter 1: Cohesity 7.4.1

Cohesity is at the forefront of AI-powered data security and management, revolutionizing how organizations protect, manage, and derive value from their data across the data center, edge, and cloud. With a robust ecosystem of partners, Cohesity offers comprehensive solutions that defend against cybersecurity threats through immutable backups, AI-based threat detection, and rapid recovery at scale. Cohesity's innovative approach ensures data security and efficient management, whether delivered as a service, self-managed, or through a Cohesity-powered partner. Dive into our documentation to explore the full breadth of our transformative solutions.

> **Note:** Features and functionalities herein may become subject to a separate license requirement/fee (even if free during an initial period).

Getting Started

If you are ready to:

- Set up the Cohesity cluster for the first time, see Hardware Universe
- Create users, see Users and Groups
- Protect objects, see Protection Overview
- View new features, see What's New?
- The list of "Fixed Issues" for all the releases are available exclusively in the Fixed Issues Online Help topic under the Release Notes. This information is not included in the offline downloadable PDF version of the Help.

Explore Our Videos

You can learn about Cohesity products and solutions by watching videos on the Cohesity How-to Videos page.

Chapter 2: Supportability

This chapter includes the following topics:

- Supported Software 16
- Supported Workflows and External Targets 72
- Supported Multitenancy Workflows 72
- Supported Multitenancy Workflows for Hybrid Extender 73
- Supported Multitenancy Workflows for Hybrid Extender Realms 73
- Supported Cloud Spin and Other Workflows 121
- Supported Platforms for Heterogeneous Cohesity cluster 121
- Agent Compatibility Within a Cohesity cluster 125
- Cohesity Cloud Edition Supportability 130
- Replication Compatibility 152
- Supportability for Qualifications 155

Supported Software

Microsoft Hyper-V

Hyper-V	Versions	Guest OS	Notes
Hyper-V Standalone, SCVMM & Cluster	2012R2, 2016, 2019, 2022, 2025	For the list of supported guest OS , see the supported operating system of Physical Servers. The following guest OS are also supported: • Oracle Linux 8.x, 7.x, 6.x • FreeBSD 12-12.1, 11.1-11.3, 11.0, 10.0-10.3, 9.1-9.3, 8.4	• For agent-based file level restores only 64-bit guest operating systems are supported. • Cohesity does not support the backup of Hyper-V VMs with shared disks. • Cohesity supports Hyper-V in the Desktop and Server Core installation modes. • Cohesity supports Hyper-V in the Nano Server installation mode with Windows Server 2016. Nano Server 2019 is used as a container and you cannot install Microsoft Hyper-V inside the container. It is primarily used for non-containerized installations. • Azure Local 23H2 is supported only on Standalone and Cluster. **Note:** This is an early-access feature. Contact your Cohesity account team to enable Azure Local 23H2.
	Azure Local 23H2		

Note: Azure Local was previously known as Azure Stack HCI.

Nutanix Acropolis Hypervisor

Application	AOS Version	File System	Notes
AHV (Acropolis Hypervisor)	5.1.1, 5.5.3.1, 5.8.2, 5.10.x, 5.15, 5.16, 5.17.x, 5.18.1, 5.19.1, 5.20, 6.5.x, 6.8, 6.10.x, 7.0.1, 7.1, 7.3.x	XFS, NTFS, Ext4, Ext3, Ext2, BTrFS	Cohesity supports application-consistent backups, VM recovery, and file and folder recovery in Nutanix Acropolis Hypervisors.

Nutanix Prism Central

Application	Versions	Note
Nutanix Prism Central	7.3.x, 7.5.x, 2024.2.x, 2024.3.x	Cohesity supports Nutanix Prism Central starting from 7.2.2_u2.

Nutanix Cloud Clusters (NC2)

AHV	Notes
Nutanix Cloud Clusters on AWS	Nutanix Cloud Clusters on AWS is supported with the following Cohesity cluster types: • Physical cluster • Virtual Edition cluster

AHV	Notes
Nutanix Cloud Clusters on Azure	Nutanix Cloud Clusters on Azure is supported with the following Cohesity cluster types: • Physical cluster • Virtual Edition cluster • Cohesity Cloud Edition (legacy) on Azure • Cohesity Cloud Edition on Azure
Nutanix Cloud Clusters on GCP	Nutanix Cloud Clusters on GCP is supported with the following Cohesity cluster types: • Physical cluster • Virtual Edition cluster

oVirt

RHV	File System	Notes
4.0, 4.1, 4.2, 4.3, 4.4	XFS, NTFS, Ext4, Ext3, Ext2, BTrFS	Cohesity supports RHV only until version 6.8.1. While you can still protect your RHV sources in 7.0 or later releases, the Cohesity Support team will not provide assistance.

OLVM (registered as RHV: oVirt Manager)	File System	Notes
4.5	XFS, NTFS, Ext4, Ext3, Ext2	

> **Note:** Using iSCSI/FC storage requires assistance from Cohesity Support.

VMware

vCenter, vSphere, ESXi	Virtual Machine Hardware	Guest OS	Virtual Disk Transport Method	vRealize Orchestrator and vRealize Automation	Notes
9.0	7, 8, 9, 10, 11, 13, 14, 15, 17, 18, 19, 20, 21	For the list of guest OS supported for VM, file-based, and volume-based backups, see Physical Servers. For the list of guest OS supported for Indexing of File Systems, see Supported File Systems for Indexing.	By default, the Cohesity cluster uses the NBDSSL (encrypted) transport method. NBDSSL uses SSL to encrypt the data.	See this link about the Cohesity vRO plugin	For Linux OSs: LVM2 For LVM2 volumes of types Linear, Striped, Mirrored or Mirrored + Striped, only file downloads are supported. RAID is not supported. File recovery to source requires VMware Tools (or Open VM Tools), rsync and NFS client. You must use a license with the vSphere Storage API when instantly recovering a VMware VM from a Cohesity Cluster to a vSphere datastore. For more information, see the Recovering a VMware VM returns the message "License not available to

vCenter, vSphere, ESXi	Virtual Machine Hardware	Guest OS	Virtual Disk Transport Method	vRealize Orchestrator and vRealize Automation	Notes
					perform the operation" KB article. The qualification of VMware vSphere 9.0 does not include the support of the new features and functions introduced in VMware vSphere 9.0. Cohesity recommends increasing the default HTTPS connection limit after upgrading to vSphere 7.0 and later. For details, see the KB article. For the list of data objects and properties Cohesity clusters can protect, see the KB article. vSAN ESA is now supported.
8.0 U3, 8.0 U2, 8.0 U1, 8.0	7, 8, 9, 10, 11, 13, 14, 15, 17, 18, 19, 20, 21				
7.0 U3	7, 8, 9, 10, 11, 13, 14, 15, 17, 18, 19				
7.0 U2	7, 8, 9, 10, 11, 13, 14, 15, 17, 18, 19				
7.0 U1	7, 8, 9, 10, 11, 13, 14, 15, 17, 18				
7.0	7, 8, 9, 10, 11, 13, 14, 15, 17				
6.7 U3	7, 8, 9, 10, 11, 13, 14, 15				
6.7 U2	7, 8, 9, 10, 11, 13, 14, 15				
6.7	7, 8, 9, 10, 11, 13, 14				

> **Note:** All vSphere or ESXi patches are supported as long as the base version or update itself is supported.

For more information, see ESXi/ESX hosts and compatible virtual machine hardware versions list (2007240).

VMware vCloud Director

vCloud Director	VMware vCenter version qualified with vCD	vCD Extension Version	Notes
10.6.1	8.0U, 7.0U	See this link.	• Contact your Cohesity account team to enable vCD 10.4.1. • You can recover vCD objects in the same location or in an alternate location in vCD. • You can recover vCD VMs and standalone VMs to VMware vCenter and vice-versa.
10.6			
10.5.1	7.0.1		
10.5			
10.4.2			
10.4.1	7.0		
10.4			

VMware Cloud Infrastructure

VMware Cloud Infrastructure	
VMware Cloud (VMC) on AWS	
Azure VMware Solution (AVS)	
Amazon Elastic VMware Service (Amazon EVS)	Supported only on Cloud Edition on AWS
Google Cloud VMware Engine (GCVE)	

> **Note:** VMware Cloud (VMC) on AWS supports stretched cluster Software-Defined Data Center (SDDC) which provides an additional layer of resiliency.

VMware Data Protection Transport Methods

Deployment	Datastore Type	NBDSSL Transport	SAN Transport	HotAdd Transport	Notes
ESXi, vCenter	VMFS	Yes	Yes	No	
	NFS	Yes	N/A	No	
	vSAN	Yes	N/A	No	
	vVOL	Yes	N/A	No	
VMware Cloud on AWS (VMC)	vSAN	No	N/A	Yes	HyX is a prerequisite for VMC
Azure VMware Solution (AVS)	VMFS	Yes	No	No	AVS Gen 2 supports Pure CBS
	NFS	Yes	N/A	No	AVS Gen 2 supports NetApp Files
	vSAN	Yes	N/A	No	
Google Cloud VMware Engine (GCVE)	vSAN	Yes	N/A	No	
Amazon Elastic VMware Service Amazon (EVS)	vSAN	Yes	N/A	No	EVS is treated as vCenter

OpenStack

> **Note:**
> - OpenStack protection is provided as a third-party integration with Storware.
> - For supported OpenStack distributions, Storware features, backup strategies, and version details, refer to the Storware OpenStack documentation.

OpenStack	Supported Versions
Canonical OpenStack (Sunbeam)	2024.1 LTS (Sunbeam)
Canonical OpenStack (Charmed)	2024.1 (Charmed)
Platform9 Managed OpenStack	PCD 2025.10-180
Red Hat OpenStack Platform	17.1, 18.0
Community OpenStack	• Wallaby • Antelope • Caracal • Dalmatian • Epoxy

\Virtual Edition

Hypervisor	Cluster Type
VMware	• Single-Node Cluster • Multi-Node Cluster
Microsoft Hyper-V	• Single-Node Cluster
Nutanix AHV	• Single-Node Cluster • Multi-Node Cluster

Hypervisor	Cluster Type
Red Hat Enterprise Linux	• Single-Node Cluster • Multi-Node Cluster
Red Hat OpenShift	• Single-Node Cluster • Multi-Node Cluster
AlmaLinux	• Single-Node Cluster • Multi-Node Cluster
OCI Linux KVM	• Multi-Node Cluster

Physical Servers

> **Note:** Applicable to both file-based and volume-based backups.

For information on prerequisites and considerations, see Prerequisites and Considerations.

Operating System	File System	CPU Architecture	Notes
Windows 10 Desktop Edition Windows 11	NTFS	x64	
Windows Server 2025 Core Windows Server 2022 Core Windows Server 2019 Core Windows Server 2016 Core	NTFS	x64	

Operating System	File System	CPU Architecture	Notes
Windows Server 2025 Windows Server 2022 Windows Server 2019 Windows Server 2016 Windows Server 2012 Windows Server 2012 R2	NTFS	x64	Cohesity supports Microsoft failover clusters that are compatible with these Windows Server versions.

Operating System	File System	CPU Architecture	Notes
CentOS 6.0+, 7.0 - 7.9, 8.0, 8.3, 9.0	ext3, ext4, xfs, cvfs, zfs	x64	
Oracle Linux (OEL) 6.x, 7.0 - 7.7, 7.9, 8.0, 8.2 - 8.10, 9.0 - 9.7, 10.1		x64	
RHEL 6.4, 6.7+, 7.0 - 7.9, 8.0 - 8.10, 9.0 - 9.8, 10.0-10.2		x64 PowerPC (RHEL 8, 8.6, 9.0) Linux on z/Architecture (9.x) ARM (aarch64) (RHEL 9.x)	• Only file-based backup is supported for Linux on PowerPC. • Cohesity supports the little-endian architecture for PowerPC. • Cohesity supports full indexing and file-level recovery for RHEL 10 VMs with XFS filesystems. The XFS parser handles the nrext64 feature, which is enabled by default in RHEL 10.

Operating System	File System	CPU Architecture	Notes
Rocky Linux 8.7 - 8.10, 9.0 - 9.7		x64	With physical clusters indexing and file-recovery through downloads are supported. This is an early-access feature. Contact Cohesity Support to use this functionality.
Alma Linux 8.10, 9.0 - 9.7		x64	
SUSE Linux Enterprise Server 11 SP4, 12 SP3,12 SP4, 12 SP5, 15, 15 SP3, 15 SP4, 15 SP5, 15 SP6, SP7, 16		x64 PowerPC (SLES 12 SP5, 15 SP3, 15 SP4, SP5, SP6, SP7)	
openSUSE 15.1, 15.3, 15.4		x64	
Ubuntu 14.04, 16.04, 18.04, 19.04, 20.04, 22.04, 23.04, 24.04		x64	
Debian 9.6, 10, 11.2, 12.x, 13		x64	
Amazon Linux 2		x64	Amazon Linux 2 is supported only for file-based backup.
Amazon Linux 2023		x64	Amazon Linux 2023 is supported only for file-based backup.

Operating System	File System	CPU Architecture	Notes
AIX 7.1 TL3 SP0 or later AIX 7.2 AIX 7.3	JFS, JFS2, CDRFS, VxFS, MMFS, GPFS, NFS (over the network), AutoFS	IBM POWER7, POWER8, POWER9, POWER10	AIX 7.3 is supported only for file-based backup.

Operating System	File System	CPU Architecture	Notes
	Note: JFS(*) - Cohesity has tested the support of JFS using the Cohesity Java Agent only on AIX version 6.1 64-bit TL7. Since AIX OS upgrades are backward compatible, JFS file system is also supported on AIX versions 7.1 64-bit and 7.2 64-bit.		

Operating System	File System	CPU Architecture	Notes
Solaris 10, 11	ZFS (native Solaris 11), UFS (native Solaris 10), VxFS, NFS (over the network), SMBFS (over the network)	SPARC V9 x86_64	
HP-UX 11i v3 (B.11.31)	VxFS file system (Version 6.1.0.000)	Intel Itanium (IA64)	See Install and Manage the Agent on HPUX Servers for installation steps.

> **Note:** ReFS file system is supported on applicable Windows Server editions, through file-based protection.

Storage Snapshot Providers

Snapshot Provider	Version	Protocol	Notes
Cisco HyperFlex	3.5+	NBDSSL	
Pure Storage FlashArray	4.10.9, 5.x, 6.x	iSCSI, FC	
Nimble Storage	5.0, 5.1, 5.2	iSCSI, FC	

Snapshot Provider	Version	Protocol	Notes
HPE Alletra 5000	6.1, 6.x	iSCSI, FC	Registering HPE Alletra 5000 and 6000 as a source is an Early Access feature.
HPE Alletra 6000	6.x	iSCSI, FC	
HPE Alletra Storage MP	10.4.x	iSCSI, FC	Registering HPE Alletra Storage MP and B10000 as a source is an Early Access feature.
HPE Alletra Storage MP B10000	10.6.x	iSCSI, FC	
NetApp ONTAP	9.x	NFS	
Nutanix	6.5.x	iSCSI	Nutanix as Snapshot Provider for VMware VMs (VMware ESXi on Nutanix AOS).
IBM Storage FlashSystem	8.6.0.2 or later	iSCSI, FC	For more information on the supported hardware models, see Hardware Models.

SAN

SAN	Version
Pure Storage FlashArray	6.1.0 or later
IBM Storage FlashSystem	Software Version: 8.6.0.2 or later. Hardware Models: • IBM Storage FlashSystem 5045 model • IBM Storage FlashSystem 5200 model • IBM Storage FlashSystem 5300 model • IBM Storage FlashSystem 9100 model • IBM Storage FlashSystem 9200 model • IBM Storage FlashSystem 9500 model • IBM SAN Volume Controller SA2 and SV2 models • IBM SAN Volume Controller SV3 model

Microsoft SQL Server

Supported OS Versions	Microsoft SQL Server Versions								
	Express 2016	Express 2017	Express 2019	Express 2022	2016	2017	2019	2022	2025
Microsoft Windows 2016 (64-bit)	✅	✅	✅	✅	✅	✅	✅	✅	✅
Microsoft Windows 2019 (64-bit)	✅	✅	✅	✅	✅	✅	✅	✅	✅

Supported OS Versions	Microsoft SQL Server Versions								
	Express 2016	Express 2017	Express 2019	Express 2022	2016	2017	2019	2022	2025
Microsoft Windows 2022 (64-bit)	✅	✅	✅	✅	✅	✅	✅	✅	✅
Microsoft Windows 2025 (64-bit)	✅	✅	✅	✅	✅	✅	✅	✅	✅
Red Hat Enterprise Linux (RHEL) 7.x	❌	❌	❌	❌	❌	❌	✅	❌	❌

- MS SQL Server 32-bit deployment is not supported with the agent-based approach. It is recommended to use the Native SQL dump to Cohesity SMB for backups.
- MS SQL Server 64-bit running on Windows 64-bit is supported.
- SQL databases and logs residing on Clustered Shared Volumes (CSV) can be protected using Cohesity's VDI based backup method.
- FILESTREAM databases can be protected using Cohesity's File, VDI, and Volume based backup method.
- Memory-optimized databases can be protected using Cohesity's VDI based backup method.
- MS SQL Server on Red Hat Enterprise Linux (RHEL) 7.x is an Early Access feature. Contact Cohesity Account team to enable this feature.
- Only VDI-based backups are currently supported on Linux OS.
- Amazon Custom RDS supports protection of SQL Server 2022 and SQL Server 2019.
- For recovering a database to an Amazon RDS Custom SQL Server, the recovery destination path must be under the **D:\rdsdbdata** folder.

SAP HANA (Cohesity Connector Agents)

Supported OS Versions	SAP HANA Version	
	SAP HANA 2.0 SPS07 or earlier	SAP HANA 2.0 SPS08
	On-Premise	On-Premise
	Platform Edition	Platform Edition
Red Hat Enterprise Linux (RHEL) 7.x	✅	❌
Red Hat Enterprise Linux (RHEL) 8.x	✅	❌
Red Hat Enterprise Linux (RHEL) 9.x	✅	✅
SUSE Linux Enterprise Server (SLES) 12.x	✅	❌
SUSE Linux Enterprise Server (SLES) 15.x	✅	✅

Note:

- SAP HANA Express Edition is not supported.
- Supported on x86-64 or IBM PowerPC (ppc64le) based systems running Red Hat Enterprise Linux or SUSE Linux Enterprise Server.
- Refer to the SAP Note **2235581** for more specific information on SAP HANA and OS version mappings.

SAP S/4HANA (Cohesity Connector Agents)

Supported OS Versions	SAP S/4HANA 2021, 2022, and 2023 Versions	
	SAP HANA 2.0 SPS07	SAP HANA 2.0 SPS08
	On-Premise	On-Premise
	Public Edition	Public Edition
Red Hat Enterprise Linux (RHEL) 8.x	✅	❌
Red Hat Enterprise Linux (RHEL) 9.x	✅	✅
SUSE Linux Enterprise Server (SLES) 12.x	✅	❌
SUSE Linux Enterprise Server (SLES) 15.x	✅	✅

SAP HANA (Remote Adapter Plug-in)

Supported OS Versions	SAP HANA Version	
	SAP HANA 2.0 SPS07 or earlier	SAP HANA 2.0 SPS08
Red Hat Enterprise Linux (RHEL) 7.x	✅	❌
Red Hat Enterprise Linux (RHEL) 8.x	✅	❌
Red Hat Enterprise Linux (RHEL) 9.x	✅	✅
SUSE Linux Enterprise Server (SLES) 12.x	✅	❌

Supported OS Versions	SAP HANA Version	
	SAP HANA 2.0 SPS07 or earlier	SAP HANA 2.0 SPS08
SUSE Linux Enterprise Server (SLES) 15.x	✓	✓

> **Note:** Supported on x86-64 or IBM PowerPC (ppc64le) based systems running Red Hat Enterprise Linux or SUSE Linux Enterprise Server.

SAP on Oracle (Remote Adapter Plug-in)

Supported OS Versions	SAP on Oracle Versions				
	11gR2	12cR1	12cR2	18c	19c
Red Hat Enterprise Linux (RHEL) 7	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 8	✓	✓	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 12	✓	✓	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 15	✓	✓	✓	✓	✓
Oracle Linux 7	✓	✓	✓	✓	✓

Supported OS Versions	SAP on Oracle Versions				
	11gR2	12cR1	12cR2	18c	19c
Oracle Linux 8	✅	✅	✅	✅	✅

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Oracle Linux.

SAP on Oracle (Cohesity Connector Agents)

Supported OS Versions	SAP on Oracle Versions				
	11gR2	12cR1	12cR2	18c	19c
Red Hat Enterprise Linux (RHEL) 7	✅	✅	✅	✅	✅
Red Hat Enterprise Linux (RHEL) 8	✅	✅	✅	✅	✅
SUSE Linux Enterprise Server (SLES) 12	✅	✅	✅	✅	✅
SUSE Linux Enterprise Server (SLES) 15	✅	✅	✅	✅	✅
Oracle Linux 7	✅	✅	✅	✅	✅

Supported OS Versions	SAP on Oracle Versions				
	11gR2	12cR1	12cR2	18c	19c
Oracle Linux 8	✓	✓	✓	✓	✓
Microsoft Windows Server 2016	✓	✓	✓	✓	✓
Microsoft Windows Server 2019	✓	✓	✓	✓	✓

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Oracle Linux, or Microsoft Windows Server.

CockroachDB

Supported OS Versions	CockroachDB Versions	
	20.1	20.2
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓
SUSE Linux Enterprise Server (SLES) 12	✓	✓

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux or SUSE Linux Enterprise Server.

SAP Sybase IQ

Supported OS Versions	SAP Sybase IQ Versions					
	16.0 SP11 PL28	16.1 SP03 PL9	16.1 SP04 PL01	16.1 SP04 PL09	16.1 SP05 PL12	16.1 SP05 PL14
Red Hat Enterprise Linux (RHEL) 6.x	✓	✓	✓	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓	✓	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 8.x	✗	✗	✗	✗	✓	✗
Red Hat Enterprise Linux (RHEL) 9.x	✗	✗	✗	✗	✗	✓
SUSE Linux Enterprise Server (SLES) 12	✗	✗	✗	✗	✓	✗

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux or SUSE Linux Enterprise Server.

SAP Sybase ASE

Supported OS Versions	SAP Sybase ASE Versions									
	16.0 SP03 PL09	16.0 SP03 PL10	16.0 SP03 PL11	16.0 SP03 PL14	16.0 SP04	16 SP04 PL03 HF1	16 SP04 PL04	16 SP04 PL05	16 SP04 PL06	16.1
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓	✓	✓	✗	✓	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓	✓	✓	✗	✓	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 12.x	✓	✓	✓	✓	✓	✗	✓	✓	✓	✗
SUSE Linux Enterprise Server (SLES) 15.x	✓	✓	✓	✓	✓	✗	✓	✓	✓	✓

Supported OS Versions	SAP Sybase ASE Versions									
	16.0 SP03 PL09	16.0 SP03 PL10	16.0 SP03 PL11	16.0 SP03 PL14	16.0 SP04	16 SP04 PL03 HF1	16 SP04 PL04	16 SP04 PL05	16 SP04 PL06	16.1
Microsoft Windows Server 2016	✓	✓	✓	✓	✓	✗	✓	✓	✓	✗
Microsoft Windows Server 2019	✓	✓	✓	✓	✓	✗	✓	✓	✓	✗
Microsoft Windows Server 2022	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Microsoft Windows Server.

SAP MaxDB

Supported OS Versions	SAP MaxDB Versions		
	7.9.10	7.9.10.12	7.9.11.5
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓

Supported OS Versions	SAP MaxDB Versions		
	7.9.10	7.9.10.12	7.9.11.5
SUSE Linux Enterprise Server (SLES) 15.x	✗	✓	✗
Microsoft Windows Server 2016	✓	✓	✗
Microsoft Windows Server 2019	✓	✓	✗

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Microsoft Windows Server.

IBM DB2

Supported OS Versions	IBM DB2 Versions			
	10.5.x	11.1.x (maintain)	11.5.x	12.1.x
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 12.x	✓	✓	✓	✗
SUSE Linux Enterprise Server (SLES) 15.x	✓	✓	✓	✓

Supported OS Versions	IBM DB2 Versions			
	10.5.x	11.1.x (maintain)	11.5.x	12.1.x

Note:

- Supported on x86-64 based systems running Red Hat Enterprise Linux or SUSE Linux Enterprise Server.
- For more specific versions of IBM DB2 and OS support mapping, refer to IBM's website for further details of which combination of OS and DB2 versions is supported by IBM.

Oracle MySQL

Supported OS Versions	Oracle MySQL Versions		
	8.0.x EE (Commercial Edition)	8.4.x	9.2.x
Red Hat Enterprise Linux (RHEL) 7.x	✓	✗	✗
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 9.x	✗	✓	✓
Red Hat Enterprise Linux (RHEL) 10.x	✗	✗	✓
SUSE Linux Enterprise Server (SLES) 15.x	✗	✗	✓
Oracle Linux 7.x	✓	✗	✗
Oracle Linux 8.x	✓	✗	✗

Supported OS Versions	Oracle MySQL Versions		
	8.0.x EE (Commercial Edition)	8.4.x	9.2.x
Rocky Linux 8.x	✗	✓	✓
Rocky Linux 9.x	✗	✓	✓

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux, Oracle Linux, or Rocky Linux.

Percona MySQL

Supported OS Versions	MySQL Versions		
	8.0.x	8.4.x	9.1.x
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux.

PostgreSQL

Supported OS Versions	PostgreSQL Versions									
	9.6	10.x	11.x	12.x	13.x	14.x	15.x	16.x	17.x	18.x
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
Red Hat Enterprise Linux (RHEL) 10.x	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 12.x	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
SUSE Linux Enterprise Server (SLES) 15.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Rocky Linux 8.x	✗	✗	✗	✓	✓	✓	✓	✓	✓	✗
Rocky Linux 9.x	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓

Supported OS Versions	PostgreSQL Versions									
	9.6	10.x	11.x	12.x	13.x	14.x	15.x	16.x	17.x	18.x
Rocky Linux 10.x	✗	✗	✗	✗	✗	✓	✗	✗	✓	✓
Ubuntu 20.04.x 22.04.x 24.04.x	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓
ALMA LINUX 8	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
ALMA LINUX 9	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Rocky Linux.

Percona Distribution for PostgreSQL

Supported OS Versions	Percona Distribution for PostgreSQL Version			
	14.x	15.x	16.x	17.x
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓	✓

EDB Postgres

Supported OS Versions	EDB Postgres Versions					
	11.x	12.x	14.x	15.x	16.x	17.x
Red Hat Enterprise Linux (RHEL) 7.1.x	✓	✓	✓	✓	✗	✗
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 15.x	✗	✗	✗	✗	✓	✓
Rocky Linux 8.x	✗	✗	✗	✗	✗	✓
Rocky Linux 9.x	✗	✗	✗	✗	✗	✓

> **Note:** Supported on x86-64 based systems running Red Hat Enterprise Linux or SUSE Linux Enterprise Server.

Oracle Adapter

Supported OS Versions	Oracle Versions						
	11gR2	12cR1	12cR2	18c	19c	21c	23c
AIX 7.3	✓	✓	✓	✓	✓	✓	✓
AIX 7.1 TL5	✓	✓	✓	✓	✓	✓	✓
AIX 7.2 TV2+	✓	✓	✓	✓	✓	✓	✓
CentOS 6.x	✓	✓	✓	✓	✓	✓	✓
CentOS 7.x	✓	✓	✓	✓	✓	✓	✓
CentOS 8.x	✓	✓	✓	✓	✓	✓	✓
CentOS 9.x	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 6.x	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 7.x	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 8.x	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 9.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 6.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓	✓	✓	✓	✓

Supported OS Versions	Oracle Versions						
	11gR2	12cR1	12cR2	18c	19c	21c	23c
SUSE Linux Enterprise Server (SLES) 11 SP4	✓	✓	✓	✗	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 12 SP2.x	✓	✓	✓	✓	✓	✓	✓
SUSE Linux Enterprise Server (SLES) 15 SP1.x	✓	✓	✓	✓	✓	✓	✓
Microsoft Windows Server 2012	✓ (11.2.0.4)	✓	✓	✓	✓	✗	✗
Microsoft Windows Server 2012R2	✗	✓	✓	✓	✓	✓	✗
Microsoft Windows Server 2016	✗	✓	✓	✓	✓	✓	✗
Microsoft Windows Server 2019	✗	✗	✓	✓	✓	✓	✗
Microsoft Windows Server 2022	✗	✗	✗	✗	✓	✓	✗
Solaris 5.11	✓	✓	✓	✓	✓	✓	✓

Note: Oracle 23c version is supported from Cohesity cluster version 6.8.2 and later.

Oracle SBT Plugin

Supported OS Versions	Oracle Versions						
	11gR2	12cR1	12cR2	18c	19c	21c	23c
AIX 7.3	✓	✓	✓	✓	✓	✓	✓
AIX 7.2 TV2+	✓	✓	✓	✓	✓	✓	✓
AIX 7.1 TL5	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 6.x	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 7.x	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 8.x	✓	✓	✓	✓	✓	✓	✓
Oracle Linux 9.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 6.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 7.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 8.x	✓	✓	✓	✓	✓	✓	✓
Red Hat Enterprise Linux (RHEL) 9.x	✓	✓	✓	✓	✓	✓	✓
Microsoft Windows Server 2012	✓	✓	✓	✓	✓	✓	✓
Microsoft Windows Server 2012R2	✓	✓	✓	✓	✓	✓	✓
Microsoft Windows Server 2016	✓	✓	✓	✓	✓	✓	✓
Microsoft Windows Server 2019	✓	✓	✓	✓	✓	✓	✓

Supported OS Versions	Oracle Versions						
	11gR2	12cR1	12cR2	18c	19c	21c	23c

> **Note:**
> - For RHEL 6 and OEL 6, ensure that the Cohesity version is 6.9 and above.
> - Oracle 23c version is supported from Cohesity cluster version 6.8.2 and later.

NoSQL and Hadoop Platforms

Platform	Supported Versions
Apache Cassandra	4.1, 4.0
Cloudera Data Platform (Private Cloud base)	7.3.1, 7.1.9, 7.1.7
Couchbase	7.2.6, 7.2, 7.1, 7.0
Datastax Enterprise	6.8.x, 5.1.x **Note:** Supports IPv6 on DataStax Enterprise 6.x only.
MongoDB	8.0, 7.0

MongoDB via Ops Manager

> **Note:** Only Enterprise MongoDB clusters are supported.

Supported OS Versions	MongoDB Ops Manager Version: 8.0.10 - 8.0.23 MongoDB Cluster/Deployment Version: 7.x_ent, 8.x_ent
Red Hat Enterprise Linux (RHEL) 8.x	✅
Red Hat Enterprise Linux (RHEL) 9.x	✅

NAS

Vendor	Version	Notes
Dell EMC Isilon	8.0, 8.1, 8.2, 9.0.x - 9.4.x, 9.5.x, 9.7.x, 9.8.x, 9.9.x, 9.10.x, 9.11.x, 9.12.x, 9.13.x, 9.14.x	• Cohesity supports backing up SMB version 2 and 3 shares. • Cohesity uses SMB version 2 for NetApp backups by default. • Cohesity does not support restoring NAS volumes and data using SMB v1.
NetApp ONTAP Cluster-Mode		
Cohesity file-runner based backup approach	8.2, 8.3, 9.3, 9.5, 9.6, 9.7, 9.8, 9.9.1, 9.10.x, 9.11.x, 9.12.x, 9.13.x, 9.14.x, 9.15.x, 9.16.x, 9.17.x, 9.18.x	
NetApp SnapDiff-based backup approach	9.8, 9.9.1, 9.10.x, 9.11.x, 9.12.x, 9.13.x, 9.14.x, 9.15.x, 9.16.x, 9.17.x, 9.18.x	
Nutanix Files	FSM and Files Server 5.1.1, Files Manager 5.1	
Pure Storage FlashBlade	2.2.4+, 3.0+, 3.2.x, 4.1.x, 4.4.x, 4.5.x, 4.8.1	
Elastifile EFS	3.1.0.31+	
IBM GPFS	4.2.x, 5.0.x, 5.1.x	

Vendor	Version	Notes
Generic NAS	All vendors supporting NFS (v3, v4.1) and SMB (v1+, v2+, v3+)	

Supported File Systems for Indexing

Cohesity supports indexing for the file systems listed in this table:

Protection Group Type	Operating System	Supported File Systems
Virtual Server	Supported OS of Physical Server	ext2, ext3, ext4, xfs, and ntfs
	Linux	btrfs **Note:** Cohesity does not support the indexing of btrfs filesystem created on multiple devices using the `mkfs.btrfs` command.
Physical Server (Block-based)	Supported OS	ext2, ext3, ext4, xfs, and ntfs
	Linux	btrfs **Note:** Cohesity does not support the indexing of btrfs filesystem created on multiple devices using the `mkfs.btrfs` command.
Physical Server (File-based)	Supported OS	Agnostic of the file system
NAS	Not applicable	NFS, SMB/CIFS

Microsoft Applications

Application	Version	Availability	Notes
Active Directory	Windows 2025 Windows 2022 Windows 2019 Windows 2016 Windows 2012, 2012 R2		Kerberos and NTLM authentication are supported when accessing shares on an AD-joined Cohesity cluster.
Exchange on virtual servers	Exchange 2019 Exchange 2016 Exchange SE 2025	DAG and standalone	Running on supported Windows servers. Requires Ontrack® software and license for granular recovery of Microsoft Exchange mailbox objects.
Exchange on physical servers	Exchange 2019 Exchange 2016 Exchange SE 2025	DAG and standalone	Ontrack® PowerControls™ 9.2 or later is required for granular recovery. Cohesity recommends to use the latest version Ontrack® PowerControls™.
Microsoft SharePoint	Microsoft SharePoint Server 2016 Microsoft SharePoint Server 2019	Ontrack® PowerControls™ 9.2 or later is required for granular recovery*. Cohesity recommends to use the latest version of Ontrack® PowerControls™.	

*Ontrack® PowerControls™ license validation requires a 2-way TCP communication over HTTPS (port 443) using TLSv1.2 or later.

Veritas Cluster Server

Veritas Cluster Server	Notes
7.4.1 running on RHEL 7.x	Cohesity supports backup and restore of Oracle DBs running on VCS.
7.4.3 running on OEL 7	

Red Hat High Availability (Pacemaker) Cluster

Pacemaker version	Notes
2.x on RHEL 8.x 2.x on RHEL 9.x	Cohesity supports Pacemaker clusters starting with the 2.1 release.

LDAP

LDAP Version
3

QStar Archive Storage Manager

QStar ASM	Notes
7.0.8061	Cohesity supports QStar ASM 7.0.8061 installed on Windows Server 2016, 2019, and 2022.

BMR

CBMR

CBMR Version	OS Version	Version	CPU Architecture	Notes
9.6	**Windows**			
	Windows Server	2022, 2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	8.0, 8.3	x64	
	OEL	8.0, 8.2 - 8.6, 9.0	x64	
	RHEL	8.0 - 8.10, 9.0 - 9.4	x64	
	Debian	9.6, 10, 11.2	x64	
	Ubuntu	22.04, 24.04	x64	
9.4	**Windows**			
	Windows Server	2022, 2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	7.0 - 7.9, 8.0, 8.3, 9.0	x64	
	OEL	6.x, 7.0 - 7.7, 7.9, 8.0 - 8.6, 9.0	x64	
	RHEL	6.4, 6.7+, 7.0 - 7.9, 8.0, 8.2 - 8.6, 9.0	x64	
	Debian	9.6, 10, 11.2	x64	
	Ubuntu	22.04, 24.04	x64	

CBMR Version	OS Version	Version	CPU Architecture	Notes
9.3	**Windows**			
	Windows Server	2022, 2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	7.0 - 7.9, 8.0, 8.3, 9.0	x64	
	OEL	6.x, 7.0 - 7.7, 7.9, 8.0 - 8.6, 9.0	x64	
	RHEL	6.4, 6.7+, 7.0 - 7.9, 8.0, 8.2 - 8.6, 9.0	x64	
	Debian	9.6, 10, 11.2	x64	
9.2	**Windows**			
	Windows Server	2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	7.0 - 7.9, 8.0, 8.3, 9.0	x64	
	OEL	6.x, 7.0 - 7.7, 7.9, 8.0 - 8.6, 9.0	x64	
	RHEL	6.4, 6.7+, 7.0 - 7.9, 8.0, 8.2 - 8.6, 9.0	x64	
	Debian	9.6, 10, 11.2	x64	

CBMR Version	OS Version	Version	CPU Architecture	Notes
8.6	**Windows**			
	Windows Server	2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	7.0 - 7.9, 8.0, 8.3, 9.0	x64	
	OEL	6.x, 7.0 - 7.7, 7.9, 8.0 - 8.6, 9.0	x64	
	RHEL	6.4, 6.7+, 7.0 - 7.9, 8.0, 8.2 - 8.6, 9.0	x64	
	Debian	9.6, 10, 11.2	x64	

CoBMR

> **Note:** Cohesity supports both file-based and block-based backups starting CoBMR version 9.3 and later.

CoBMR Version	OS Version	Version	CPU Architecture	Notes
9.6.5	**Windows**			
	Windows Server	2025	x64	All operating support statements from 9.6 continue to apply to 9.6.5 version as well.

CoBMR Version	OS Version	Version	CPU Architecture	Notes
9.6	**Windows**			
	Windows Server	2022, 2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	8.0, 8.3	x64	
	OEL	8.0, 8.2 - 8.6, 9.0	x64	
	RHEL	8.0 - 8.10, 9.0 - 9.4	x64	
	SLES	12 SP5, 15 SP3	x64	
	Debian	9.6, 10, 11.2	x64	
	Ubuntu	22.04, 24.04	x64	

CoBMR Version	OS Version	Version	CPU Architecture	Notes
9.5	**Windows**			
	Windows Server	2022, 2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	7.0 - 7.9, 8.0, 8.3, 9.0	x64	
	OEL	6.x, 7.0 - 7.7, 7.9, 8.0 - 8.6	x64	
	RHEL	6.4, 6.7+, 7.0 - 7.9, 8.0, 8.2 - 8.6, 9.0	x64	
	SLES	12 SP5, 15 SP3	x64	
	Debian	9.6, 10, 11.2	x64	
	Ubuntu	22.04, 24.04	x64	

CoBMR Version	OS Version	Version	CPU Architecture	Notes
9.4	**Windows**			
	Windows Server	2022, 2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	7.0 - 7.9, 8.0, 8.3, 9.0	x64	
	OEL	6.x, 7.0 - 7.7, 7.9, 8.0 - 8.6	x64	
	RHEL	6.4, 6.7+, 7.0 - 7.9, 8.0, 8.2 - 8.6, 9.0	x64	
	SLES	12 SP5, 15 SP3	x64	
	Debian	9.6, 10, 11.2	x64	
	Ubuntu	22.04, 24.04	x64	

CoBMR Version	OS Version	Version	CPU Architecture	Notes
9.3	**Windows**			
	Windows Server	2022, 2019, 2016, 2012, 2012 R2	x64	
	Windows Desktop edition	10, 11	x64	
	Linux			
	CentOS	7.0 - 7.9, 8.0, 8.3, 9.0	x64	
	OEL	6.x, 7.0 - 7.7, 7.9, 8.0 - 8.6, 9.0	x64	
	RHEL	6.4, 6.7+, 7.0 - 7.9, 8.0, 8.2 - 8.6, 9.0	x64	
	Debian	9.6, 10, 11.2	x64	

Kubernetes Clusters

Kubernetes Cluster	Version
Vanilla Kubernetes	1.14 - 1.21, 1.23 - 1.36
Red Hat OpenShift	4.4 - 4.20
VMware Tanzu Kubernetes Grid (TKGm, multi-cloud version)	1.2.1, 1.3.0, 1.4.0, 1.5.3, 1.5.4, 1.6.0, 1.6.1
VMware Tanzu Kubernetes Grid Service	1.19

Kubernetes Cluster	Version
Rancher (RKE2)	1.32, 1.34, 1.35
Nutanix Kubernetes Platform (NKP)	2.17.x
Azure Kubernetes Service (AKS)	1.33

> **Note:**
> - Cohesity supports Red Hat OpenShift only on the x86_64 architecture.
> - Backup and recovery of Azure Kubernetes Service (AKS) is supported only on Cohesity Cloud Edition on Azure.
> - Rancher (RKE2), Azure Kubernetes Service (AKS), and Nutanix Kubernetes Platform (NKP) managed Kubernetes clusters are not listed as separate options during registration. These clusters are supported as upstream Kubernetes distributions. Select **Mainline** to proceed with the registration.

KubeVirt

Platform	Version
Red Hat OpenShift	4.17 - 4.19

> **Note:** Cohesity supports Red Hat OpenShift only on the x86_64 architecture.

Velero

Velero Version	Kubernetes Version	Notes
v1.15.0	v1.18 and later	
V1.8.1	Earlier than v1.18	

S3-Compatible Storage

Cohesity has qualified the following S3-compatible storage systems for data protection. However, any S3-compatible storage that adheres to the AWS S3 protocol is supported for data protection.

S3-Compatible Storage	Version
Dell EMC ECS	3.8.1.2.140142.3df35ef5802
Pure Storage Flashblade	3.3.7, 4.5.6, 4.8.1
NetApp StorageGrid	11.9.0.3
MinIO	20241107005220.0.0
IBM Cloud Object Storage	N/A

Cloud Platforms

Supported Cloud Platforms for Backup and Recovery

Cloud Platforms	Workloads	Notes
AWS	EC2	
	S3 Bucket	
	RDS for MySQL	Ingest-based protection is supported only on Cloud Edition on AWS.
	RDS for PostgreSQL	
	RDS for Oracle	
	Oracle Database	
	RDS for Microsoft SQL Server	
	Aurora MySQL	
	Aurora PostgreSQL	
	RDS Custom for Oracle	
	RDS Custom for SQL Server	
	DocumentDB	Supported only on Cloud Edition on AWS.
	DynamoDB	
	RedShift Cluster	

Cloud Platforms	Workloads	Notes
Azure	VM	
	MySQL	Supported only on Cloud Edition on Azure.
	Cosmos DB for Cassandra	
	Cosmos DB for MongoDB	
	Cosmos DB for NoSQL	
	Cosmos DB for Table	
	SQL Database	
	Oracle Database	
	SQL Managed Instance	
	Table Storage	
	PostgreSQL	
	Blob Storage and Data Lock Storage	

Cloud Platforms	Workloads	Notes
GCP	VM	
	BigQuery	Supported only on Cloud Edition on GCP.
	BigTable	
	Spanner	
	Firestore	
	Oracle Database	
	Cloud SQL for MySQL	
	Cloud SQL for PostgreSQL	
	Cloud SQL for AlloyDB PostgreSQL	
	Cloud SQL for SQL Server	
OCI	OCI VM	Oracle adapter backup and recovery supports both configurations (Image Copy and SBT adapter.)
	OCI Bare Metal	
	OCI ExaCS	

Supported Operating Systems for Cloud VM Backup and Recovery

Operating System	File System	Notes
Windows 10 Desktop Edition **Note:** Cohesity supports the Enterprise Edition of Windows 10 Desktop Edition.	NTFS	Backup using Cloud Snapshot Management (CSM) is independent of operating systems. Full backup and restore using native snapshot is independent of operating systems. The operating systems mentioned in the Operating System column are applicable only for File/Folder Level Recovery.
Windows 2019 64-bit Windows 2016 64-bit Windows 2012, 2012 R2 64-bit	NTFS	
CentOS 5.10 64-bit, 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit	ext3, ext4, xfs	
OEL 5.11 - 5.8 64-bit, 6.x 64-bit, 7.0 - 7.9 64-bit, 8.2 64-bit		
RHEL 5.11 - 5.8 64-bit, 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit		
SUSE Linux Enterprise Server 11 SP4 64-bit, 12 SP4 64-bit, 12.3 64-bit, 15.0 64-bit		
openSUSE 15.1 64-bit		
Ubuntu 14.x 64-bit, 16 64-bit, 18 64-bit, 19.x 64-bit, 20 64-bit		
Debian 6.0.0-6.0.8 64 bit, 7.0.0-7.8.0 64 bit, 8.0.0 64 bit, 9.6 64-bit, 10 64-bit		
Fedora Server 19-21 64 bit		

Failover from VMware to AWS

Operating System	File System
Windows 2022 64-bit Windows 2019 64-bit Windows 2016 64-bit Windows 2012, 2012 R2 64-bit Windows 10 Desktop Edition	NTFS
Ubuntu 14.x 64-bit, 16 64-bit, 18 64-bit, 20 64-bit	ext3, ext4, xfs
RHEL 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit	
CentOS 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit	
OEL 5.11 - 5.8 64-bit, 6.x 64-bit, 7.0 - 7.9 64-bit	
SUSE Linux Enterprise Server 11 64 bit, 12 64 bit	
Debian 6.0.0-6.0.8 64 bit, 7.0.0-7.8.0 64 bit, 8.0.0 64 bit	
Fedora Server 19-21 64 bit	

Failover from VMware to Azure

Operating System	File System
Windows 2022 64-bit Windows 2019 64-bit Windows 2016 64-bit Windows 2012, 2012 R2 64-bit Windows 10 Desktop Edition	NTFS
RHEL 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit	ext3, ext4, xfs
CentOS 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit	

Azure Stack Hub

Supported Azure Stack Hub Versions
• 2102 • 2005

CloudSpin

Source	Destination	Supported OS
VMware	AWS	**BIOS boot** Microsoft Windows Server 2019, 2016, 2012, 2012 R2. Microsoft Windows 7 (Home, Professional, Enterprise, Ultimate), Windows 8 (Home, Professional, Enterprise), Windows 8.1 (Professional, Enterprise), Windows 10 (Home, Professional, Enterprise, Education). CentOS 5.1-5.11, 6.1-6.8, 7.0-7.9, 8.0-8.2. Debian 6.0.0-6.0.8, 7.0.0-7.8.0, 8.0.0. Oracle Linux 5.10-5.11, 6.1-6.10, 7.0-7.6. RHEL 5.1-5.11, 6.1-6.9, 7.0-7.9, 8.0-8.2. SUSE Linux Enterprise Server 11 with Service Pack 1, Service Pack 2 and Service Pack 3, SUSE Linux Enterprise Server 12 with Service Pack 1 - Service Pack 5, SUSE Linux Enterprise Server 15 with Service Pack 1 and Service Pack 2. Ubuntu 12.04, 12.10, 13.04, 13.10, 14.04, 14.10, 15.04, 16.04, 16.10, 17.04, 18.04, and 20.04. **UEFI boot** Windows Server 2012, 2016, 2019 Ubuntu 16, 18

Source	Destination	Supported OS
VMware	Azure	**BIOS boot** Windows Server 2012, 2016, 2019. CentOS 5.10 64-bit, 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit. RHEL 5.8 - 5.11 64-bit, 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit. **UEFI boot** Windows Server 2012, 2016, 2019.
Hyper-V	Azure	**BIOS and UEFI boot** Windows 2012, 2016, 2019. CentOS 5.10 64-bit, 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit. RHEL 5.8 - 5.11 64-bit, 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit. Ubuntu 14, 16, 18, 20.

Browsers

The Cohesity Dashboard supports the following browsers.

Browser	Version
Chrome	Newest version
Edge	
Firefox	
Safari	

Supported Workflows and External Targets

For more information, see Supported Workflows and External Targets in the online help.

Supported Multitenancy Workflows

For more information, see Supported Multitenancy Workflows in the online help.

Supported Multitenancy Workflows for Hybrid Extender

For more information, see Supported Multitenancy Workflows for Hybrid Extender in the online help.

Supported Multitenancy Workflows for Hybrid Extender Realms

This sections lists the supported multitenancy workflows for Hybrid Extender Realms.

Protection Group and External Targets Workflows (Hybrid Extender Realms)

- Virtual Server (VMs) and Physical Server Workflows (Hybrid Extender Realms)
  - Protection Based Workflow
  - Archival Based Workflow
  - Replication Based Workflow
- Databases Workflows (Hybrid Extender Realms)
  - Protection Based Workflow for MS SQL VM
  - Archival Based Workflow for MS SQL VM
  - Replication Based Workflow for MS SQL VM
  - Protection Based Workflow for MS SQL Volume-base, File-based, and Native Backup
  - Archival Based Workflow for MS SQL Volume-base, File-based, and Native Backup
  - Replication Based Workflow for MS SQL Volume-base, File-based, and Native Backup
  - Protection Based Workflow for Oracle
  - Archival Based Workflow for Oracle
  - Replication Based Workflow Oracle

Legend for Charts Below:

Yes = Supported

No = Not Supported

N/A = Not Applicable

CA = Cluster Administrator

TA = Tenant Administrator

Create = Create for an Org

Delete = Delete from Cluster

Assign = Assign to an Org

Un-assign = Un-assign from an Org

Virtual Server (VMs) and Physical Server Workflows (Hybrid Extender Realms)

Protection Based Workflow

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
REGISTRATION	Register a source	Supported (CA, TA)	Supported (TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	N/A	N/A	N/A	N/A	N/A
	Auto-protect new objects	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CA,TA)
	Full backup with CBT	N/A	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CA,TA)
	Incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A	N/A	N/A
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	CloudSpin to AWS and Azure	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	No	No	No	No
	Recover a SQL DB to the original server and instance	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A	N/A	N/A	N/A
	Recover DB to Alternate Instance with overwite	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A	N/A	N/A	N/A
	Instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No
	Crash-consistent recover for apps	Supported (CA,TA)	No	No	No	No
CLONE	Clone an object of a job run	Supported (CA,TA)	No	No	No	No
	Clone all the objects of a job run	Supported (CA,TA)	No	No	No	No
	Clone db to FCI instance	N/A	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A
TEARDOWN	Tear down objects	Supported (CA,TA)	No	No	No	No
	Tear down instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for a backed-up object	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download files/folders	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file	Supported (TA)	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	Supported (CA,TA)	N/A	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	Supported (CA,TA)	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	Supported (CA,TA)	N/A	N/A	N/A	N/A

Archival Based Workflow

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A
	Archive a full backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (CA,TA)	No	No	No	No
	Recover an entire backup job run from an archive to a source	Supported (CA,TA)	No	No	No	No
	Instant volume mount from an archive	Supported (CA,TA)	No	No	Supported (CA,TA)	No
	Crash-consistent recover for apps from an archive	Yes	No	No	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	No	No	No	No
	Clone all objects of a job run from an archive	Supported (CA,TA)	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	Supported (CA,TA)	No	No	No	No
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	Supported (CA,TA)	No	No	No	No
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders from an archive	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for files or folders from an archive snapshot	No	No	No	Supported (CA,TA)	Supported (CA,TA)
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	No	Supported (CA,TA)	Supported (CA,TA)	No	No
	Download n VMX file from an archive	No	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	No	No	No	No	No
	Clone all objects of a job run from an archive	No	No	No	No	No
	clone DB to FCI	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	No	No	No	No	No
	Recover all objects of job run from an archive to a source	No	No	No	No	No
	Instant volume mount from an archive	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Search files or folders from the metadata of a retrieved job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse files or folders from a retrieved snapshot	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse files or folders from the metadata of a downloaded job run	Supported (CA,TA)	No	No	No	No
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover files from an archive to a source (to the same and different sources)	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from an archive	No	N/A	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	No	N/A	N/A	N/A	N/A

Replication Based Workflow

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	N/A	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	Recover an entire backup job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A	N/A	N/A	N/A
	Instant volume mount of replicated data	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Crash-consistent recover for apps from a remote cluster	Supported (CA,TA)	No	No	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	Clone all objects of a job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	No
	Crash-consistent clone for apps from a remote cluster	Supported (CA,TA)	No	No	No	

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	No	No	No	No
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for replicated object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover replicated files from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from a remote cluster	Supported (TA)	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	Supported (CA,TA)	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	Supported (CA,TA)	N/A	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A	N/A	N/A

Databases Workflows (Hybrid Extender Realms)

Protection Based Workflow for MS SQL VM

Categories	Workflows	Sources
		MS SQL VM
REGISTRATION	Register a source	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)
	Refresh Application registration	Yes using VM edit workflow
	Auto-protect new objects	Supported (CA,TA)
BACKUP	Incremental backup with CBT	Yes as log-based incremental
	Full backup with CBT	Supported (CA,TA)
	Incremental backup with SQL server differential	N/A
	Full backup without CBT	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A
	Backup, archive and replicate now	Supported at DB and server level both
	CloudSpin to AWS and Azure	N/A

Categories	Workflows	Sources
		MS SQL VM
RECOVER	Recovering an entity in the Job Run	Yes
	Recover an entire backup job run	Yes as crash-consistent VM not as SQL restore
	Recover a SQL DB to the original server and instance	No
	Recover a SQL DB to the original server but a different instance	No
	Recover a SQL DB to a different server that has the same SQL version	No
	Recover DB to Alternate Instance with overwite	No
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	No
	Instant volume mount	No
	Crash-consistent recover for apps	Supported (CA,TA)
CLONE	Clone an object of a job run	No
	Clone all the objects of a job run	No
	Clone db to FCI instance	No
	Crash-consistent recover for apps	Supported (CA,TA)
DB Migration	SQL DB Migration without auto sync	N/A
	SQL DB Migration with auto sync	N/A
TEARDOWN	Tear down objects	Supported (CA,TA)
	Tear down instant volume mount	Supported (CA,TA)

Categories	Workflows	Sources
		MS SQL VM
CANCEL	Cancel a running job run	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)
FILE RECOVERY	Search for files or folders	Supported (CA,TA)
	Browse for a backed-up object	Supported (CA,TA)
	Download files/folders	Supported (TA)
	Recover files	Supported (CA,TA)
	Download a VMX file	Supported (TA)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A
	Register AWS source and clone VMs to AWS	N/A

Archival Based Workflow for MS SQL VM

Categories	Workflows	Sources
		MS SQL VM
ARCHIVAL	Archive an incremental backup	Supported (Not for log backup)(CA,TA)
	Archiving an incremental backup with SQL server differential	N/A
	Archive a full backup	Supported (Not for log backup)(CA,TA)

Categories	Workflows	Sources
		MS SQL VM
RECOVER	Recover a single object of a job run from an archive to a source	Supported (No Databases, Only Files and VM)(CA,TA)
	Recover an entire backup job run from an archive to a source	No
	Instant volume mount from an archive	No for SQL (use VM or physical)
	Crash-consistent recover for apps from an archive	Yes for files & VMs but not DBs
DB Migration	SQL DB Migration without auto sync	N/A
	SQL DB Migration with auto sync	N/A
CLONE	Clone objects of a job run from an archive	No
	Clone all objects of a job run from an archive	No
	clone db to FCI	N/A
	Crash-consistent clone for apps from an archive	No
CANCEL	Cancel a running archival task	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	No
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)

Categories	Workflows	Sources
		MS SQL VM
FILE RECOVERY	Search for files or folders from an archive	Supported (CA,TA)
	Browse for files or folders from an archive snapshot	No
	Download files and folders from an archive	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	No
	Download n VMX file from an archive	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported for volume based job(CA)
	Download metadata for all runs of a job	Supported for volume based job(CA)
	Download a snapshot (archived data) for a job run	Supported for volume based job(CA)
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported for volume based job
	Clone all objects of a job run from an archive	No
	clone DB to FCI	N/A
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported for volume based job
	Recover all objects of job run from an archive to a source	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical, volume based)(CA,TA)

Categories	Workflows	Sources
		MS SQL VM
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported for volume based job
	Search files or folders from the metadata of a retrieved job run	Supported for volume based job (CA,TA)
	Browse files or folders from a retrieved snapshot	Supported for volume based job (CA,TA)
	Browse files or folders from the metadata of a downloaded job run	Supported for volume based job (CA,TA)
	Download files and folders from an archive	Supported for volume based job(TA)
	Recover files from an archive to a source (to the same and different sources)	Supported for volume based job (CA,TA)
	Download a VMX file from an archive	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A

Replication Based Workflow for MS SQL VM

Categories	Workflows	Sources
		MS SQL VM
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Yes as log-based incremental
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)

Categories	Workflows	Sources
		MS SQL VM
FAILOVER	Failover an inactive Job on a remote cluster	No
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)
	Recover an entire backup job run from a remote cluster	Yes as crash-consistent, not as SQL restore
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	Supported (CA,TA)
	Instant volume mount of replicated data	Supported (CA,TA)
	Crash-consistent recover for apps from a remote cluster	Supported (CA,TA)
DB Migration	SQL DB Migration without auto sync	N/A
	SQL DB Migration with auto sync	N/A
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)
	Clone all objects of a job run from a remote cluster	Supported as Crash Consistent VM not as SQL Clone(CA,TA)
	clone db to FCI	N/A
	Crash-consistent clone for apps from a remote cluster	Supported (CA,TA)
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)
CANCEL	Cancel a running replication task	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)

Categories	Workflows	Sources
		MS SQL VM
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)
	Browse for replicated object on a remote cluster	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Supported (TA)
	Recover replicated files from a remote cluster	Supported (CA,TA)
	Download a VMX file from a remote cluster	Supported (TA)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A
	Register AWS source and clone VMs to AWS	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A
	Create a snapshot in Azure and replicate to another region	N/A

Protection Based Workflow for MS SQL Volume-base, File-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Auto-protect new objects	Supported (CA,TA)	Supported exclusion is not supported (CA,TA)	Supported exclusion is not supported (CA,TA)
BACKUP	Incremental backup with CBT	Yes as log-based incremental	Supported (CA,TA)	N/A
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)	N/A
	Incremental backup with SQL server differential	N/A	N/A	Yes
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	Supported (CA,TA)	Supported (CA,TA)
	Backup, archive and replicate now	Supported at DB and server level both	Supported at DB and server level both	Supported at DB and server level both
	CloudSpin to AWS and Azure	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
RECOVER	Recovering an entity in the Job Run	No	Yes at DB level	Yes at DB level
	Recover an entire backup job run	No	No	No
	Recover a SQL DB to the original server and instance	No	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to the original server but a different instance	No	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has the same SQL version	No	Supported (CA,TA)	Supported (CA,TA)
	Recover DB to Alternate Instance with overwite	No	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	No	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount	No	N/A	N/A
	Crash-consistent recover for apps	No	N/A	N/A
CLONE	Clone an object of a job run	No	Supported (CA,TA)	No
	Clone all the objects of a job run	No	N/A	N/A
	Clone db to FCI instance	N/A	Supported (CA,TA)	N/A
	Crash-consistent recover for apps	No	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
DB Migration	SQL DB Migration without auto sync	N/A	Supported (CA,TA)	N/A
	SQL DB Migration with auto sync	N/A	Supported (CA,TA)	N/A
TEARDOWN	Tear down objects	N/A	Supported (CA,TA)	N/A
	Tear down instant volume mount	N/A	N/A	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported at server level (CA,TA)	Supported at server level (CA,TA)
FILE RECOVERY	Search for files or folders	No	N/A	N/A
	Browse for a backed-up object	Supported (CA,TA)	N/A	N/A
	Download files/folders	Supported (TA)	N/A	N/A
	Recover files	Supported (CA,TA)	N/A	N/A
	Download a VMX file	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A

Archival Based Workflow for MS SQL Volume-base, File-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
ARCHIVAL	Archive an incremental backup	Supported (Not for log backup)(CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	Supported (CA,TA)
	Archive a full backup	Supported (Not for log backup)(CA,TA)	Supported (CA,TA)	Supported (CA,TA)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (No Databases, Only Files)(CA,TA)	Supported at DB level (CA,TA)	Supported at DB level (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No	No
	Instant volume mount from an archive	No for SQL (use VM or physical)	N/A	N/A
	Crash-consistent recover for apps from an archive	Yes for files but not DBs	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
CLONE	Clone objects of a job run from an archive	No	Yes	No
	Clone all objects of a job run from an archive	No	No	No
	clone db to FCI	N/A	Yes	N/A
	Crash-consistent clone for apps from an archive	No	N/A	N/A
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	No	Yes	Yes
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Yes	Yes

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
FILE RECOVERY	Search for files or folders from an archive	Supported (CA,TA)	N/A	N/A
	Browse for files or folders from an archive snapshot	Supported (CA,TA)	N/A	N/A
	Download files and folders from an archive	Supported (TA)	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	No	N/A	N/A
	Download n VMX file from an archive	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported for volume based job(CA)	No	No
	Download metadata for all runs of a job	Supported for volume based job(CA)	No	No
	Download a snapshot (archived data) for a job run	Supported for volume based job(CA)	No	No
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported for volume based job	No	No
	Clone all objects of a job run from an archive	No	No	No
	clone DB to FCI	N/A	Yes	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported for volume based job	No	No
	Recover all objects of job run from an archive to a source	No	No	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical, volume based)(CA,TA)	No	No
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported for volume based job	N/A	N/A
	Search files or folders from the metadata of a retrieved job run	Supported for volume based job(CA,TA)	N/A	N/A
	Browse files or folders from a retrieved snapshot	Supported for volume based job(CA,TA)	N/A	N/A
	Browse files or folders from the metadata of a downloaded job run	Supported for volume based job(CA,TA)	N/A	N/A
	Download files and folders from an archive	Supported for volume based job(TA)	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	Supported for volume based job(CA,TA)	N/A	N/A
	Download a VMX file from an archive	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Replication Based Workflow for MS SQL Volume-base, File-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Yes as log-based incremental	Supported (CA,TA)	N/A
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	Supported (CA,TA)
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
RECOVER	Recover a single object of a job run from a remote cluster	Yes only for DBs & not full server for physical	Yes	Supported (CA,TA)
	Recover an entire backup job run from a remote cluster	No	No	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount of replicated data	Supported (CA,TA)	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	No	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Only DB is supported (CA,TA)	No
	Clone all objects of a job run from a remote cluster	No	No	No
	clone db to FCI	N/A	Supported (CA,TA)	N/A
	Crash-consistent clone for apps from a remote cluster	No	No	No

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Only DB is supported (CA,TA)	No
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)	N/A	N/A
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	N/A	N/A
	Browse for replicated object on a remote cluster	Supported (CA,TA)	N/A	N/A
	Download replicated files and folders from a remote cluster	Supported (TA)	N/A	N/A
	Recover replicated files from a remote cluster	Supported (CA,TA)	N/A	N/A
	Download a VMX file from a remote cluster	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A

Protection Based Workflow for Oracle

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	Supported (CA,TA)	Supported (CA,TA)
	Auto-protect new objects	No	No

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)
	Incremental backup with SQL server differential	N/A	N/A
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	No	No
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)
	CloudSpin to AWS and Azure	N/A	N/A

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run	No	No
	Recover a SQL DB to the original server and instance	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to the original server but a different instance	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	Supported (CA,TA)	Supported (CA,TA)
	Recover DB to Alternate Instance with overwite	Supported (CA,TA)	N/A
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	No	No
	Instant volume mount	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A
CLONE	Clone an object of a job run	Supported (CA,TA)	Supported (CA,TA)
	Clone all the objects of a job run	No	No
	Clone db to FCI instance	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
TEARDOWN	Tear down objects	Supported (CA,TA)	Supported (CA,TA)
	Tear down instant volume mount	N/A	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders	N/A	N/A
	Browse for a backed-up object	N/A	N/A
	Download files/folders	N/A	N/A
	Recover files	N/A	N/A
	Download a VMX file	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A

Archival Based Workflow for Oracle

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A
	Archive a full backup	Supported (CA,TA)	Supported (CA,TA)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No
	Instant volume mount from an archive	N/A	N/A
	Crash-consistent recover for apps from an archive	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A
CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from an archive	Not supported	Not supported
	clone db to FCI	N/A	N/A
	Crash-consistent clone for apps from an archive	N/A	N/A

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	Not supported	Not supported
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders from an archive	N/A	N/A
	Browse for files or folders from an archive snapshot	N/A	N/A
	Download files and folders from an archive	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A
	Download n VMX file from an archive	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	No	Not supported
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from an archive	No	Not supported
	clone DB to FCI	N/A	N/A

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)
	Recover all objects of job run from an archive to a source	No	No
	Instant volume mount from an archive	N/A	N/A
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	N/A	N/A
	Search files or folders from the metadata of a retrieved job run	N/A	N/A
	Browse files or folders from a retrieved snapshot	N/A	N/A
	Browse files or folders from the metadata of a downloaded job run	N/A	N/A
	Download files and folders from an archive	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A
	Download a VMX file from an archive	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A

Replication Based Workflow Oracle

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run from a remote cluster	Not supported	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	No	No
	Instant volume mount of replicated data	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from a remote cluster	No	No
	clone db to FCI	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	N/A	N/A
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)
	Tear down instant volume mount on a remote cluster	N/A	N/A
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	No	No
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders on a remote cluster	N/A	N/A
	Browse for replicated object on a remote cluster	N/A	N/A
	Download replicated files and folders from a remote cluster	N/A	N/A
	Recover replicated files from a remote cluster	N/A	N/A
	Download a VMX file from a remote cluster	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A

Categories	Workflows	Sources	
		Oracle Standalone	Oracle RAC
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A

Supported Cloud Spin and Other Workflows

For more information, see Supported Cloud Spin and Other Workflows in the online help.

Supported Platforms for Heterogeneous Cohesity cluster

Cohesity supports heterogeneous Cohesity cluster that spans nodes of different Cohesity hardware platforms and Cohesity-qualified partner hardware platforms.

Within a heterogeneous Cohesity cluster, you can have a combination of:

- Multiple Cohesity hardware platforms.
- Cohesity hardware platforms and hardware platforms from any one partner.
- Hardware platforms from a partner.

You can create a new Cohesity cluster with heterogeneous nodes or add heterogeneous nodes to an existing Cohesity cluster and combine all the nodes' storage capacities and performance in a Cohesity cluster.

Hardware Compatibility Within a Cohesity cluster

> **Important:** Cohesity recommends using the Sizer tool for guidance on mixing nodes within a cluster.

The heterogeneous Cohesity cluster support has evolved since Cohesity version 6.1. From 6.1 release, more platforms can be combined to form a heterogeneous Cohesity cluster.

Note:

- Creating a new heterogeneous cluster, or expanding an existing cluster to make it Heterogeneous must only be done after consulting your Cohesity accounts team.
- Creating incorrect combinations of heterogeneous clusters without prior consultation with your Cohesity accounts team risks undesirable consequences. Reverting this can be complicated and time-consuming.
- Use the Cohesity Sizer tool for checking the heterogenous cluster options for both Cohesity C-Series and OEM Hardware platforms.

Before you create a heterogeneous Cohesity cluster, review and follow the guidelines below:

- Cohesity does not recommend mixing All-Flash (nodes with SSDs) and hybrid nodes (nodes with HDDs) in the same Cluster.
- Hardware platforms from only one partner can be combined with Cohesity hardware platforms to form a heterogeneous Cohesity cluster. For example, you cannot form a heterogeneous Cohesity cluster with Cohesity hardware, Cisco, and HPE hardware platforms.
- Performance sizing should be done at 104% of the performance-gated nodes suggested by the sizer, using the highest-performance node.
- Mixing Hybrid Platforms
    - Equivalent Node Type - Two hybrid node types are considered equivalent if:
        - Their CPU/SSD ratio is within 10% of each other.
        - They have the same number of hard disks (HDDs).
        - The total HDD capacity of the nodes should be within 25% of each other.
    - Mixing Equivalent Nodes - In a heterogeneous cluster, if two platforms are equivalent, the combined number of total nodes in the cluster should be greater than or equal to the recommended number of nodes per the EC configuration.
    - Mixing Non-Equivalent Nodes - In a heterogeneous cluster, if two platforms are non-equivalent, then the minimum number of nodes for each equivalence-node type must be as per the EC configuration recommendation.

      For more details on the recommended number of nodes, see the Disk and Node Resiliency in Cohesity Fault Tolerance—Data Integrity for Modern Web-scale Environments.

Hardware Compatibility Matrix

The following table lists the Cohesity and partner hardware platforms compatible with Cohesity-Series and all other OEM platforms to form a heterogeneous cluster:

> **Note:**
> - New nodes that are added to the cluster should have the same or more CPU cores and the same or more SSD capacity (total SSD capacity) than the existing nodes in the cluster.
> - Hardware platforms from only one partner can be combined with Cohesity hardware platforms to form a heterogeneous cluster. However, you can combine any models from the same partner hardware.

Vendor	Cohesity C-Series	Cisco	Dell	HPE	Lenovo	Intel	Fujitsu	Super Micro
Cohesity C-Series	✓	✓	✓	✓	✓	✓	✓	✓
Cisco	✓	✓	✗	✗	✗	✗	✗	✗
Dell	✓	✗	✓	✗	✗	✗	✗	✗
HPE	✓	✗	✗	✓	✗	✗	✗	✗
Lenovo	✓	✗	✗	✗	✓	✗	✗	✗
Intel	✓	✗	✗	✗	✗	✓	✗	✗
Fujitsu	✓	✗	✗	✗	✗	✗	✓	✗
SuperMicro	✓	✗	✗	✗	✗	✗	✗	✓

> **Note:** Pure FlashRecover is not supported in a heterogeneous cluster configuration.

> **Important:**
> Refer the Hardware Universe Help for the following information:
>
> - For a complete list of all the supported C-Series and OEM platforms.
> - For a complete list of all the End-Of-Sale C-Series and OEM platforms, see the End-Of-Sale Platforms section.
> - Minimum supported Cohesity versions for C-Series and OEM platforms.

You can combine compute or general-purpose nodes with Cohesity hardware platforms or any models from the same partner hardware to form a heterogeneous Cohesity cluster. For more information on compute nodes, see Manage Compute Nodes.

Node Coexistence Compatibility Within a Chassis

The table below shows which node hardware models can coexist within one chassis. For example, a C2510 node can be physically added to a chassis with C2500 nodes. However, it is not supported for a C2515 node to be physically added to a chassis with C2500 nodes.

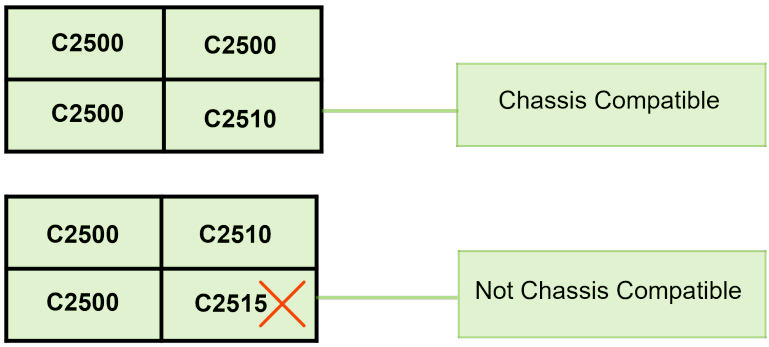

The green check ✓ in the table cell below indicates compatibility. Hardware combinations with no checks are not compatible.

> **Note:** For C4000, C5000, C6000, and CX8000 the node hardware models cannot coexist within the same chassis.

Hardware Model	C2100	C2300	C2500	C2510	C2105	C2305	C2505	C2515	C2605
C2100	✓								

Hardware Model	C2100	C2300	C2500	C2510	C2105	C2305	C2505	C2515	C2605
C2300		✓							
C2500			✓	✓					
C2510			✓	✓					
C2105					✓				
C2305						✓			
C2505							✓	✓	✓
C2515							✓	✓	✓
C2605							✓	✓	✓

> **Note:** For C4000, C5000, CX8000 and CNG2, the node hardware models cannot coexist within the same chassis. This limitation does not apply to the C6000, since C6000 is a single-node chassis.

Agent Compatibility Within a Cohesity cluster

This page details the Linux and Windows agent compatibility within a Cohesity cluster.

You must install the Cohesity Agent on Physical servers, SQL servers, Exchange servers, and other Hypervisor environments you want to protect.

The Cohesity versions and Cohesity Agent versions specified in the following matrix include their respective minor release versions. For example, Cohesity 6.8.1 version support is also applicable for 6.8.1_u1, 6.8.1_u2, and 6.8.1_ux versions, where x denotes the minor/update release version number.

Similarly, Cohesity Agent 6.8.1 version support is also applicable for 6.8.1_u1, 6.8.1_u2, and 6.8.1_ux versions.

> **Note:**
> Legends:
> ✅ - Supported
> ❌ - Not Supported

Cohesity Linux Agent Compatibility Matrix

Cohesity Agent Versions	Cohesity Versions									
	7.4.1	7.4	7.3.2	7.3.1	7.3	7.2.2	7.2.1	7.2	7.1.2	6.8.2
7.4.1	✅	❌	❌	❌	❌	❌	❌	❌	❌	❌
7.4	✅	✅	❌	❌	❌	❌	❌	❌	❌	❌
7.3.2	✅	✅	✅	❌	❌	❌	❌	❌	❌	❌
7.3.1	✅	✅	✅	✅	❌	❌	❌	❌	❌	❌
7.3	✅	✅	✅	✅	✅	❌	❌	❌	❌	❌
7.2.2	✅	✅	✅	✅	✅	✅	❌	❌	❌	❌
7.2.1	✅	✅	✅	✅	✅	✅	✅	❌	❌	❌
7.2	✅	✅	✅	✅	✅	✅	✅	✅	❌	❌
7.1.2	✅	✅	✅	✅	✅	✅	✅	✅	✅	❌
6.8.2	✅	✅	✅	✅	✅	✅	✅	✅	✅	✅

> **Note:** Cohesity recommends upgrading the Agent version to the same version as the Cohesity cluster version. For more information, see Install and Manage the Agent on Linux Servers.

Cohesity Windows Agent Compatibility Matrix

Cohesity Agent Versions	Cohesity Versions									
	7.4.1	7.4	7.3.2	7.3.1	7.3	7.2.2	7.2.1	7.2	7.1.2	6.8.2
7.4.1	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗
7.4	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
7.3.2	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
7.3.1	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
7.3	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗
7.2.2	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
7.2.1	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗
7.2	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
7.1.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
6.8.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

> **Note:** Cohesity recommends upgrading the Agent version to the same version as the Cohesity cluster version. For more information, see Install and Manage the Agent on Windows Servers.

Cohesity AIX Agent Compatibility Matrix

Cohesity Agent Versions	Cohesity Versions									
	7.4.1	7.4	7.3.2	7.3.1	7.3	7.2.2	7.2.1	7.2	7.1.2	6.8.2
7.4.1	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗
7.4	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
7.3.2	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
7.3.1	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
7.3	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗
7.2.2	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
7.2.1	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗
7.2	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
7.1.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
6.8.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

> **Note:** Cohesity recommends upgrading the Agent version to the same version as the Cohesity cluster version. For more information, see Install and Manage the Cohesity Java Agent on AIX Servers.

Cohesity Solaris Agent Compatibility Matrix

Cohesity Agent Versions	Cohesity Versions									
	7.4.1	7.4	7.3.2	7.3.1	7.3	7.2.2	7.2.1	7.2	7.1.2	6.8.2
7.4.1	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗
7.4	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
7.3.2	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
7.3.1	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
7.3	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗
7.2.2	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
7.2.1	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗
7.2	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
7.1.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
6.8.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

> **Note:** For more information, see Install and Manage the Agent on Solaris Servers.

Cohesity HP-UX Agent Compatibility Matrix

Cohesity Agent Versions	Cohesity Versions									
	7.4.1	7.4	7.3.2	7.3.1	7.3	7.2.2	7.2.1	7.2	7.1.2	6.8.2
7.4.1	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗

Cohesity Agent Versions	Cohesity Versions									
	7.4.1	7.4	7.3.2	7.3.1	7.3	7.2.2	7.2.1	7.2	7.1.2	6.8.2
7.4	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
7.3.2	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
7.3.1	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
7.3	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗
7.2.2	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
7.2.1	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗
7.2	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
7.1.2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
6.8.2	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗

Cohesity Cloud Edition Supportability

This section covers the supportability information of the Cohesity Cloud Edition:

- Cohesity Cloud Edition on AWS Supportability
- Cohesity Cloud Edition on Azure Supportability
- Cohesity Cloud Edition on GCP Supportability
- Cohesity Cloud Edition on OCI Supportability
- Cohesity Cloud Edition on OCI Supportability

Cohesity Cloud Edition on AWS Supportability

In Cohesity Cloud Edition on AWS, backed-up data is down-tiered to object storage configured with the Standard, Cold, or Archive storage tier.

Supported Sources and Workloads for Standard Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on AWS when the object storage is configured with the Standard tier.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Relational Database Service for MySQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both snapshot-based and ingest-based protection. Cross-cloud restore is not supported.
Amazon Aurora MySQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Cluster-level (With snapshot-based protection.)	Supports both snapshot-based and ingest-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service for PostgreSQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (with snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Relational Database Service for PostgreSQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Cluster-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service for Microsoft SQL Server	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service for Oracle	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Instance-level	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Relational Database Service Custom for Oracle			Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service Custom for SQL Server	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon RDS and Amazon Aurora	Yes	Yes	No	No	Instance-level	Supports only snapshot-based protection
Amazon Redshift Cluster	Yes	Yes	Yes	Yes	Database-level	Supports only ingest-based protection.
Amazon DocumentDB	Yes	Yes	Yes	Yes	Database-level	Supports only ingest-based protection.
Amazon DynamoDB	Yes	Yes	Yes	Yes	Table-level	Supports only ingest-based protection.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Elastic Compute Cloud Instances	Yes	Yes	Yes	Yes	File-level	File-level recovery is not supported for EC2 instances with XFSv5 file system.
Amazon Simple Storage Service Bucket	Yes	Yes	Yes	Yes	File-level	Granular recovery of S3 objects with prefixes is supported.
VMware Cloud (VMC)	Yes	Yes	Yes	No	File-level	Multi-tenancy should be enabled and hybrid extender needs to be configured.
Amazon Elastic VMware Service	Yes	Yes	Yes	No	File-level	
Nutanix AHV Prism Central	Yes	Yes	Yes	Yes	File-level	
Acropolis: Acropolis Standalone Cluster	Yes	Yes	Yes	Yes	File-level	
Physical Servers	Yes	Yes	Yes	Yes	File-level	Supports both file-based and block-based backups.
Microsoft SQL Server	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based and VDI-based backups.
Microsoft 365	Yes	Yes	Yes	Yes	File-level	

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
NAS	Yes	Yes	Yes	Yes	File-level	• Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions. • NetApp: Supports only File-Runner-based backup. SnapDiff-based backup is not supported. • Isilon: Supports both File-runner and CFT-based backups.
SAP HANA on Linux (x86)	Yes	Yes	Yes	Yes	Not Applicable	
MongoDB	Yes	Yes	Yes	Yes	Not Applicable	
Cassandra	Yes	Yes	Yes	Yes	Not Applicable	

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Hadoop	Yes	Yes	Yes	Yes	File-level (Only Hadoop HDFS)	Supports protection of HDFS, Hive, and HBase objects.
Oracle Database	Yes	Yes	Yes	Yes	Not Applicable	Supports both Oracle Adapter and Oracle SBT Plug-in.
IBM DB2	Yes	Yes	Yes	Yes	Not Applicable	
PostgreSQL	Yes	Yes	Yes	Yes	Not Applicable	Supports instant restore.
SAP on Oracle for Linux	Yes	Yes	Yes	Yes	Not Applicable	
SAP on Oracle for Microsoft Windows	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Linux	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Windows	Yes	Yes	Yes	Yes	Not Applicable	
MySQL EE	Yes	Yes	No	No	Database-level	

*Cohesity Cloud Edition on AWS supports only Cloud Archive with **Incremental Forever** archival format.

- Cohesity Cloud Edition on AWS supports Write Once Read Many (WORM) for Cloud Archive with incremental forever archival format. The WORM model ensures that the objects written to the WORM-supported targets cannot be tampered with (deleted or overwritten). The archives created are immutable for the duration configured. For the list of external targets with WORM support,

For the list of supported external targets for Cloud Archive, see Supported Workflows and External Targets.

Supported Sources and Workloads for Archive and Cold Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on AWS when the object storage is configured with the Archive and Cold Tier tiers.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Elastic Compute Cloud Instances	Yes	Yes	No	No	File-level	• File-level recovery is not supported for EC2 instances with XFSv5 file system. • VM indexing disabled.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
NAS	Yes	Yes	No	No	File-level	• Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions. • NetApp: Supports only File-Runner-based backup. SnapDiff-based backup is not supported. • Isilon: Supports both File-runner and CFT-based backups.
VMware VM	Yes	Yes	No	No	File-level	VM indexing disabled.

Supported Features and Services

Features and Services	Standard Tier	Cold Tier	Archive Tier	Notes
SmartFiles Support	Yes	Yes	No	Cohesity Cloud Edition on AWS configured with Standard Tier and Cold Tier support SmartFiles for SMB, NFS, and S3 protocols, with functionality limited to View Snapshots and Replication.

Features and Services	Standard Tier	Cold Tier	Archive Tier	Notes
Marketplace Support	Yes	No	No	You can install Marketplace Apps on Cohesity Cloud Edition on AWS configured with Standard Tier. For more information on Marketplace App, see Apps
FortKnox Support	Yes	No	No	Cohesity FortKnox supports vaulting data from Cohesity Cloud Edition on AWS configured with Standard Tier to a cloud vault. For more information, see Clusters.
Multitenancy Support	Yes	No	No	Cohesity Cloud Edition on AWS configured with Standard Tier supports the multitenancy feature from 7.2.2_u1 onwards. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. For more information, see Organizations (Multitenancy) .

Cohesity Cloud Edition on Azure Supportability

In Cohesity Cloud Edition on Azure, backed-up data is down-tiered to Azure storage configured with the Hot, Cold, or Archive storage tier.

Supported Sources and Workloads for Hot Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on Azure when the object storage is configured with the Hot tier.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
Azure Blob Storage and Azure Data Lake Storage	Yes	Yes	Yes	Yes	Yes	Yes	File-level	
Azure SQL Managed Instance	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure SQL Database	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure Table Storage	Yes	No	Yes	Yes	Yes	Yes	Table-level	
Azure Cosmos DB - Table	Yes	No	Yes	Yes	Yes	Yes	Table-level	
Azure Cosmos DB - Cassandra	Yes	No	Yes	Yes	Yes	Yes	Keyspace -level	
Azure Cosmos DB - NoSQL	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure Cosmos DB - MongoDB	Yes	No	Yes	Yes	Yes	Yes	Database-level	

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
Azure PostgreSQL	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure MySQL	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure VMware Solution (AVS) Gen1 and Gen2	Yes	Yes	Yes	Yes	Yes	Yes	File-level	
Azure Kubernetes Service (AKS)	Yes	No	Yes	Yes	Yes	Yes	File-level	
Azure	Yes	No	Yes	Yes	Yes	Yes	File-level	
Nutanix AHV Prism Central	Yes	No	Yes	Yes	Yes	Yes	File-level	
Acropolis: Acropolis Standalone Cluster	Yes	No	Yes	Yes	Yes	Yes	File-level	
Nutanix Cloud Clusters on Azure	Yes	No	Yes	Yes	Yes	Yes	File-level	

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
Physical Servers	Yes	No	Yes	Yes	Yes	Yes	File-level	
Microsoft SQL Server (Standalone and FCI)	Yes	Yes	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based and VDI-based backups.
Microsoft 365	Yes	Yes	Yes	Yes	Yes	Yes	File-level	
NAS	Yes		Yes	Yes	Yes	Yes	File-level	Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions.
SAP HANA	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
MongoDB	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
Cassandra	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
IBM DB2	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP on Oracle for Linux	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP on Oracle for Microsoft Windows	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
Hadoop	Yes	No	Yes	Yes	Yes	Yes	Yes (Only Hadoop HDFS)	Supports protection of HDFS, Hive, and HBase objects.
PostgreSQL	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supports instant restore.
SAP Sybase ASE on Microsoft Windows	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supported from 7.2.2_u1 version onwards.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
SAP Sybase ASE on Linux	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supported from 7.2.2_u1 version onwards.
Oracle Database	Yes	Yes	Yes	Yes	Yes	Yes	Not Applicable	Supports both Oracle Adapter and Oracle SBT Plug-in.
SAP MaxDB on Linux	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Windows	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP Sybase IQ	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supported from 7.2.2_u1 version onwards.
MySQL EE	Yes	No	Yes	Yes	No	No	Database-level	

* Cohesity Cloud Edition on Azure supports both **Cloud Archive with periodic full** and **Cloud Archive with Incremental Forever** archival formats.

For the list of supported external targets for Cloud Archive, see Supported Workflows and External Targets.

Supported Sources and Workloads for Archive and Cold Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on Azure when the object storage is configured with the Archive and Cold Tier tiers.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Azure	Yes	No	Yes	Yes	Yes	Yes	File-level	VM indexing disabled.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
NAS	Yes	No	Yes	Yes	No	No	File-level	• Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions. • NetApp: Supports only File-Runner-based backup. SnapDiff-based backup is not supported. • Isilon: Supports both File-

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
								runner and CFT-based backups.
VMware VM	Yes	Yes	Yes	Yes	No	No	File-level	VM indexing disabled.

Supported Features and Services

Features and Services	Hot Tier	Cold Tier	Archive Tier	Notes
SmartFiles Support	Yes	Yes	No	Cohesity Cloud Edition on Azure configured with Hot Tier and Cold Tier support SmartFiles for SMB, NFS, and S3 protocols, with functionality limited to View Snapshots and Replication.
Marketplace Support	Yes	No	No	You can install Marketplace Apps on Cohesity Cloud Edition on Azure configured with Hot Tier. For more information on Marketplace App, see Apps
FortKnox Support	Yes	No	No	Cohesity FortKnox supports vaulting data from Cohesity Cloud Edition on Azure configured with Hot Tier to a cloud vault. For more information, see Clusters.
Multitenancy Support	Yes	No	No	Cohesity Cloud Edition on Azure configured with Hot Tier supports the multitenancy feature from 7.2.2_u1 onwards. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. For more information, see Organizations (Multitenancy) .

Cohesity Cloud Edition on GCP Supportability

Supported Sources and Workloads

The following table provides the mapping of the workflows that are supported on the sources Cohesity Cloud Edition on GCP supports, as per the latest release:

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Bigtable	Yes	Yes	Yes	Yes	Database-level	
Firestore	Yes	Yes	Yes	Yes	Database-level	
BigQuery	Yes	Yes	Yes	Yes	Database-level	
Spanner	Yes	Yes	Yes	Yes	Database-level	
AlloyDB for PostgreSQL	Yes	Yes	Yes	Yes	Database-level	
Cloud SQL for SQL Server	Yes	Yes	Yes	Yes	Database-level	
Cloud SQL for MySQL	Yes	Yes	Yes	Yes	Database-level	
Cloud SQL for PostgreSQL	Yes	Yes	Yes	Yes	Database-level	
GCP VM	Yes	Yes	Yes	Yes	File-level	File-level recovery is not supported for GCP VMs with XFSv5 file system.
Cassandra	Yes	Yes	Yes	Yes	Not Applicable	

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Hadoop	Yes	Yes	Yes	Yes	File-level (Only Hadoop HDFS)	Supports protection of HDFS, Hive, and HBase objects.
Microsoft SQL Server	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based volume-based, and VDI-based backups.
MongoDB	Yes	Yes	Yes	Yes	Not Applicable	

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Google Cloud VMware Engine (GCVE)	Yes	Yes	Yes	No	File-level	On the GCVE, the default user (CloudOwner@gve.local) is missing the permission to access the Host -> Configuration -> Storage Partition Configuration. As a result, instant recovery fails, with the following error message: `Could not mount datastore on any host. Additional details: SOAP 1.1 fault 'ServerFaultCode [no subcode] - Permission to perform this operation was denied.` As a workaround, ensure that the default user is granted the Host -> Configuration -> Storage Partition Configuration permission.
Physical Servers	Yes	Yes	Yes	Yes	File-level	
Cohesity Views	Yes	Yes	Yes	Yes	File-level	
NAS(Generic NAS, NetApp, Isilon, Pure Storage FlashBlade)	Yes	Yes	Yes	Yes	File-level	Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Red Hat OpenShift Virtualization	Yes	Yes	Yes	Yes	File-level	Supports only 4.14 version.
PostgreSQL	Yes	Yes	Yes	Yes	Not Applicable	Supports instant restore.
Oracle Database	Yes	Yes	Yes	Yes	Not Applicable	Supports both Oracle Adapter and Oracle SBT Plug-in.
SAP MaxDB on Linux	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Windows	Yes	Yes	Yes	Yes	Not Applicable	
SAP HANA	Yes	Yes	Yes	Yes	Not Applicable	
MySQL EE	Yes	Yes	No	No	Database-level	

*Cohesity Cloud Edition on GCP supports only Cloud Archive with **Incremental Forever** archival format.

For the list of supported external targets for Cloud Archive, see Supported Workflows and External Targets.

SmartFiles Support

Cohesity Cloud Edition on GCP supports SmartFiles for SMB, NFS, and S3 protocols, with functionality limited to View Snapshots and Replication.

> **Note:** NGCE cluster maintains a local cache of snapshots for a maximum of 7 days to optimize recovery. As a result, only snapshot versions from the last 7 days are available.

Marketplace Support

You can install Marketplace Apps on Cohesity Cloud Edition on GCP. For more information on Marketplace App, see Apps.

Multitenancy

Cohesity Cloud Edition on GCP supports the multitenancy feature. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. For more information, see Organizations (Multitenancy) .

Cohesity Cloud Edition on OCI Supportability

Supported Sources and Workloads

The following table provides the mapping of the workflows that are supported on the sources supports, as per the latest release:

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Physical Servers	Yes	Yes	Yes	Yes	File-level	
NAS (Generic NAS, NetApp, Isilon, Pure Storage FlashBlade)	Yes	Yes	Yes	Yes	File-level	Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions.
Microsoft 365	Yes	Yes	Yes	Yes	File-level	
Microsoft SQL Server	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based and VDI-based backups.
Oracle Database	Yes	Yes	Yes	Yes	Not Applicable	Supports Oracle Adapter-based protection.

Replication Compatibility

Replication automatically makes copies of snapshots captured by a Protection Group on one Cohesity cluster and copies them to a second Cohesity cluster.

The following table indicates the replication compatibility between Cohesity clusters running different versions.

	To 6.8.2	To 7.1.2	To 7.2	To 7.2.1	To 7.2.2	To 7.3	To 7.3.1	To 7.3.2	To 7.4	To 7.4.1
Replicating from 6.8.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.1.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.2.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.2.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.3	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.3.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.3.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.4	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Replicating from 7.4.1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note:

- The table above reflects replication compatibility only. It does not indicate restore compatibility between versions.
- If the source cluster adds support for a new workload or a new category within an existing workload and attempts replication, the target cluster must also support that workload or category. If the target cluster is running an older version that does not recognize this capability, replication of such backups fails.

Restore Limitations in Cross-Version Replication Scenarios

Replication between Cohesity clusters running different software versions is supported, as shown in the compatibility matrix. However, restore operations can fail when a restore is initiated on a cluster running an older software version using backups replicated from a newer-version cluster.

Restore failures occur because newer Cohesity releases can capture additional metadata for protected sources. Although this metadata replicates successfully to older clusters, the older software may not be able to interpret it during restore operations. This behavior varies by workload type, adapter, and Cohesity release.

Once the target cluster is upgraded to the same or a newer software version than the cluster that created the backup, the replicated backups become fully usable for restore operations.

To ensure reliable recovery, Cohesity recommends performing restore operations on a cluster running the same or a newer software version than the version that created the backup.

If restore operations must be performed while the cluster is running an older software version, Cohesity Support must review the scenario on a case-by-case basis.

For guidance specific to your environment, contact Cohesity Support.

Replication 1.0 and 2.0

Replication 1.0 and 2.0 are replication algorithms used by Cohesity to compare and transmit data between the source cluster and the target cluster.

- Replication 1.0 computes the changed files and directories by scanning the entire volume and then transmits the changed metadata to the target cluster.
- Replication 2.0 leverages SnapTree to efficiently determine the changed entries and transmits the changes to the target cluster.

Note:

- Replication 2.0 was introduced in Cohesity version 6.3 and has been the default replication protocol since then. It is also used for S3 view replication, a capability added in versions 6.5.x and 6.6 for object-store views.
- By default, when you upgrade a cluster from a version earlier than 6.3 to version 6.3 or later, a Protection Group begins using Replication 2.0 only after two replication runs complete on the upgraded version. If you want to start using Replication 2.0 immediately after upgrading from a pre-6.3 version, cancel any pending or in-progress replication jobs for backups taken before the upgrade. This ensures that Replication 2.0 is applied immediately, without waiting for two post-upgrade cycles.

Supportability for Qualifications

This topic provides information regarding Cohesity's aims around time frames for qualifying newly released versions of operating systems, applications, and databases. This document is for information purposes only. It does not amend or supplement any contract and is not a warranty, commitment, or obligation of any kind. The information in this document is subject to change without notice, and no third party should rely on it to make any decisions.

Workloads

The workloads covered by this document are categorized into the following tiers:

- Tier 1
- Tier 2

Note: Qualifications of new releases of workloads will only include certification of current Cohesity Generally Available capabilities. New capabilities are part of feature releases. Cohesity will only qualify new third-party vendor releases as far back as the most recent Cohesity LTS release. The third-party vendors determine major and minor release designations.

Cohesity Estimates Based on Vendor Release GA Date

Workload	Major Release	Minor Release
Tier 1	1 Quarter	2 Months
Tier 2	2 Quarters	1 Quarter

Tier 1

Categories	Supported Software
Operating Systems	Windows, Oracle Linux, Red Hat Linux, AIX, SLES/SuSe Linux
Applications	Microsoft SQL Server, Oracle, Cassandra, MySQL, SAP HANA, SAP Oracle, DB2, PostgreSQL, Active Directory
Hypervisors	VMware, Nutanix AHV, Kubernetes (Vanilla and Red Hat Openshift)
NAS	Dell Isilon, NetApp, Pure Storage FlashBlade

Tier 2

Categories	Supported Software
Operating Systems	Debian, Ubuntu, FreeBSD, Solaris, HP-UX
Applications	Couchbase, SAP Sybase, SAP MaxDB, MariaDB, Hadoop, MongoDB, Microsoft Exchange (on-prem), CockroachDB, HDFS
Hypervisors	KVM/RHV/OLVM, Microsoft Hyper-V, Kubernetes (VMware Tanzu)
NAS	Elastifile, Nimble, IBM GPFS

Chapter 3: Data Protection Key Concepts

This chapter includes the following topics:

- Cohesity Cluster Architecture 157
- Protection Overview 161
- About Scale Out 173
- About Cloud Archival 175
- About Cloud Tier 178
- About Tape Archival 179
- About Fault Tolerance 183
- Secure Software Development Practices 193
- Cohesity Security Features 200
- MegaFile 226
- Data Retention for Audits 229
- Dynamic Disk Support 232

Cohesity Cluster Architecture

This section describes the components that make up a Cohesity cluster.

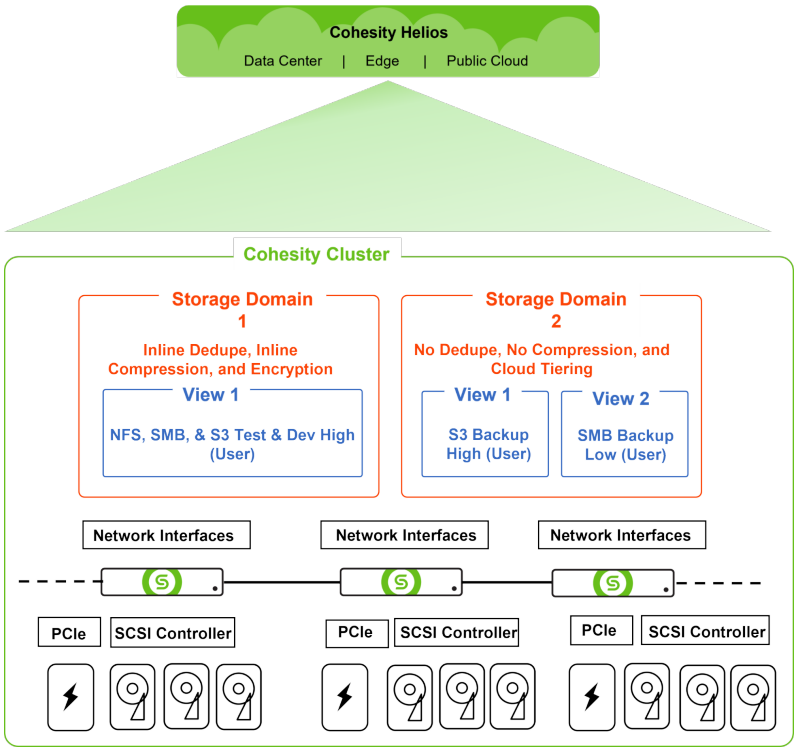

Nodes

A Cohesity node contains CPU with spinning disk or flash storage that is preinstalled with the Cohesity software. A Cohesity cluster can be hosted on:

- Datacenters: Running on x86-based server consisting of three or more physical nodes.
- Edge: A single node ROBO running on x86-based physical servers or a Virtual Edition cluster running on hypervisor with a single node or multiple nodes, depending on the hypervisor.
- Public Cloud: Cohesity Cloud Edition Cohesity cluster deployed on public clouds like Amazon Web Services, Microsoft Azure and Google Cloud Platform.

A node can be in one of the following states:

- Available on the Network but not assigned to the cluster
- Assigned to the cluster

Storage Domains

A Storage Domain, formerly known as a View Box, is a named storage location on a cluster. At the storage domain level, settings like fault tolerance, encryption, and storage efficiencies are applied. A Storage Domain contains Views. When you configure a Protection

Group, you specify a Storage Domain and during a Protection Group Run, the Cohesity cluster stores the Snapshots in that Storage Domain.

If Deduplication is on, the Cohesity cluster eliminates duplicate blocks of repeating data stored on the Cohesity cluster, thus reducing the amount of storage space needed to store data. You may want to disable deduplication and here are some possible reasons:

- The speed of backing up on the Storage Domain is more important than conserving the disk space used to store the Snapshots or files. For example the Storage Domain contains a View and you want fast writes to that View (Datastore).
- You are storing backups on the Cohesity cluster that were created by another backup product that already deduplicates the data.

If Inline Deduplication is on, deduplication occurs when the Cohesity cluster writes the data to the Cohesity View. If Inline Deduplication is off, deduplication occurs after the Cohesity cluster writes the data to the Cohesity View.

Views

A Cohesity view provides a storage location with one or more of these protocols namely NFS, SMB, S3 and SWIFT in a storage domain on the Cohesity cluster.

The NFS mount path can be mounted onto systems using the NFS protocol. The SMB mount path can be mounted onto systems as an SMB share.

You can use the view to store data such as files, backup snapshots, and cloned VMs. Views can only be mounted on any machines specified in the allowlist subnets. A view also acts as a datastore during the VM cloning process and during VM recovery. You can select the SWIFT protocol and a create SWIFT-only view for OpenStack environments.

You can create a view for the following three categories:

- File Shares
- Backup Target
- Object Services

You can select the **Read/Write** protocol and, optionally, select the **Read-only** protocol.

You can create a Cohesity view only by using a template. The template can be a predefined template or a custom template. For more information, see:

- Create Views Using a Predefined Template
- Create Views Using a Custom Template

> **Note:** If you are using Kerberos (AD or MIT), you must create a view using the same storage domain that is mapped to AD/MIT-Kerberos + LDAP.

For more information, see Create Views Using a Predefined Template and Add Subnets to Allowlists.

Views can be used to store data such as files and directories. For example, other backup products can store data on a view that is hosted on the Cohesity cluster. To access a view from your system, you must mount the view.

The table below summarizes the features at the storage domain and view levels, and how they are inherited.

Features	Storage Domain	View
Fault tolerance	Erasure Coding (EC) Replication Factor (RF)	Inherits storage domain fault tolerance settings and cannot be changed for the View.
Storage efficiency	Deduplication (inline, post-process, or disabled) Compression (inline, post-process, or disabled)	Inherits storage domain efficiencies OR Can be overwritten at the View level to be post-process if at the storage domain level is inline, but cannot be disabled.
Encryption	Enabled or disabled	Inherits storage domain's settings and cannot be changed for the View.
Cloud Tier	Inherits cluster settings, unless overwritten at the storage domain level.	Inherits storage domain's settings and cannot be changed for the View.
QoS Policy	N/A	Set for each View.

Protection Overview

Cohesity consolidates data protection, management and web-scale storage onto a unified platform. Easy integration with the VMware ecosystem provides a plug-and-play experience for virtualization admins and significantly reduces the complexity of protecting data. With physical workload support as well, Cohesity is an end-to-end data protection solution that enables businesses to instantly see a full index of their environment and choose the objects they want to protect. These objects can then be protected with easy-to-use protection policies that can be customized based on retention period or application group.

Protection is the process of safeguarding your organization's data. Cohesity uses the following terminology for protection components.

- **Source:** An object in your organization that you want to protect, such as a VM, physical server, Pure Storage volume, MS SQL Server and NAS.
- **Agent:** The Cohesity software that you install on physical servers and certain VM hypervisors to enable Cohesity to protect them.
- **Registration:** The process of registering the source with the Cohesity cluster.
- **Policy:** A collection of reusable settings that define how and when sources are protected, replicated and archived. You select a policy when configuring a Protection Group.
- **Protection Group:** A backup job that runs repeatedly, based on an associated policy, to back up data from a source and store it on the cluster.
- **Remote Adapter Protection Group:** A backup job that runs scripts or executables on a remote Linux machine, typically for sources that are not supported by Cohesity.
- **External Target:** A software environment outside the cluster, such as a cloud or tape management system, which can be used for archiving inactive data for long-term retention on less expensive storage.

For various workloads (sources), Cohesity's incremental-forever strategy relies on change block technology that is either developed by Cohesity or by the vendor Cohesity is integrating with. There are battled-tested technologies and it is highly unlikely that you will ever face data corruption. However, in the case of a mission-critical workload, Cohesity already enables you to mitigate against this in multiple ways with the following existing functionality:

- By performing a backup validation. For example, using Database consistency checker.
- As a part of your regular database restore runs
- As a part of your regular VM or database clones
- You may take a periodic full backup

Cohesity recommends a first full and incremental forever approach to backup all workloads (sources) with the exception for Oracle(SBT based backup) and SQL (VDI-based) where Cohesity recommends periodic full backups.

The common high-level steps to protect sources are:

1. Download and install the agent software on the source. This is required only for physical servers and some hypervisors.
2. Register the source with the cluster.
3. Choose a policy provided by Cohesity or create your own.
4. Create a Protection Group to back up the source.

The Backup Manager Service is responsible for handling all the requests from the Cohesity Data Cloud (Self-Managed) and APIs. A multi-node cluster environment with a large amount of data increases the load on the primary instance of the Backup Manager Service which increases the response time of Cohesity Data Cloud (Self-Managed and API calls. For enabling a read-only replica of the Backup Manager Service, see Backup Manager Service Replica.

About Recovering VMs

Cohesity provides the ability to recover Protected Objects (such as VMs) from a Snapshot created earlier by a Protection Group. You can recover VMs from a Cohesity cluster or a currently registered archive. You can recover VMs to their original location or a new location. Follow these procedures to recover VMs from a Cohesity cluster: Recover VMs to the Original LocationRecover VMs to a New LocationYou can also watch the VMware VM Recovery video to learn how to recover a VMware VM to an original or new location..

The Recover task extracts the VM files (such as the VMDK files) stored in Snapshots and creates new instances of the VMs in their original locations or in a new location depending on the options you choose during recovery:

- Recover to the Original Location—Recover the VM files (such as the VMDK files) to their original datastores and create new instances of the VMs in the original Resource Pool available in the original Source.
- Recover to a New Location—Recover the VM files (such as the VMDK files) to an alternate datastore and create new instances of the VMs in an alternate Resource Pool of a registered Source.
- Recover from a Cloud Archive—Recover the VM files from a registered archive (External Target) and place them in a View on the Cohesity cluster.

During recovering VMs, one of the following methods is used to identify the ESXi hosts that have access to the required Datastore:

- If all the hosts have access to the required Datastore then VMware is responsible to select the ESXi hosts based on load balancing.
- If only some of the hosts have access to the required Datastore, a random host from the available list of hosts that have access to the required Datastore is selected.

For more details about recovering VMs using the Cohesity Dashboard, see Recover VMs.

Recover to the Original Location

While recovering VMs, you can choose to recover the VM files (such as the VMDK files) to their original datastores and create new instances of the VMs in the original Resource Pool of the original Source as described in the figure and procedure below.

The following example shows one VMware VM being recovered, but you can select multiple VMs to recover if the Snapshots of the VMs are stored in the same Storage Domain. The action steps in the following procedure are represented in the figure below with blue numbered time squares: T1 .

The VM Recover task consists of the following actions:

T1) The Cohesity cluster extracts the VM files (such as the VMDK files) stored in Snapshots and places the new VM files in a temporary View. In this example, a new VMDK2 is created from the Snapshot of VM2 made on 03/21/2025 at 02:00 AM.

T2) New instances of the VMs are created in the original Resource Pool by using the new VMDKs from the temporary View. In this example, the new VM2, using VMDK2, is created in the Resource Pool. The View temporarily becomes the datastore for the VMs as represented by the black line between the Resource Pool and View (datastore). The VMs can be used immediately and you do not have to wait until the VMDK files are completely recovered to the original datastore.

T3) The Cohesity cluster migrates the VM files (such as the VMDKs) using Storage vMotion to the original datastore and any changes made to the data on the VM during the recovery process are preserved. The temporary View is deleted when Storage vMotion is completed.

T4) After the VM files are recovered, the VMs use the VM files in the original datastore. In this example, after VMDK2 is fully recovered, VM2 uses VMDK2 located in the original datastore.

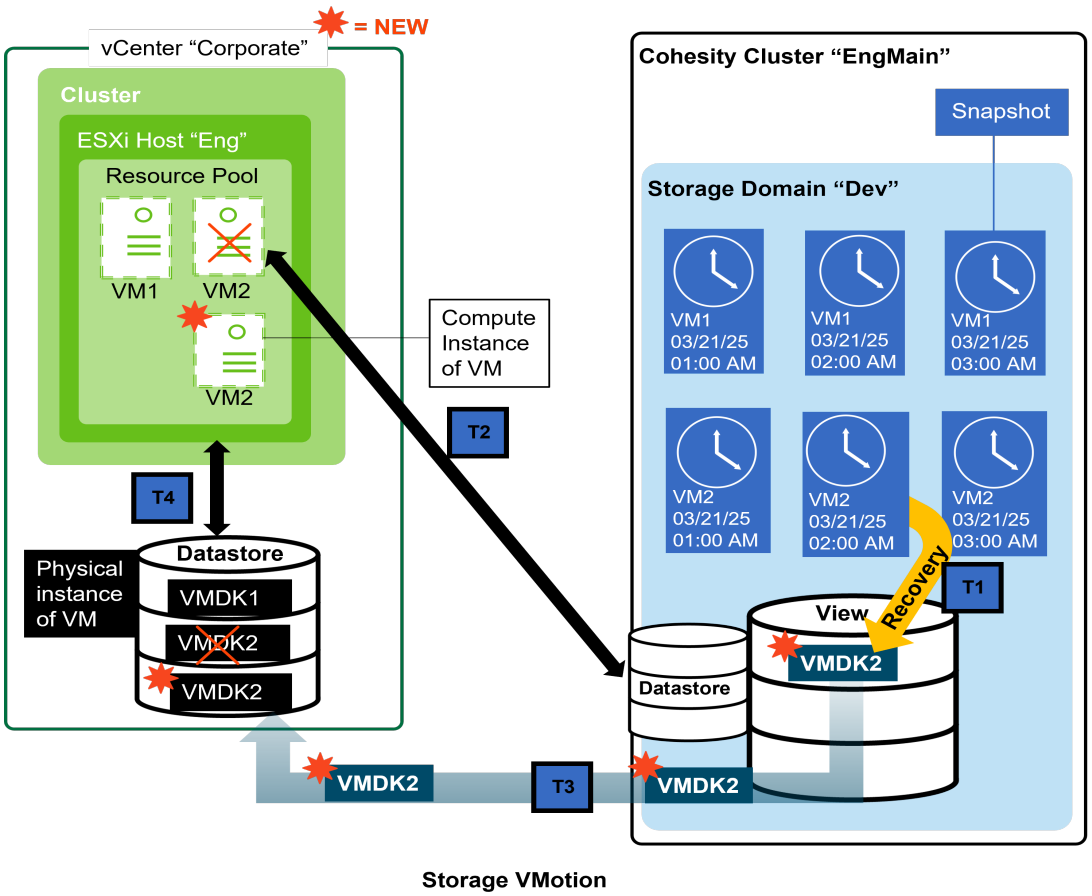

Storage VMotion

Recover to a New Location

While recovering VMs, you can choose to recover the VM files (such as the VMDK files) to a specified datastore and create new instances of the VMs in a specified Resource Pool of a registered Source as described in the figure and procedure below. The VMs do not need to be recovered to their original Resource Pool and the VM files do not need to be recovered to the original datastore.

The following example depicts a VMware VM being recovered. The action steps in the following procedure are represented in the figure below with blue numbered time squares: T1 .

T1) The Cohesity cluster extracts the VM files (such as the VMDK files) stored in Snapshots and places the new VM files in a temporary View. In this example, a new VMDK2 is created from the Snapshot of VM2 made on 03/21/2025 at 02:00 AM and a new VMDK3 is created from the Snapshot of VM3 made on 03/21/2025 at 01:00 AM.

T2) New instances of the VMs are created in the specified Resource Pool by mounting the new VMDKs from the temporary View. In this example, a new VM2 is created by mounting VMDK2 and a new VM3 is created by mounting VMDK3. The View temporarily becomes the datastore for the VM as represented by the black line

between the Resource Pool and View (datastore). The VM can be used immediately and you do not have to wait until the VMDK is completely recovered to the datastore. You can optionally choose to add a suffix and a prefix to the new VM name. For example you can specify the Prefix `Test_` and the Suffix `_QA` and when VM2 is recovered, the new VM that is created is called `Test_VM2_QA`.

T3) The Cohesity cluster migrates the VM files (such as the VMDKs) using Storage vMotion to the specified datastore and any changes made to the data on the VMs during the recovery process are preserved. The temporary View is deleted when Storage vMotion is completed.

T4) After the VM files are recovered, the VMs use the VMDK files stored in the datastore on the Source. In this example, after VMDK2 is fully recovered, `Test_VM2_QA` uses VMDK2 located in the datastore of the vCenter "London" Source. In addition, after VMDK3 is fully recovered, `Test_VM3_QA` uses VMDK3 located in the datastore of the vCenter "London" Source.

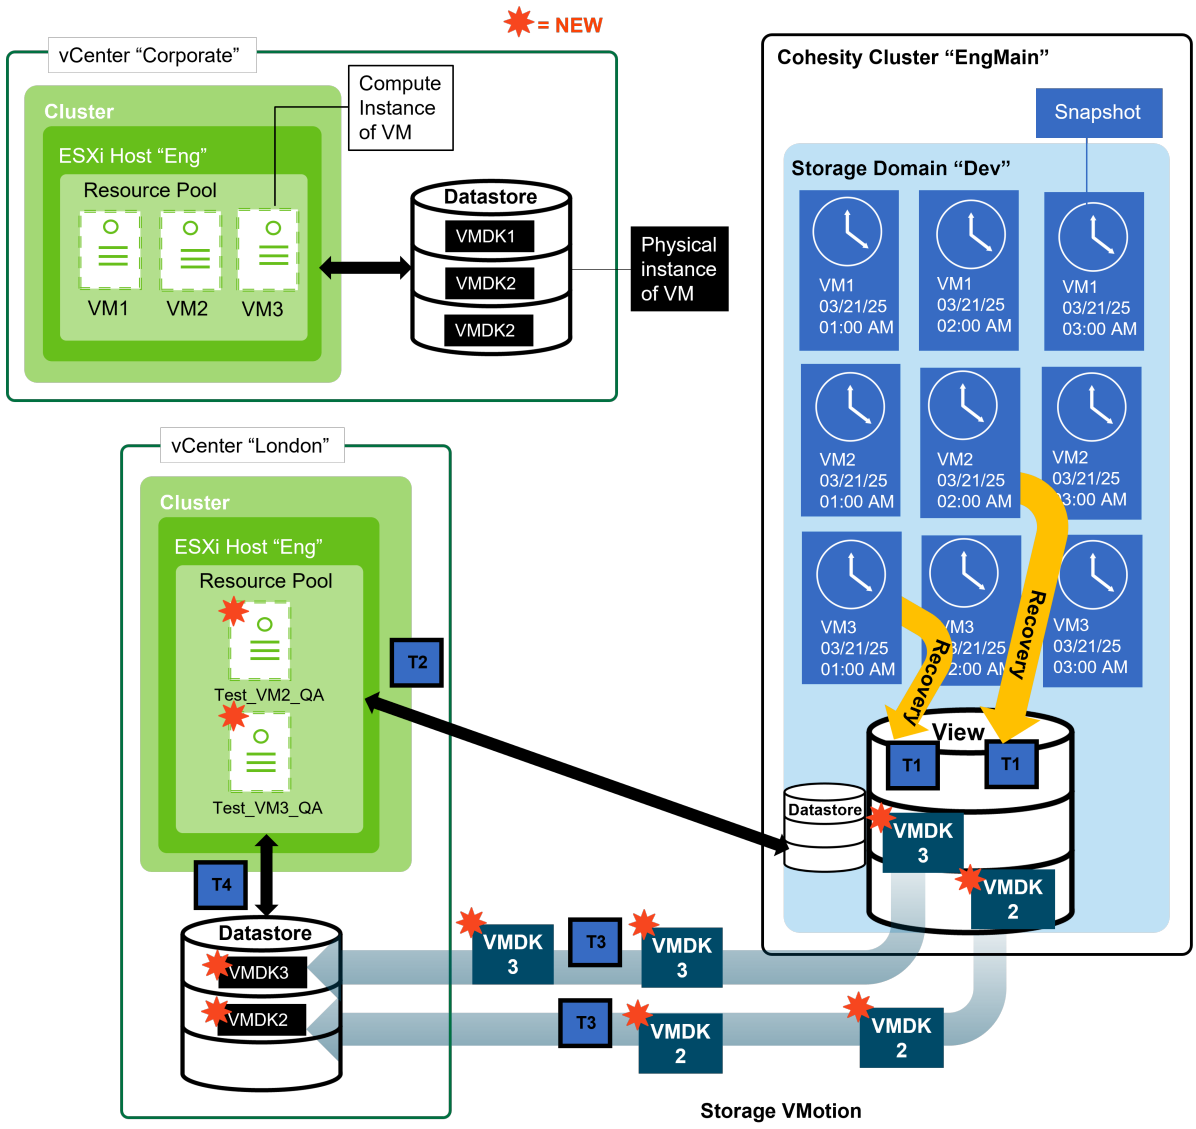

Recover VMs to Multiple Datastores

When you recover VMs to a new location, you can choose multiple target datastores from a specified resource pool of a registered source.

Before you recover multiple VMs to multiple datastores, review and understand these factors:

- You can select multiple datastores only when recovering VMs to a new location.
- The number of target datastores you can select is limited to the number of VMs selected for recovery.
- The VMs are recovered to the selected datastores based on the size of the VMs, free space available in the datastores, successful restore attempts, and other factors.
- One VM will be restored to one datastore.

For more details about recovering VMs using the Cohesity Dashboard, see Recover VMs.

Recover VMs to vCloud Director Using Storage Profiles

You can recover VMs to an alternate vCloud Director location based on the vCloud Director Storage Profiles.

Before you recover VMs to the vCloud Director, review and understand these factors:

- You can select one vCloud Director Storage Profile for a recovery job.
- VMs are recovered to datastores that are compatible with the Storage Profile.
- The datastore to which the VMs are recovered is selected based on the size of the VMs, free space available in the datastores, successful restore attempts, and other factors.
- Instant Recovery is only supported if the target recovery location is explicitly specified while creating a recovery job (recover to a new location).

For more details about recovering VMs using the Cohesity Dashboard, see Recover VMs.

Recover from a Cloud Archive

While recovering VMs, you can choose to recover the VM files from a currently registered archive (External Target). The recovery process is the same as recovering from the Cohesity cluster.

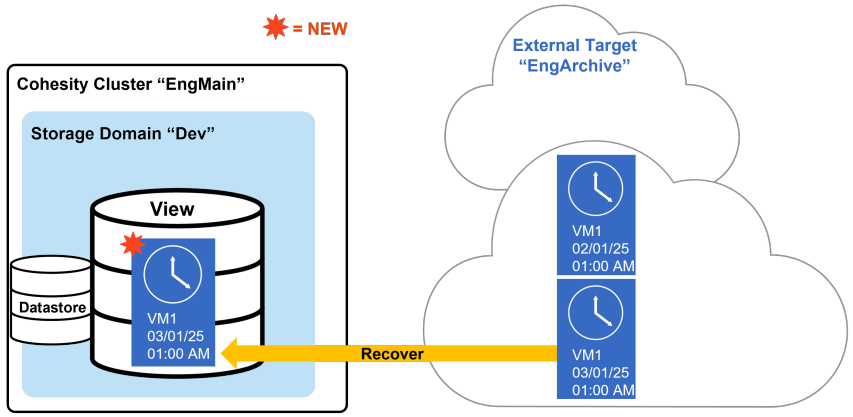

Recover from a Tape Archive

While recovering VMs, you can choose to recover the VM files from a currently registered archive (External Target). The recovery process is the same as recovering from the Cohesity cluster.

When recovering from a tape archive, the Cohesity Dashboard displays the list of tapes required for recovery and their respective details. You can still create a recovery task if some tapes are offline. Recovery can also progress if some tapes are offline. When an offline tape is required for recovery to continue, Cohesity raises an alert and displays a message with the recovery task that needs attention.

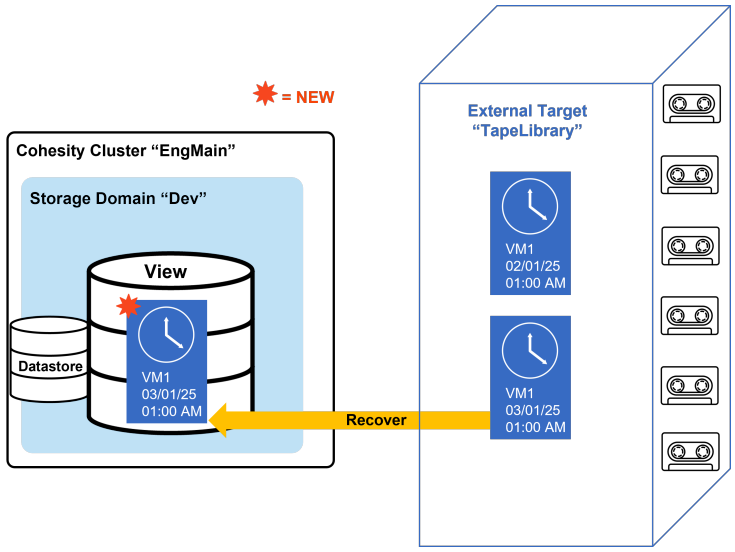

About File Recovery

Cohesity provides the ability to recover files and folders from a Snapshot created earlier by a Protection Group. Files and folders can be recovered to their original location or to a newly specified location, which can be within the original Source or a different one. You can choose to retain the recovered files' and folders' original (at the time of the backup) permissions and attributes. You can also download files and folders from selected Snapshots that were created by a Cohesity Protection Group. However, only items that were indexed when the Snapshot was created can be downloaded (this does not apply to NAS). Recovering and searching for files and folders in Snapshots created by other backup software is not supported.

This Recover task extracts the files stored in Snapshots and creates new instances of them in their original location or a new location depending on the options you choose during recovery. You can also choose to download a file or folder.

- Recover Files or Folders—Recover files or folders to the original location or a new location.
- Download a File or Folder—Download a file or folder from an existing Snapshot.

Refer to the respective source sections for more information about recovery and downloads.

Recover Files or Folders

While recovering files, you can choose to recover to the original location or a new location.

In the figure below, one file is recovered to its original location on a VM, but you can select multiple files to recover if they are from the same Snapshot.

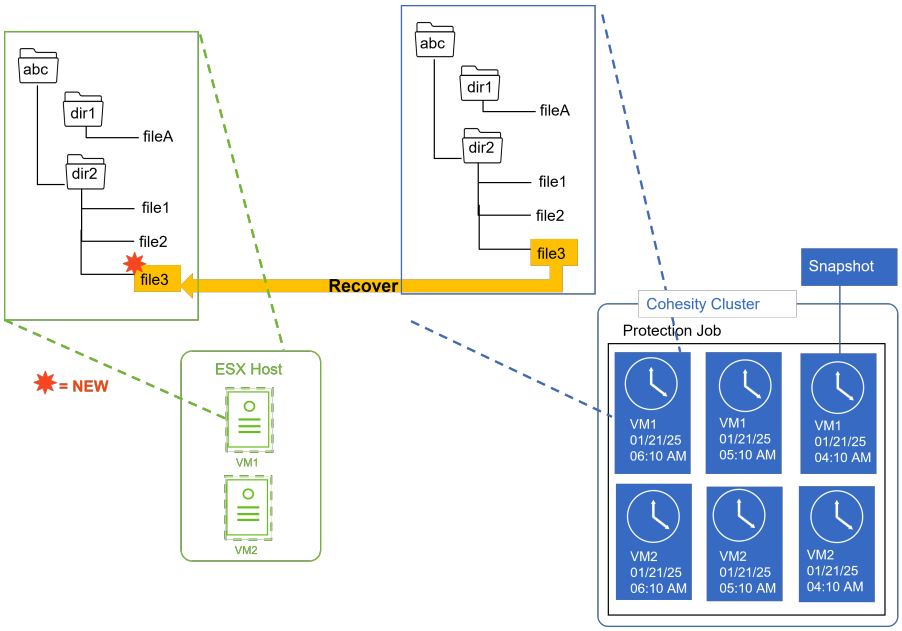

Download a File or Folder

In the figure below, one file is downloaded, but you can download multiple files and folders at once if they are from the same Snapshot. The Snapshot can reside on the local Cohesity cluster or on archives.

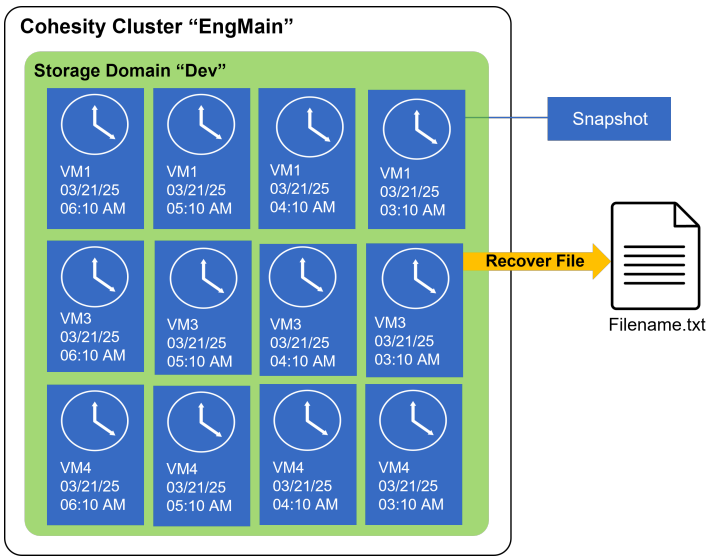

About Instant Volume Mount

Cohesity provides a fast and reliable way to access your data via the instant volume mounting feature. Some benefits include granular recovery of Microsoft Exchange and SharePoint data. Cohesity's simple workflows can instantly mount backup volumes to the Ontrack® PowerControls™ application, enabling application owners to perform granular recovery.

The following table lists the Cohesity backed-up solutions and the associated Ontrack® PowerControls™ application versions required to perform granular recoveries.

Cohesity Backed Up Applications	Supported Ontrack PowerControls Versions
Microsoft Exchange 2010 and later	9.2 and later

> **Note:** Microsoft Exchange (on-premises) requires Ontrack® software and license for granular recovery of Microsoft Exchange mailbox objects.

Instant volume mounting is supported for virtual and physical environments. This feature allows the selected backup volumes to be made available at the target location where you can then complete the desired operations. Instant Volume Mount is available for backed-up volumes, whether they are stored locally on the Cohesity cluster or archived to a registered external target.

For instructions on how to use the Cohesity Dashboard to instantly mount volumes, see Instant Volume Mount.

Mount Volumes

Backup volumes from a Snapshot can be mounted to a target in a registered Source or to a registered physical Server.

The following example depicts mounting a volume from a backup of a Microsoft Exchange VM to an Ontrack VM in the same vCenter to recover data. The action steps in the following procedure are represented in the figure below with blue numbered time squares:

T1) The Cohesity cluster clones the volume stored in a Snapshot and places them in a temporary View.

T2) The volume in the temporary View is mounted to the Ontrack_VM as represented by the black line between the Resource Pool and View.

T3) Now you can use the Ontrack_VM to perform the desired Granular Level Recovery (GLR) operations and recover data to the Exchange_Prod_VM.

Tear down the mounted volume when you are finished working with it.

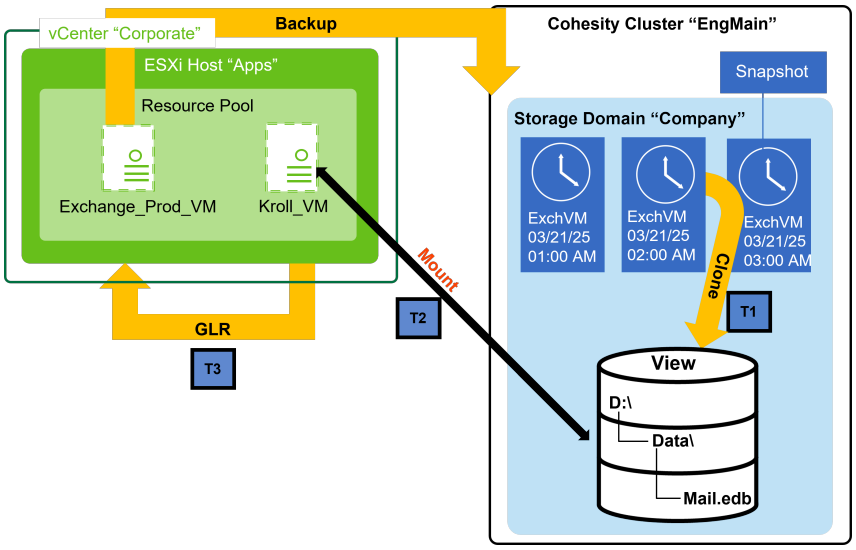

About Recovering Oracle Database

Cohesity provides the ability to recover the entire Oracle Database to its original location or to a different Oracle server that is registered on the Cohesity cluster. When restoring to the database's original location, you can perform a point-in-time recovery. You can also overwrite the original database.

An Oracle Database Recovery Task performs the following sub-tasks:

1. Determines which files to restore based on the specified recovery options.
2. Creates a View and copies the required files to the View.
3. Mounts the View on the Oracle server.
4. Runs the Cohesity-managed RMAN recovery scripts to perform the restore.
5. Optionally applies transaction logs to the database for the specified point in time.
6. Unmounts and deletes the View.

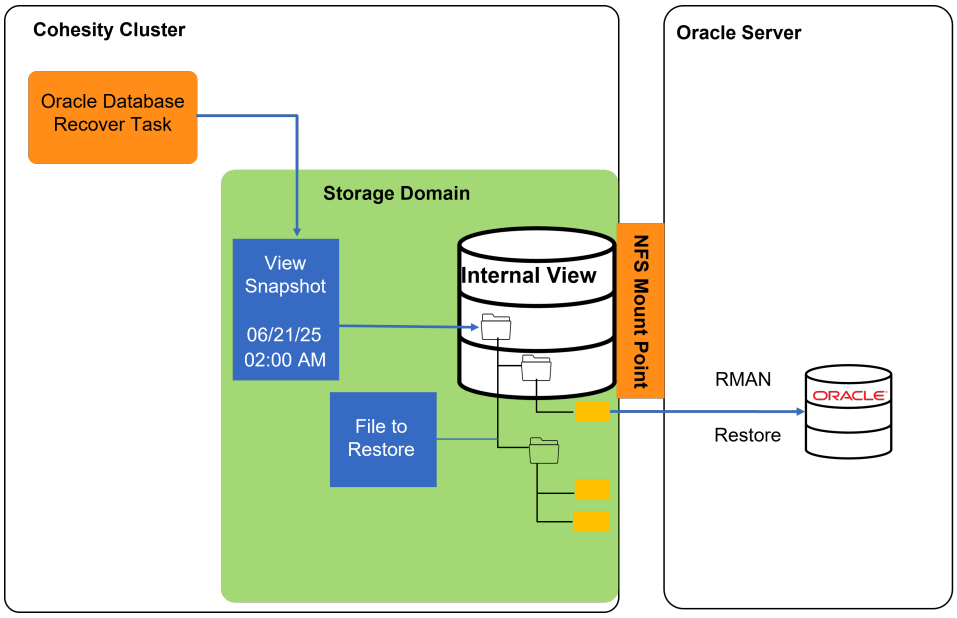

About Recovering Storage Volumes

Cohesity provides the ability to recover protected storage volumes from a Cohesity Snapshot created earlier by a Protection Group. The options to recover storage volumes are as follows:

- Recover Storage Volumes to the Original Location
- Recover Storage Volumes to a New Location
- Recover Storage Volumes as a Cohesity View

For instructions on how to recover storage volumes, see Recover NAS Storage Volumes.

Recover Storage Volumes on NAS

The following example depicts recovering data from a backup of NAS volumes and recovering one Snapshot.

During recovery, the Cohesity cluster utilizes SpanFS to instantly clone a volume's backed up data to a new View within the Cluster.

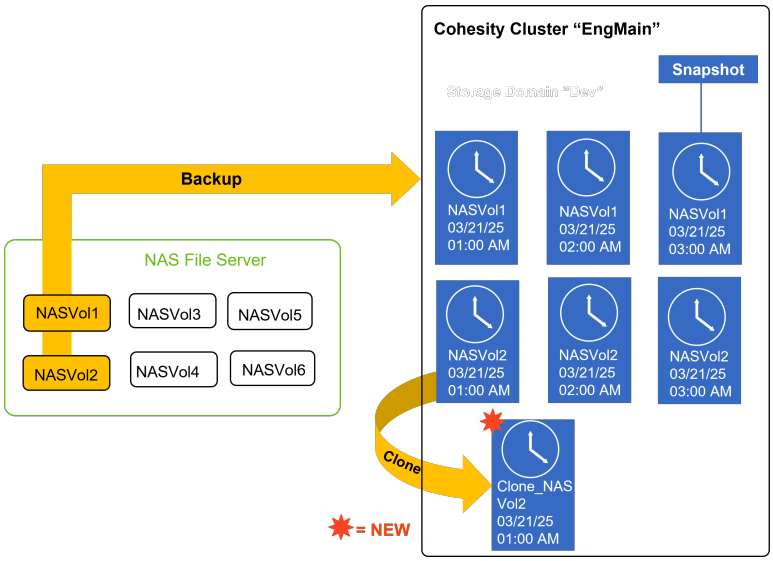

Recover Storage Volumes on Pure Storage

The following example depicts recovering data from a backup of Pure Storage volumes and recovering one volume back to the same array.

During recovery, the Cohesity cluster copies data from the Cluster to the storage array as a new storage volume.

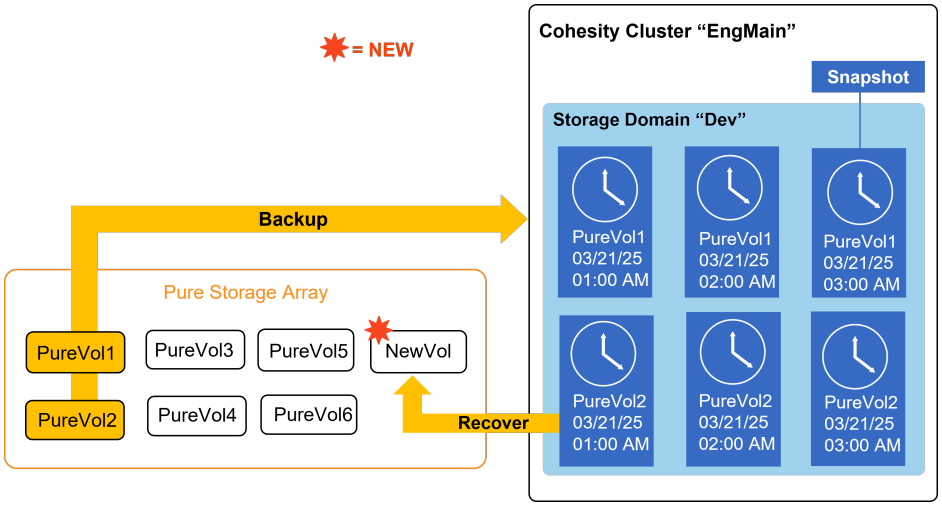

About Scale Out

Cohesity's scale out capability allows organizations to start with the right-size data platform that is ideal for their environment, and scale out as needed, on their own terms, one node at a time. No more pre-paying hefty upfront expenses for anticipated growth or incurring hidden costs associated with forklift upgrades of traditional scale-up systems.

The Cohesity cluster scales out horizontally. To increase the storage capacity and improve performance, you can add additional Nodes to the Cluster as represented by the procedure and figure below. The steps in the following procedure are represented in the figures with black numbered time squares: T1 .

T1) In this example, all the Nodes in Chassis#1 are assigned to the Cohesity cluster.

T2) Using the Cohesity Dashboard, two Nodes from Chassis#2 are selected to be added to the Cluster.

T3) The Cohesity cluster adds the selected Nodes from Chassis#2 to the Cluster. After new Nodes are added, the disk capacity of the Cluster is immediately increased and available for use by Protection Groups, Clone tasks and Recovery tasks. This addition of more Nodes improves performance because additional CPUs, memory, and disk bandwidth are available to the Cluster. Also, the Cluster distributes the blocks from existing Nodes to the new Nodes. In addition, the Cluster replicates the metadata about the block distribution onto these new Nodes.

The remaining Nodes in Chassis#2 are unassigned. These unassigned Nodes can be added to this Cluster at a later date or used to create another Cluster with the addition of another Node from another Chassis. (A Cluster requires at least three Nodes.)

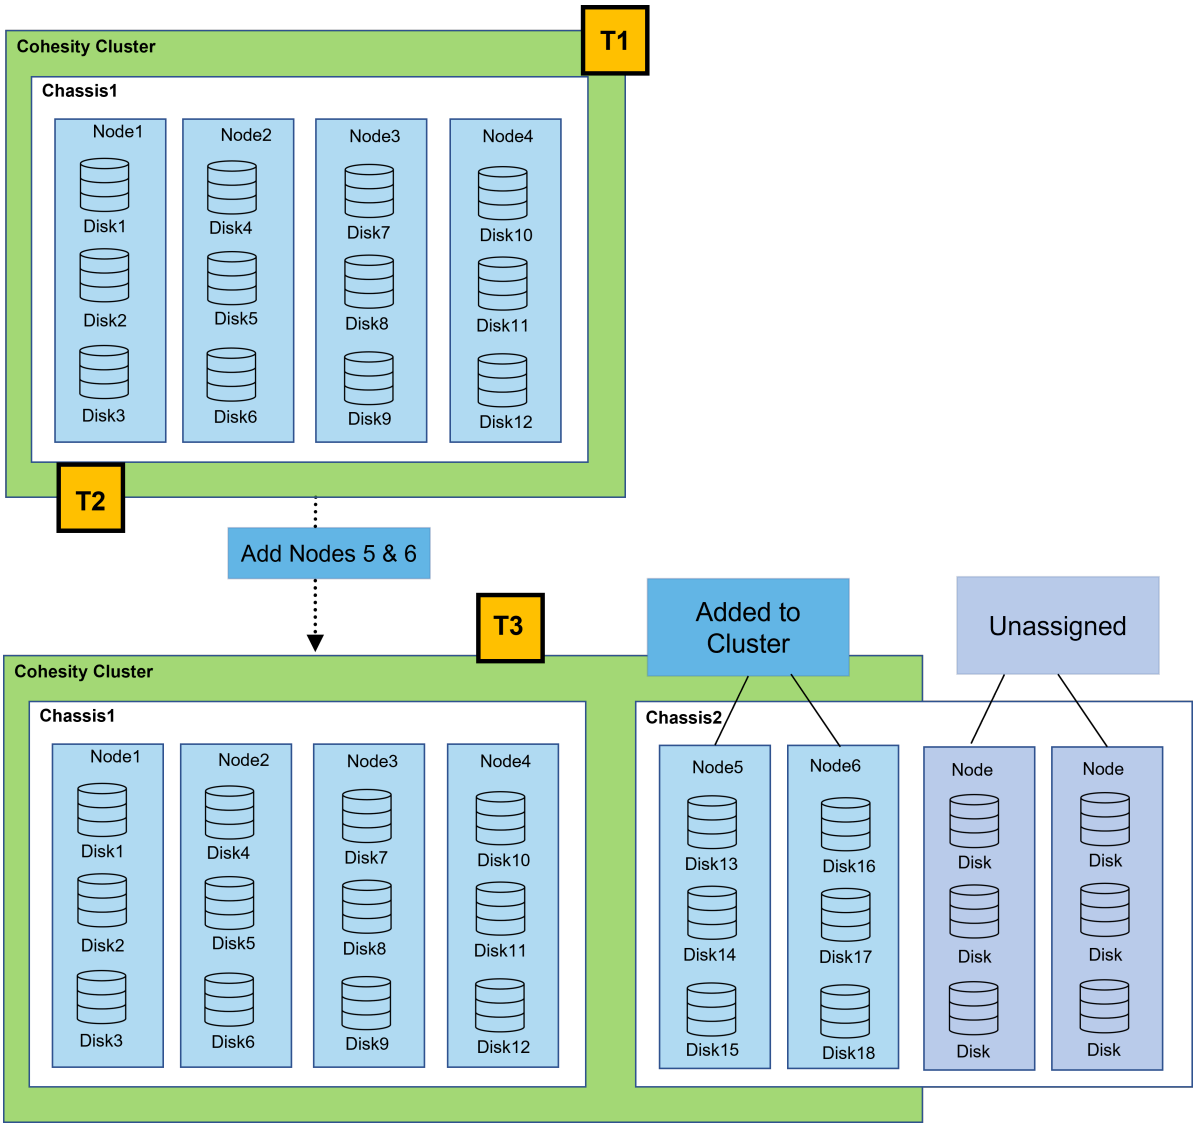

About Cloud Archival

The Cohesity cloud archive feature enables you to simplify off-site data management by integrating your archive solution with the Cohesity cluster. Archive's benefits include long-term data retention on low-cost storage to meet compliance and retention requirements. Optional incremental archival and source-side deduplication result in lower cloud capacity consumption, faster archive times and lower network bandwidth consumption during archival and retrieval.

Cohesity archive automatically copies an existing Snapshot created by Protection Groups located in a Cohesity cluster and stores it on a registered External Target.

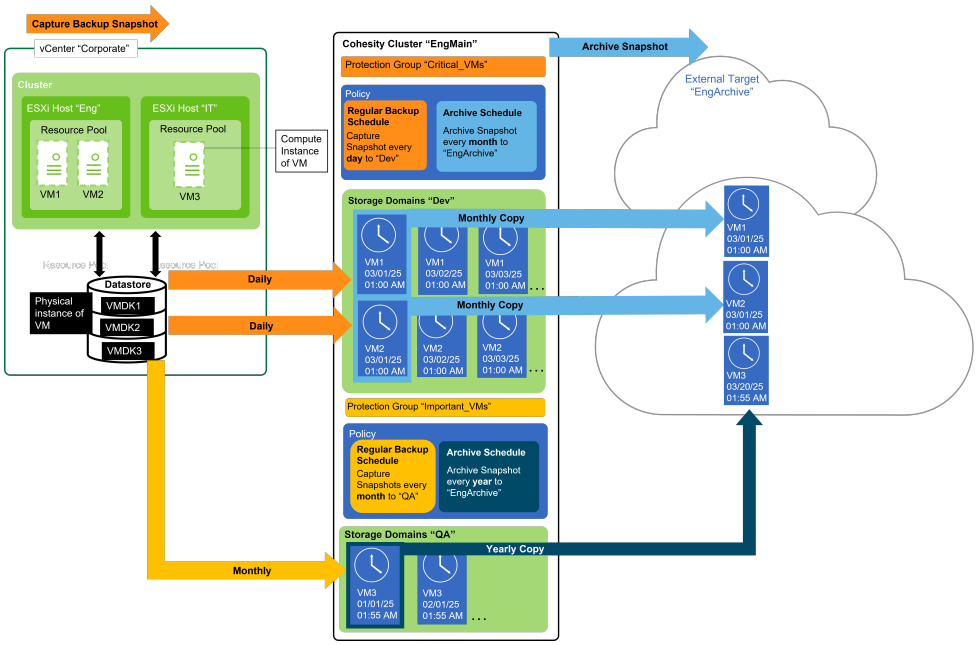

You specify a protection schedule in a Policy and that protection schedule defines how often a Protection Group using that Policy runs. In the example above, the "Critical_VMs" Protection Group, using a policy with a daily protection schedule, is capturing Snapshots of VM1 and VM2 and storing those Snapshots in the "Dev" Storage Domain. The "Important_VMs" Protection Group, using a different policy with a monthly protection schedule, is capturing Snapshots of VM3 and storing those Snapshots in the "QA" Storage Domain.

You can optionally specify an archive schedule in a Policy. The archive schedule defines how often Snapshots are copied from a Storage Domain in the Cohesity cluster and saved in an External Target by a Protection Group using that Policy. In the example above, the "Critical_VMs" Protection Group is making monthly copies of the Snapshots in the "Dev" Storage Domain and storing those Snapshots in the External Target named "EngArchive". The "Important_VMs" Protection Group is making yearly copies of the Snapshots in the "QA" Storage Domain and also storing those Snapshots in the External Target named "EngArchive".

The archive schedule for copying snapshots cannot be more often than the protection schedule for capturing snapshots. For example if snapshots are captured daily, the Cohesity cluster cannot archive them hourly but it can archive them daily or any value less often than daily.

Register an External Target

To specify an archive schedule for a Policy, you must register an External Target in the Cohesity cluster. Registering an External Target makes it eligible to be used by the Cohesity cluster. Data on an External Target resides outside of the Cohesity environment and is separate from the data stored on the local and remote Cohesity clusters. External Targets

can be cloud or tape and can be used for archiving or as a Cloud Tier. The Cohesity cluster requires the External Target's login credentials so a handshake can occur.

For detailed instructions, see Register an External Target.

External Target Encryption

When registering an External Target so it can receive archive data from your Cohesity cluster, you can choose to turn on encryption. When you turn on encryption, the data is encrypted before it is transferred to the registered Target. Using its own internal key management system, the Cohesity cluster generates a key. This key and data are encrypted and sent to the External Target over a secure channel (SSL).

With this configuration, to accomplish decryption, you do not need to upload keys, nor does the Cohesity cluster integrate with KMS or a similar key management system. When retrieving data from an encrypted External Target, the Cohesity cluster generates a key dynamically (which exists in memory only) that the Cohesity cluster uses to decrypt the key that is stored on the External Target. Only Cohesity knows how to decrypt the data because the key is based on markers within the stored data. Now the Cohesity cluster, and only a Cohesity cluster, can read data from the External Target. Even if an outside entity has access to the External Target, they cannot read the data because they cannot generate a key to decrypt the key stored on the External Target.

You can also restore a Cohesity cluster's archived data from the External Target if you lose access to the Cohesity cluster or data that is on the Cluster. To do this, you must have the lost Cluster's ID and your External Target credentials.

Cohesity recommends you store the cluster ID by registering the cluster at Cohesity Support Portal. Note that performing "clean data" on a Cohesity cluster and then re-installing changes the Cluster ID. Remember to store this new Cluster ID.

Additional Security

If you have selected **Internal KMS** or **KMIP Compliant** as the **Key Management Service (KMS) Type** for encryption, then you can enable additional security for the external target. By enabling additional security, you can manage the key manually instead of Cohesity managing the encryption key for data that is transferred out. If you choose to manage the key manually, then after registering the external target with the Cohesity cluster, you must download the key from the external target's details page and store it in a secure location outside of the Cohesity cluster. With additional security, only the Cohesity clusters where you have explicitly provided the encryption key can download data from that external target. However, the source Cohesity cluster (the cluster that archived the data) does not require the encryption key of the external target.

About Cloud Tier

Cohesity Cloud Tier enables an additional storage tier where cold data can be stored. After you register supported existing cloud storage as a Cloud Tier target with the Cohesity cluster, you can enable Cloud Tier to move rarely used and inactive data to the cloud when HDD used capacity exceeds a set threshold (typically 80%). To reduce the storage cost and network utilization, down-tiered data is always compressed in addition to being encrypted regardless of the cluster and Storage Domain settings. For deduplication, however, the down-tiered data in the External Target reflects the same deduplication settings as in the Storage Domain on the cluster. Cloud Tier benefits include utilizing low-cost cloud storage for rarely used and inactive data, allowing you to reclaim Cohesity cluster capacity for often used data. The actual amount of data that can be tiered to Cloud Tier depends on the deduplication ratio of your data. If the deduplication ratio is high, less data can be tiered to the cloud. As a general guideline, however, Cloud Tier roughly doubles your capacity.

The action steps in the following procedure are represented in the figure below with blue numbered time squares: T1 . The procedure assumes cloud storage is already registered as Cloud Tier and the HDD used capacity is beyond the threshold.

The Cloud Tier process consists of the following actions:

T1) A Cohesity cluster's Storage Domain contains two files, FileA and FileB. Both files contain data blocks that the Cohesity cluster recognizes as cold so it moves the data blocks to the Cloud Tier. Note that files may share data blocks, as indicated by Block#4 in the diagram.

T2) A system accesses FileB data Block#4 and Block#6. The Cohesity cluster recognizes that Block #6 is located on the Cluster and Block#4 is located on the Cloud Tier. The difference in block location is transparent to both the Storage Domain and the system accessing the file. This enables seamless access and simplifies data retrieval.

T3) If, based on additional file access, the Cohesity cluster determines that Block#4 is no longer cold, that data block is moved from the Cloud Tier to the Cohesity cluster.

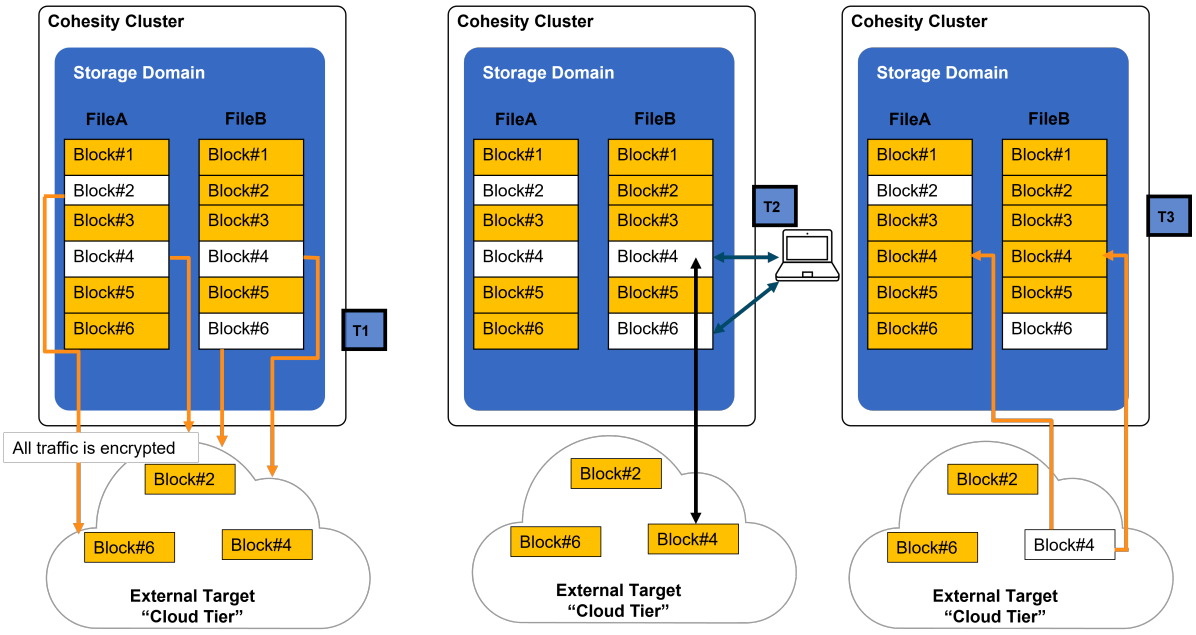

In the above diagram, white blocks represent data blocks stored on the local Cohesity cluster, and yellow blocks represent data blocks that have been moved to the Cloud Tier external target.

Register an External Target

To enable Cloud Tier for a Storage Domain, you must register an External Target in the Cohesity cluster. Registering an External Target makes it eligible to be used by the Cohesity cluster.

For detailed instructions, see Register an External Target.

Enable Cloud Tier

You enable Cloud Tier at the Storage Domain level. For detailed instructions, see Create or Edit Storage Domains. To view more details of an external target, open **Infrastructure > External Targets**. You can check the **Storage Consumed** value, which shows how much data the cluster has moved to that Cloud Tier target. You can also view an overview of Cloud Tier usage in the Storage section of the cluster dashboard once Cloud Tier is enabled.

About Tape Archival

The Cohesity tape archive feature enables you to create tape archives for long-term data retention.

The tape archival feature is supported with the QStar software only.

When archiving to tape, Cohesity automatically copies an existing Snapshot created by Protection Groups located in a Cohesity cluster and stores on a registered tape External Target. The Snapshot is stored in an Integral Volume that has already been defined in QStar. Cohesity software encryption can optionally be turned on. Compression is turned on by default.

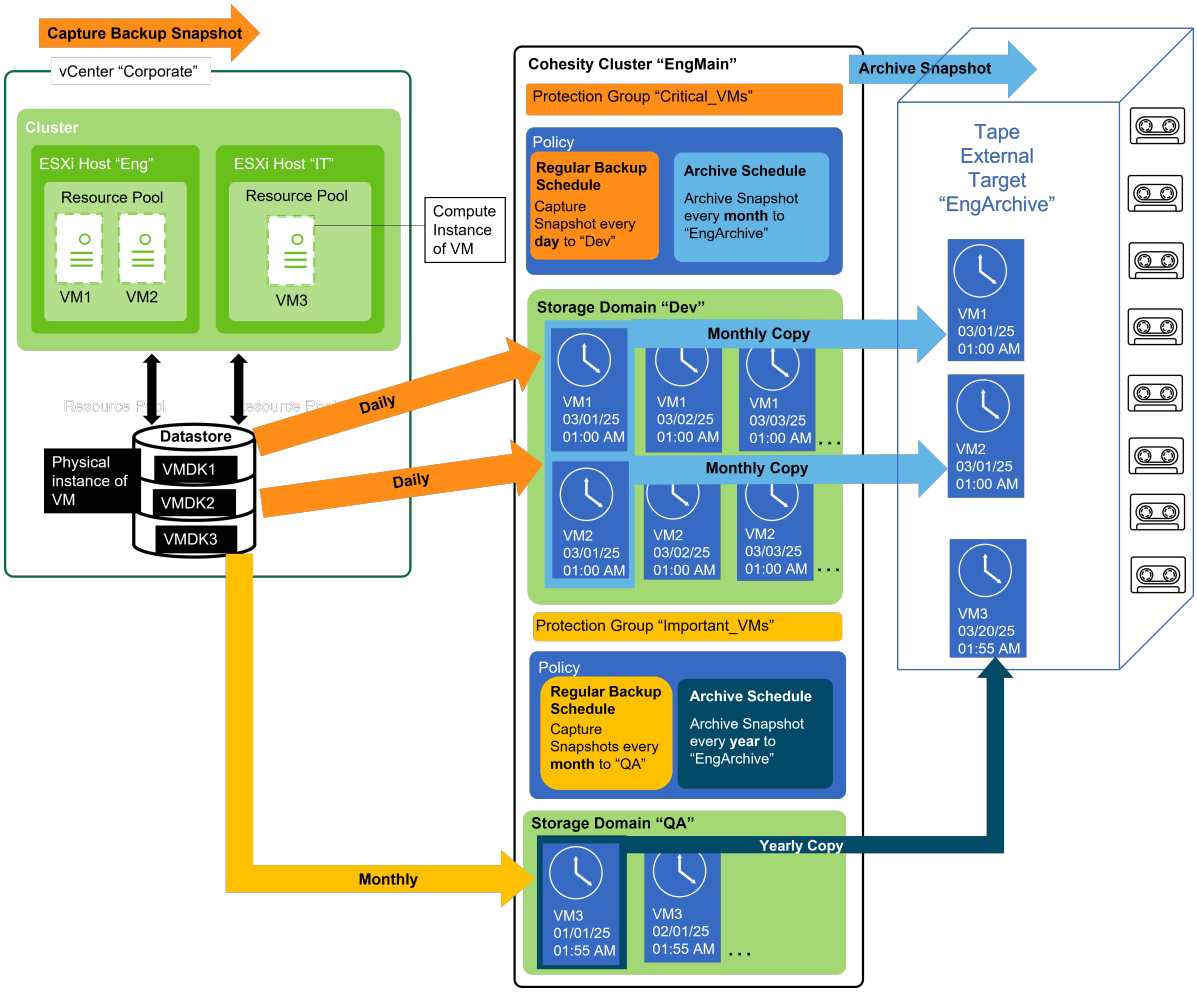

You specify a protection schedule in a Policy and that protection schedule defines how often a Protection Group using that Policy runs. In the example above, the "Critical_VMs" Protection Group, using a policy with a daily protection schedule, is capturing Snapshots of VM1 and VM2 and storing those Snapshots in the "Dev" Storage Domain. The "Important_VMs" Protection Group, using a different policy with a monthly protection schedule, is capturing Snapshots of VM3 and storing those Snapshots in the "QA" Storage Domain.

You can optionally specify an archive schedule in a Policy. The archive schedule defines how often Snapshots are copied from a Storage Domain in the Cohesity cluster and saved in an External Target by a Protection Group using that Policy. In the example above, the "Critical_VMs" Protection Group is making monthly copies of the Snapshots in the "Dev" Storage Domain and storing those Snapshots in the External Target named "EngArchive". The "Important_VMs" Protection Group is making yearly copies of the Snapshots in the "QA"

Storage Domain and also storing those Snapshots in the External Target named "EngArchive".

The archive schedule for copying snapshots cannot be more often than the protection schedule for capturing snapshots. For example if snapshots are captured daily, the Cohesity cluster cannot archive them hourly but it can archive them daily or any value less often than daily.

Register a Tape External Target

To specify an archive schedule for a Policy, you must register an External Target in the Cohesity cluster. Registering an External Target makes it eligible to be used by the Cohesity cluster. Data on an External Target resides outside of the Cohesity environment and is separate from the data stored on the local and remote Cohesity clusters. External Targets can be cloud or tape and can be used for archiving or as a Cloud Tier. The Cohesity cluster requires the External Target's login credentials so a handshake can occur.

For detailed instructions, see Register an External Target.

How to Reuse a Tape

To reuse a tape (QStar), perform the steps below:

1. In the QStar Administration Interface, select **Integral Volumes** > **Compaction**.
2. Click **Media Usage**. This will populate the list of tapes that are online and offline.

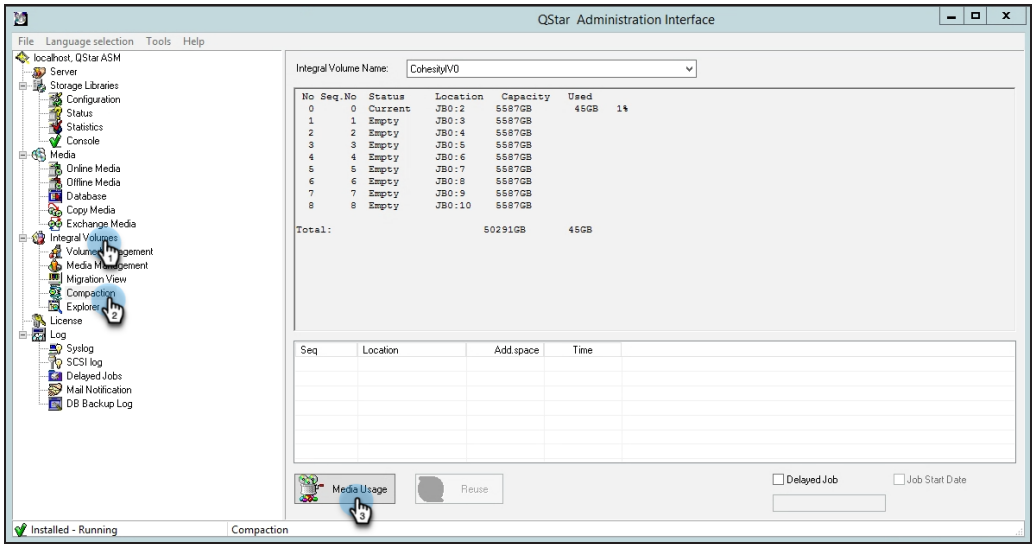

3. In the "**Used**" column, look for 0GB (0% utilization). This indicates that there is no active data on the disk. In the example below, JB0:13, shows 0% utilization.

> **Note:** Cohesity recommends to reuse tapes that are 0% utilized.

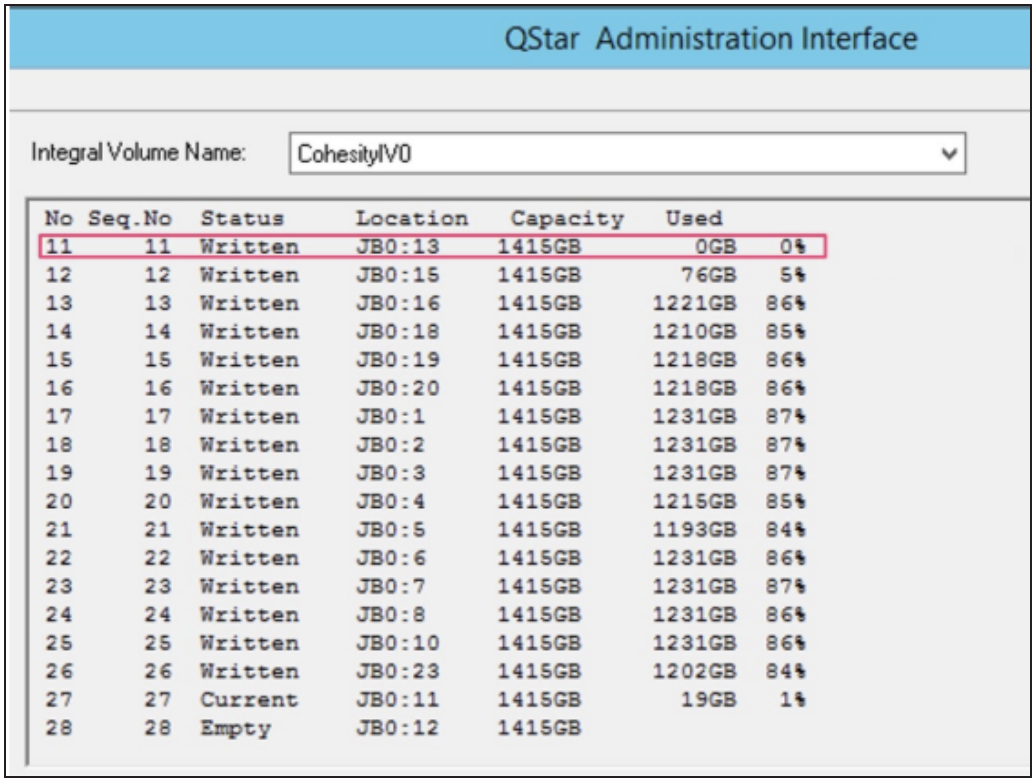

No	Seq.No	Status	Location	Capacity	Used	
11	11	Written	JB0:13	1415GB	0GB	0%
12	12	Written	JB0:15	1415GB	76GB	5%
13	13	Written	JB0:16	1415GB	1221GB	86%
14	14	Written	JB0:18	1415GB	1210GB	85%
15	15	Written	JB0:19	1415GB	1218GB	86%
16	16	Written	JB0:20	1415GB	1218GB	86%
17	17	Written	JB0:1	1415GB	1231GB	87%
18	18	Written	JB0:2	1415GB	1231GB	87%
19	19	Written	JB0:3	1415GB	1231GB	87%
20	20	Written	JB0:4	1415GB	1215GB	85%
21	21	Written	JB0:5	1415GB	1193GB	84%
22	22	Written	JB0:6	1415GB	1231GB	86%
23	23	Written	JB0:7	1415GB	1231GB	87%
24	24	Written	JB0:8	1415GB	1231GB	86%
25	25	Written	JB0:10	1415GB	1231GB	86%
26	26	Written	JB0:23	1415GB	1202GB	84%
27	27	Current	JB0:11	1415GB	19GB	1%
28	28	Empty	JB0:12	1415GB		

4. If the tape is offline:
    1. Select **Integral Volumes** > **Compaction**.
    2. From the "**Seq No**" column, make a note of the sequence number of the offline tape.
    3. Select **Media** > **Offline Media**.
    4. Match the "**Seq No**" with the "**Barcode**" of the tape.
    5. Load the tape that matches the barcode to the tape library.

    > **Note:** Once the tape is successfully loaded, an offline tape becomes online.

5. To reuse the tape:
    1. Select **Integral Volumes** > **Compaction**.
    2. From the bottom pane, select the tape.
    3. Click **Reuse**. This will add the tape to the queue for writing.

> **Note:** Reusing a tape will erase the existing data on the tape.

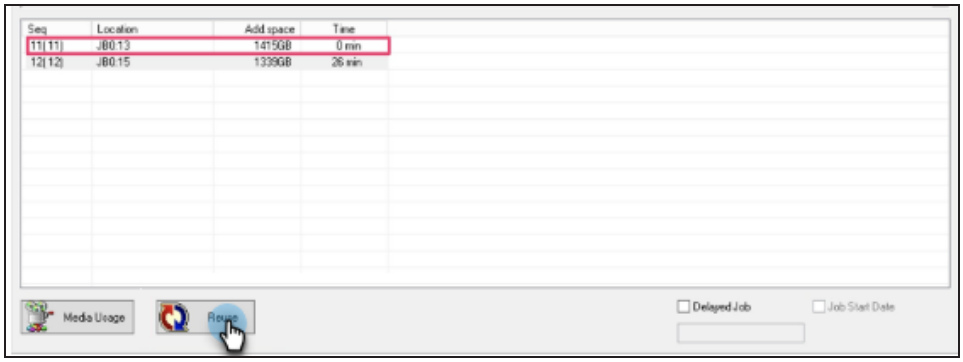

For the supported workflows for QStar, see Supported Workflows and External Targets.

About Fault Tolerance

Cohesity's highly resilient and fault-tolerant platform ensures organizations can continue conducting business during the following events:

- Hardware or software failure
- Maintenance and upgrades with non-disruptive rolling software upgrades

This topic describes how the Cohesity cluster is highly resilient and fault tolerant.

Cohesity provides fault tolerance at the following levels:

- **Cluster:** The failure domain settings at the cluster level protect the cluster against node, chassis or rack failures. You can specify how many node, chassis, or rack failures the cluster can tolerate while continuing to run and maintain strict consistency. For example, a cluster can survive during a catastrophic hardware (motherboard) failure on a node, a chassis backplane failure, or a power outage on a rack.
- **Storage Domain:** You can configure fault tolerance to survive disk, node, chassis or rack failures at a storage domain level (granular) than at the cluster level (global). A storage domain configured to survive three-disk failures will not lose data even if there are simultaneous three disk failures. In addition, erasure coding (EC) and redundancy factor (RF) levels allow for a further file data resiliency at the stripe level.
- **Metadata**: The Distributed Metadata Store is a fast key-value store containing all the metadata in the Cohesity cluster. It is a fault-tolerant store that Cohesity uses for keeping track of all the data and other stats in the system. Metadata is stored on SSD for performance reasons.

  How metadata is replicated depends on the configured resiliency:

- **1 Failure**: Tolerates either a single node or a single SSD failure.
- **2 Failures**: Tolerates either two nodes or two SSD failures.

For more details on Fault Tolerance and Best Practices, see Cohesity Fault Tolerance—Data Integrity for Modern Web-scale Environments.

Plan and configure fault tolerance based on your requirements. Use the alerts and advanced diagnostics features to monitor and protect your cluster against failures.

Prerequisites and Considerations

Review the following prerequisites and considerations to configure fault tolerance.

Prerequisites

Before you configure fault tolerance, ensure the following prerequisites:

- To enable rack as a failure domain at the global cluster level, every available chassis must be associated with a rack ID.
- All the nodes on a chassis must have the same rack ID.
- Minimum Component Requirements.

  The following table provides the minimum number of components (node, chassis, or rack) required in the Cohesity cluster to tolerate one failure of the respective component.

Component	Minimum Number of Components	Failures Tolerated
Rack	3	1
	5	2
Chassis	3	1
	5	2
Node	3	1
	5	2

Considerations

Before you configure fault tolerance, ensure that you understand the following conditions and considerations:

- You cannot delete a rack after you add and assign it to one or more chassis.
- After a rack is associated with a chassis, you cannot remove a chassis or reassign it to another rack.
- Maximum Number of Failures Tolerated.

  The following table provides the maximum number of component failures tolerated by a Cohesity cluster.

Component	Maximum Number of Simultaneous Failures Tolerated
Rack	2
Chassis	2
Node	2
HDD	3
SSD	2

For more information about configuring fault tolerance, RF, and EC options at a storage domain level, see Create or Edit Storage Domains.

Data Redundancy

The Cohesity cluster provides data redundancy to protect against disk, node, chassis, or rack failures. When files are stored on the Cluster, the Cluster breaks up the data into blocks that are distributed across the Cluster with a replication factor of two. This means that the Cluster stores two copies of each block on different disks and different Nodes. In the example below, one copy of Block#1 is stored on Disk#1 of Node#1 and the second copy is stored on Disk#7 of Node#3.

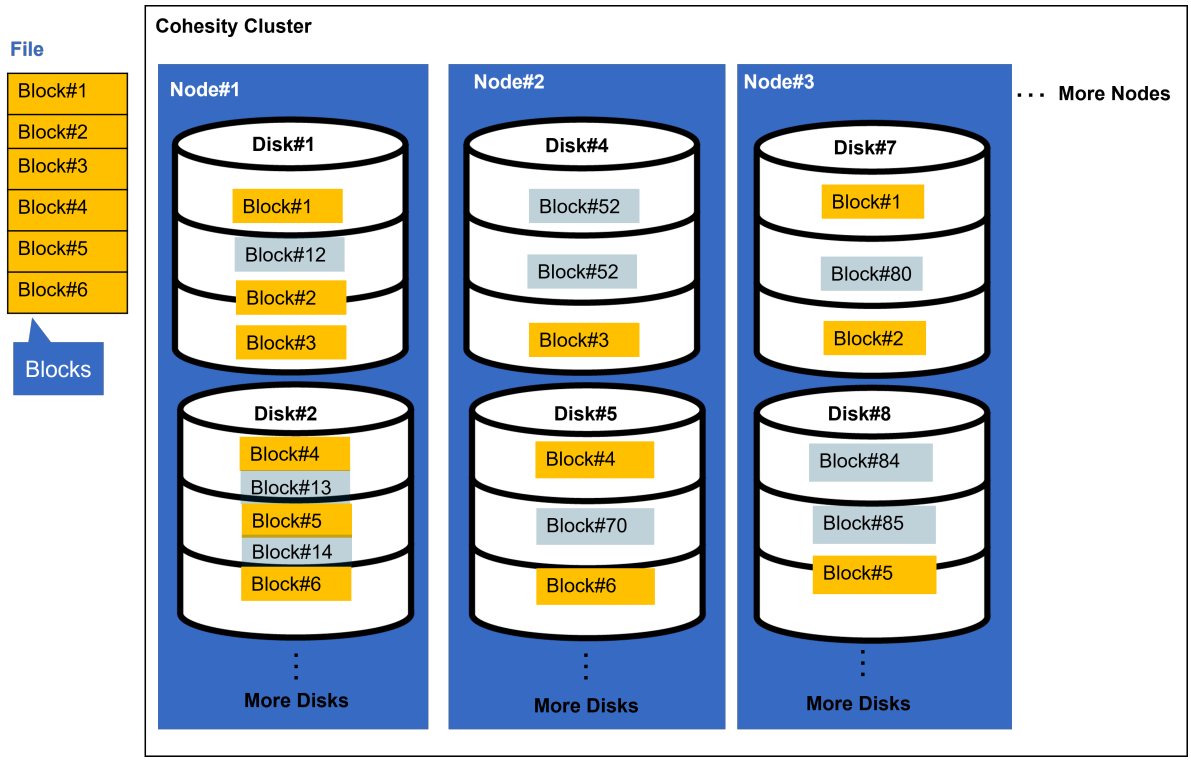

If a block is unavailable because a disk or a Node is unavailable and enough disk space exists on the remaining Nodes, the Cluster creates a new block on a different Node and a different disk from the existing copy of the block —always maintaining the replication factor of two. The Cluster makes itself fault tolerant again and no data is lost during this block restoration process.

> **Note:** The Cohesity rebalancing algorithm always creates a new copy of a block on a different disk and a different Node. Therefore, if a disk or Node becomes unavailable, another copy of the block exists on a different disk and Node or disk.

Disk Resilience

If a disk becomes unavailable, the blocks in that disk are unavailable to the Cluster. The Cluster detects that these blocks are under-replicated and replicates them from existing copies of the blocks in the Cluster as shown in the figures and procedure below. The steps in the following procedure are represented in the figures with orange numbered time squares: T1.

The data restoration process is described in the following example steps:

T1) When the Cohesity cluster determines that Disk#10 has errors, the Cohesity cluster generates a `DiskNotHealthy` Alert informing the Administrator of the problem. The

administrator triggers the disk replacement process by marking the disk offline and removing the disk.

T2) The Cohesity cluster periodically scans the blocks and determines that some blocks are under-replicated and replicates them. In this example, the Cohesity cluster determines that there is a single instance of Blocks#5 and Block#6 (because Disk#10 is unavailable) and so it finds the copies of these blocks and replicates them on other Nodes.

T3) An administrator replaces Disk#10 with a new Drive.

T4) The Cohesity cluster assimilates the new disk into the Cluster to restore the disk space capacity. The Cluster also distributes blocks to the new disk to improve performance.

No data is lost during this data restoration process.

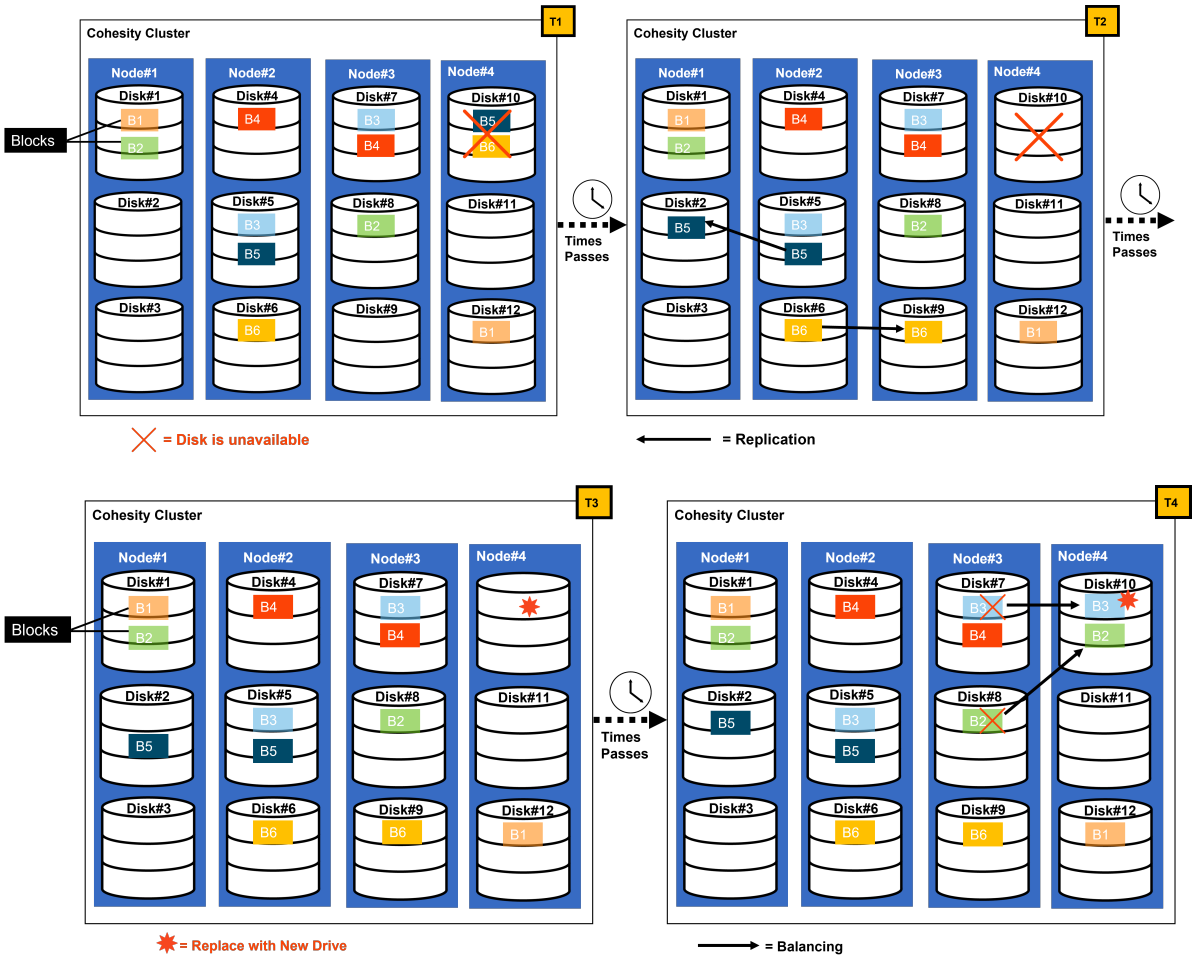

Node Resilience

If a Node becomes unavailable, the Disks in that Node are also unavailable to the Cluster. The blocks in the Disks are unavailable and the Cluster detects that these blocks are under-

replicated and replicates them from existing copies of the blocks in the Cluster as shown in the figures and procedure below. The steps in the following procedure are represented in the figures with orange numbered time squares: T1.

The data restoration process is described in the following example steps:

T1) Node#3 and all its disks (Disk#9, Disk#10 and Disk#11) are unavailable to the Cluster. Therefore, blocks (Block#3, Block#4 and Block#2) are unavailable to the Cluster.

T2) The Cluster periodically scans the blocks and determines that some blocks are under-replicated and replicates them. In this example, the Cluster determines that there is a single instance of Block#2, Block#3 and Block#4 (because Node#3 is unavailable) and so it finds the copies of these blocks and replicates them on other Nodes.

T3) An administrator replaces Node#3 with a new Node.

T4) The Cluster assimilates the new Node into the Cluster to restore the disk space capacity. The Cluster also rebalances the blocks to the new Node to improve performance.

No data is lost during this data restoration process.

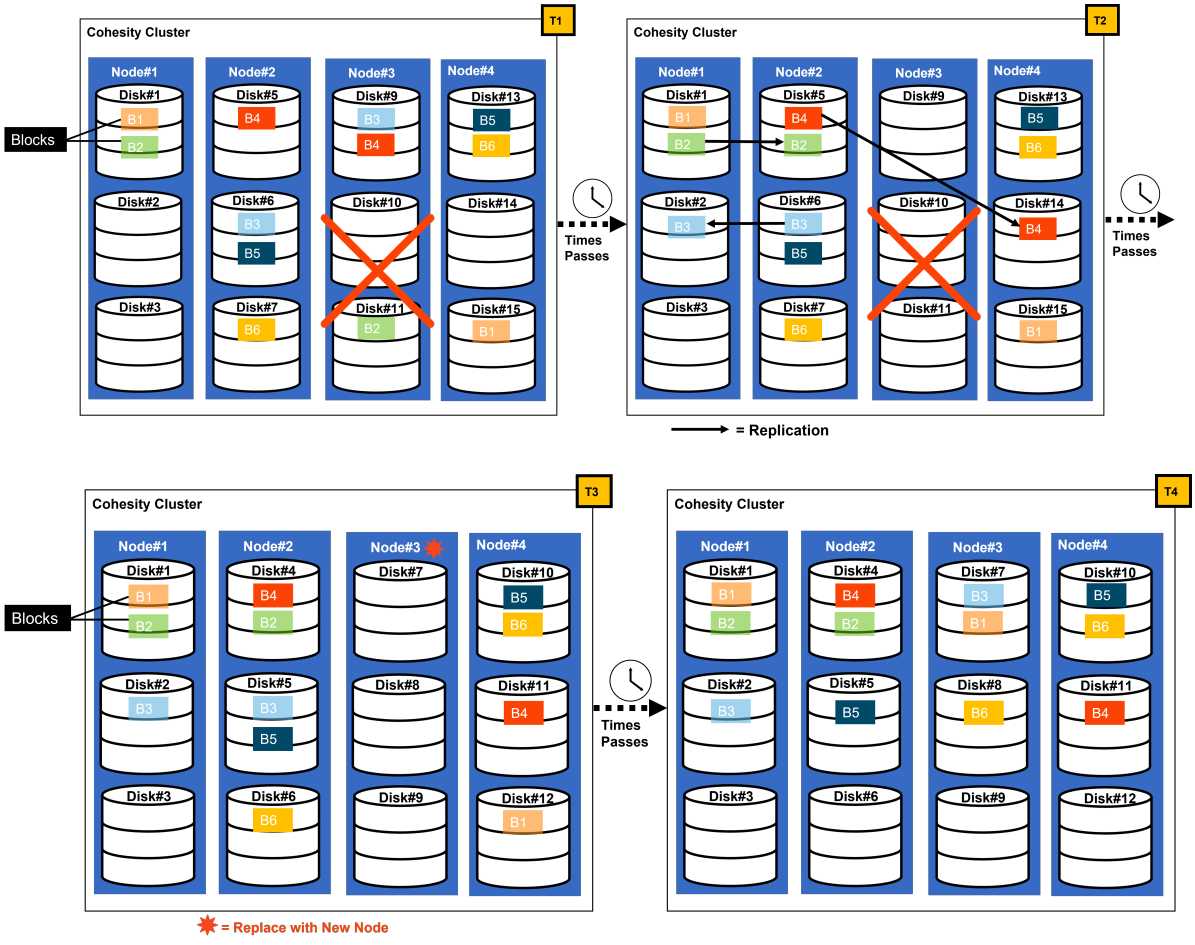

Node Failure or Node Down on a Three-Node Cohesity cluster

A three-node Cohesity cluster can sustain a node failure or node down and continues to satisfy the I/O read and write requests. The Cohesity cluster has the ability to automatically switch to an appropriate redundancy during a node failure or node down, and switch back when the cluster is restored back to the original state.

For example, if EC2:1 is configured with a fault tolerance setting of 1D:1N, the cluster switches to RF2 to maintain similar levels of data redundancy. Once the node is replaced, the cluster switches back to EC2:1.

A node failure or node down on a three-node Cohesity cluster will cause I/Os to fail under the following scenarios:

- **The cluster is in the process of healing from a previous disk/node failure** - If the data is not fully healed, turning off another node will result in data unavailability.
- **A node on the cluster is reaching its full capacity** - Turning off the node will result in the cluster triggering the healing and filling up space on the other two nodes.

> **Note:** During the above scenarios, the Cohesity cluster can tolerate a quick reboot.

For more details about node resiliency settings, see Disk and Node Resilience Settings.

Chassis Resilience

A chassis is a collection of one or more Cohesity nodes in a single form-factor. A chassis can contain up to four nodes. Cohesity provides protection levels for chassis resiliency that can be set separately for each storage domain. The allowed values depend on the fault tolerance settings. You can select the minimum failures to tolerate and the system adjusts the available options for RF and EC accordingly. For more details about chassis resiliency settings, see Chassis Resilience Settings.

Rack Resilience

A rack is typically a single tower of Cohesity-qualified appliances connected to a top-level switch. A single rack can include one or more chassis. It can also contain client or server machines from other vendors. Cohesity provides protection levels for rack resiliency that can be set separately for each storage domain. The allowed values depend on the fault tolerance settings. You can select the minimum failures to tolerate and the system adjusts the available options for RF and EC accordingly. For more details about rack resiliency settings, see Rack Resilience Settings.

Usable Storage

The following figure shows the formatted usable capacity based on various EC and RF settings. For example, an EC X:2 scheme provides a lower usable capacity than an EC X:1 scheme but a higher resilience by allowing for two stripe unit failures.

Node VIP Fail-over

The Cohesity cluster maintains a pool of Virtual IP addresses that it assigns to the Nodes as needed. If a Node goes down, the Cluster fails over the VIP to a different Node. Network traffic continues to be routed to the same VIP address but processing occurs on the new Node now associated with the VIP. Therefore, any Protection Groups or Recovery tasks running on the original Node continue to run on the new Node without interruption.

Network Redundancy

Each Node in a Cohesity cluster contains two 10 Gigabit Network Interface Controller (NIC) Cards with either SFP+ or RJ45 ports. If both NIC Cards are cabled into switches and one of the NIC Cards fails, the Node can continue to send and receive network requests using the other NIC card because the two NIC Cards are network bonded. Cohesity recommends cabling the two ports of a Node to different switches, so that if one switch goes down, the network traffic can reach the Node on the other switch.

Memory Redundancy

A dual in-line memory module (DIMM) failure either partial or full is considered a hardware component failure. In such cases, the node cannot function and therefore the failed DIMM must be replaced with a new DIMM for the node to resume memory operations.

CPU Redundancy

Each Node in a Cohesity cluster contains two CPU units housed in two CPU sockets. If one CPU unit fails, the Node can continue to process using the second CPU unit. The CPU units can be easily replaced because they are housed in CPU sockets.

Software and Hardware Failures on Cohesity cluster

The following table describes common software or hardware component failures and their Cohesity remedies. During any of these failures, all operations of the cluster, such as backups and recovery, file IO, replication, and archival are unaffected. For most, the self-

healer automatically heals the failure and brings the Cohesity cluster back to a normal state. The scenarios are listed in the increasing order of ease of resolution.

Failure	Remedy
Rack Failure	Cohesity can sustain up to two simultaneous rack failures. A minimum of three racks in your cluster is required to tolerate one rack failure. File data is reconstructed from the surviving rack that have units from the same stripes as the failed rack. This reconstructed data is placed in racks that have available (spare) capacity.
Chassis Failure	Cohesity can sustain up to two simultaneous chassis failures. A minimum of three chassis in your cluster is required to tolerate one chassis failure. File data is reconstructed from the surviving chassis that have units from the same stripes as the failed chassis. This reconstructed data is placed in chassis that have available (spare) capacity.
Node Failure	Cohesity cluster can sustain up to two simultaneous node failures. A minimum of three nodes in your cluster is required to tolerate one node failure. Some aspects for a node or a disk failure are: • Node failure results in both file data and metadata failures. • File data is recovered from the erasure coding or replication scheme. Metadata is recovered from one of the replicas.
HDD Failure	Cohesity cluster can sustain up to three simultaneous HDD failures at any point in time. HDDs are used to store the Distributed File Data. The number of HDD failures that can be tolerated depends on the RF or EC settings. Chunk files are units of replication. Cohesity cluster ensures that there is always a chunk file replica on other nodes, and internal chunk files representing the same user object (or user file) are spread across multiple disks.
SSD Failure	Cohesity cluster can sustain up to two simultaneous SSD failures. SSDs are used to store the Distributed Metadata and the Distributed Journal, and are replicated on either 2 or 3 nodes. Parts of the Distributed Metadata and Journal are stored on every node of the Cohesity cluster. An SSD failure can be considered the same as a node failure.

Failure	Remedy
Network Failure	Each node gets its own Node IP address, a VIP address, and a hardware management IPMI address. The DNS is set up to do round-robin, which gives each VIP a turn to satisfy client read or write requests and equally distributes them across all the nodes. **VIP Failover**: Cohesity cluster has a heartbeat mechanism that detects a node failure. If the directly connected network link goes down, the VIP of the failed node gets reassigned to another node, and node re-balancing similar to a node failure is initiated. Cohesity cluster nodes are typically configured with 2x10GbE connected to two separate switches. This allows for failure of a single switch or a single network port, while still allowing continuous operation. For more details, see Optimal Network Designs with Cohesity.
Boot Device Failure	Cohesity cluster boot device is isolated to a separate partition on flash. Since the Cohesity File System is already saved within persistent hard drives, there is no impact on a boot device crash. There is no inconsistency or file system repair required.
Power Supply	Cohesity cluster supports having a redundant power supply per chassis so that if one goes down, the other can still power the nodes.
Fans	Cohesity support a fan redundancy of N+1 (there is one extra fan per chassis). Hence, they will support one per node and an extra one to cover for failures.

Other Scenarios

The following table covers several other likely scenarios and how Cohesity cluster responds to each of the scenarios.

Event	Response
Node Addition	Cohesity cluster re-balancing is initiated by the self-healer. Every time a new node is added, part of the existing file data and metadata from the remaining nodes is rebalanced to the new node. This is done at a chunk-file level. The larger the cluster size, the faster the re-balancing. Common scenarios for node addition: • Cohesity cluster expansion, for example from 4 to 5 nodes. • Failed node replacement, when the Cohesity cluster size remains the same. For example, one node in a 4-node cluster fails and is replaced.

Event	Response
Node Removal	Node removal involves the removal of a node from the Cohesity cluster. When this happens, the Cohesity cluster re-balancing is initiated. Similar to a node addition, after a node failure or a node removal, the Cohesity cluster is restored. The file data and metadata are rebalanced on the surviving nodes. If the number of nodes falls to 2, there is an alert and the write policy changes from an EC-to an RF-based one, and writes continue. If a replacement node is added, the file data and metadata are re-balanced back to the replacement node. If a node goes down, other replicas ensure that newer disks are selected to store the data chunks that were lost from that node. Common scenarios for node removal from a Cohesity cluster are: • Node removal (graceful) during a hardware refresh. • Node failure. For example, one node in a 4-node cluster fails and needs to be replaced.
Node is Temporarily Down	This occurs when the node is assumed to be non-responsive after a timeout and re-balance has already kicked-in. The self-healer will stop healing (redistributing) when the node comes back. If the healing process is interrupted, some chunks will be over-replicated compared to others, but those chunks will be cleaned up later by the garbage-collection process
Attached Network is Temporarily Down	The node that is not reachable will be considered to be down. The self-healing process kicks in and re-balances the active nodes. When the network to the node is restored, the extra data and metadata replicas are cleaned up by the garbage collector.
Node Failure and Insufficient Space	If there is a node failure and the remaining nodes do not have sufficient available space to support the file data and metadata from the failed node, the self-healer will not perform re-balancing. In such cases, an alert is raised until a new node is added.
Non-Disruptive and Rolling Upgrades	Cohesity cluster architecture allows for a sequential node-by-node upgrade. Nodes can be upgraded to the latest versions without any interruption to client IOs. Any active IO access through the VIP (Virtual IP address) is transferred as the VIP is seamlessly transferred to another node.
Fault Tolerance with Cohesity for the Edge	Cohesity for the Edge is the virtual platform running as a VM on either VMware vSphere or Microsoft Hyper-V hypervisors using commodity hardware. An environment such as ESX or VMware vSAN with shared infrastructure can provide VM/storage migration capabilities on the underlying host failures.

Secure Software Development Practices

Cohesity reviews security at each phase of the software development life cycle. Vulnerability scanners are run to assess the state of the operating system, networking layer and services running on the Cohesity cluster. The scan results give the current open ports

on the Cohesity cluster. Any unused open ports and services are vulnerable. Hence, by default, these unused ports are closed. The kernel or system libraries are also updated for known vulnerabilities.

In addition to using a third-party software to conduct security assessments, Cohesity has its own in-house, dedicated Security Team to assess and improve security hardening in the product for each release.

Cohesity does not use outsourced or subcontracted components in the product. Some system integrators are used to build integration plugins, but that code is quality assurance tested and certified by our Engineering and QA Teams.

Cohesity is certified for FIPS 140-2 compliance.

Our Software Release Lifecycle

Cohesity regularly releases software updates, supporting a variety of capabilities—from new features to fixes. Cohesity software release types and their lifecycles are summarized in the below table.

Release Type	Description	Lifecycle
Feature Releases	Feature releases deliver new features. The numbering format for them follows this pattern 6.8, 7.0, etc.	Feature releases are supported for a minimum of 6 months from the first availability of it. A feature release is available via the controlled rollout process (see section below) and it progresses through various stages before it may earn LTS designation.
Maintenance Releases	Maintenance releases deliver fixes and minor enhancements in a release. They are indicated by the third digit of the release number such as 6.8.1. We may deliver intermediate update releases which contain fixes and minimal enhancements. These releases are represented by "_u<YYYYMMDD>" after the 3 digits, such as 6.8.1_u2 0240415.	The lifecycle of maintenance releases is aligned to that of a feature release and/ or LTS release. Urgent fixes might be addressed with a patch that is always included in all subsequent patch or maintenance releases.

Release Type	Description	Lifecycle
Patch Releases	Patches are urgent fixes delivered for maintenance releases. They are indicated by 'p<YYYYMMDD>' for product fixes and 'p<YYYYMMDD>s1' for product fixes plus CVEs (Common Vulnerabilities and Exposures) fixes. They are provided as needed (about 2 to 6 weeks).	Patches are rolled up cumulatively into the next maintenance release.
LTS (Long-Term Support)	LTS is a designation made for a feature release when specific stringent quality and stability requirements are met. From the moment a feature release earns an LTS designation, it is referred to as an LTS release. Our goal is to have two LTS releases available to customers at any given time.	LTS releases are supported for 12 months from the date of LTS designation and minimum of 6 months of overlapping coverage between two LTS releases and are available to all customers. Customers using LTS releases are encouraged to upgrade to the latest maintenance release as each becomes available.

The diagram below illustrates the software string key.

Release	Software Versioning Scheme	Naming Scheme	Example
Major	SemVer	X.Y _release-YYYYMMDD_SHA	7.0_release-20230925_2fecc890
Minor	SemVer	X.Y.Z_release-YYYYMMDD_SHA	7.1.1_release 20231128_900a8df7
Unified Release "U"	CalVer	X.Y_release-uYYYYMMDD-SHA X.Y.Z_release-uYYYYMMOD-SHA	6.8_release-u20231219-b55c185c 6.8.1_release-u20231219-b55c185c
Product Patch "P"	CalVer	X.Y-pYYYYMMDD ##-SHA X.Y.Z-pYYYYMMDD-##-SHA	7.1-p20240204-00-f210474 7.1.1-020240204-00-210474e
Security Patch "S"	CalVer	X.Y-pYYYMMDDs1-##-SHA X.Y.Z-pYYYYMMDDs1-##-SHA	7.1-p20240204s1-00-f210474e 7.1.1-p20240204s1-00-f210474e
Hotfix	CalVer	«service name»-X.Y_pYYYYMMDD-patch-##-SHA or «service name>-X.Y.Z_pYYYYMMDD-patch-##-SHA	atom-7.1_p20240204-patch-00-f210474e atom-7.1.1_p20240204-patch-00-f210474e

Release Controlled Rollout Process

Every feature release is controlled for a while and requires Cohesity approval for access. A controlled rollout process provides staged quality gates to ensure the release is stable and suitable for all production environments. Feature releases follow the standard support processes and customers can request access through the Cohesity account team or Cohesity Support. These releases are only visible to approved customers on downloads.cohesity.com.

Feature Release Stages

A feature release progresses through four stages:

1. **Beta:** This release is intended for non-production environments for feature verification and testing
2. **Initial Production Release (IPR):** This release is intended for non-mission critical workloads.
3. **Production Release (PR):** This release is intended for non-mission critical workloads and proof-of-concepts.
4. **Production Release Mission Critical (PRMC):** This release is intended for all workloads including mission-critical workload.

The diagram below illustrates an example of a feature release progressing through the stages and earning LTS designation:

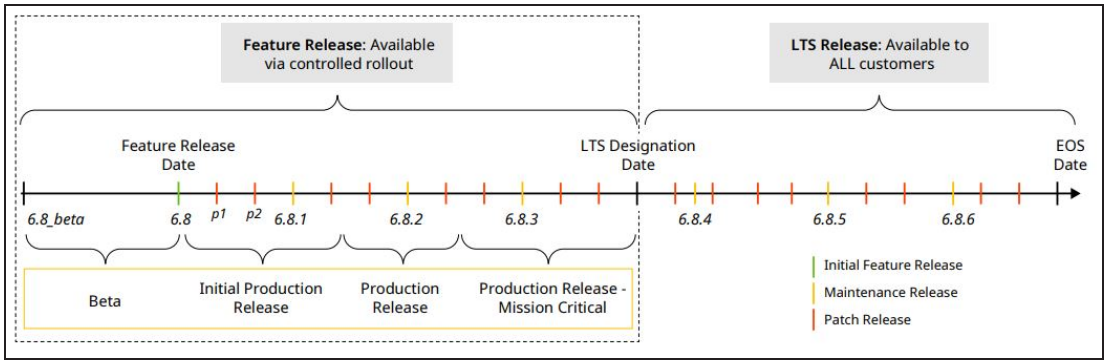

You can refer to the Cohesity Global Support and Services Handbook for more information LTS, Feature, and Maintenance releases.

Security Vulnerabilities

Cohesity's dedicated Security Team assesses and responds to security vulnerabilities reported in the products. We proactively inform customers of vulnerabilities via Support Portal alerts and emails. Knowledge Base (KB) articles are published to document the impact of the vulnerability and any required next steps.

Critical vulnerabilities are prioritized and fixes are generally implemented in the next product release. Cohesity resolves vulnerabilities based on their impact (as evaluated by Cohesity). We provide:

- **15-day resolution**. If the CVE score is between 9.0 and 10.0.
- **30-day resolution**. If the CVE score is between 7.0 and 8.9.
- **90-day resolution**. If the CVE score is between 4.0 and 6.9.
- **180-day resolution**. If the CVE score is between 0.1 and 3.9.

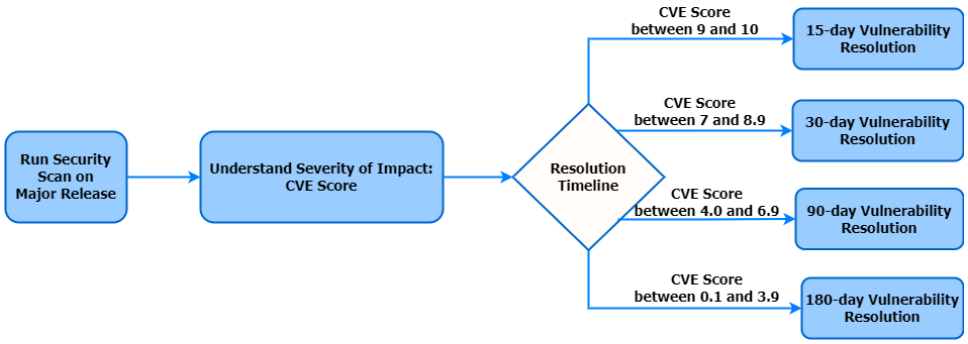

Cohesity Testing Stages

Cohesity performs extensive feature testing, system testing and release readiness testing. The following figure demonstrates the activities performed for each testing stage.

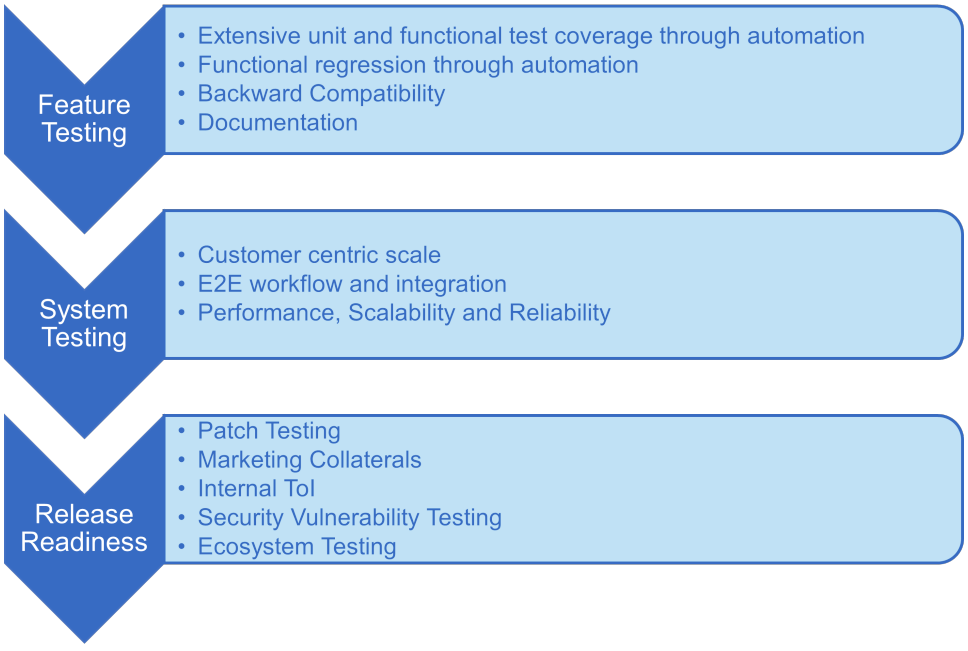

Secure Product Development Lifecycle

Cohesity follows a Secure Product Development Lifecycle to deliver secure products with the security vulnerabilities contained throughout the product's lifecycle. Cohesity follows the following six practices for the Secure Product Development Lifecycle (SPDL).

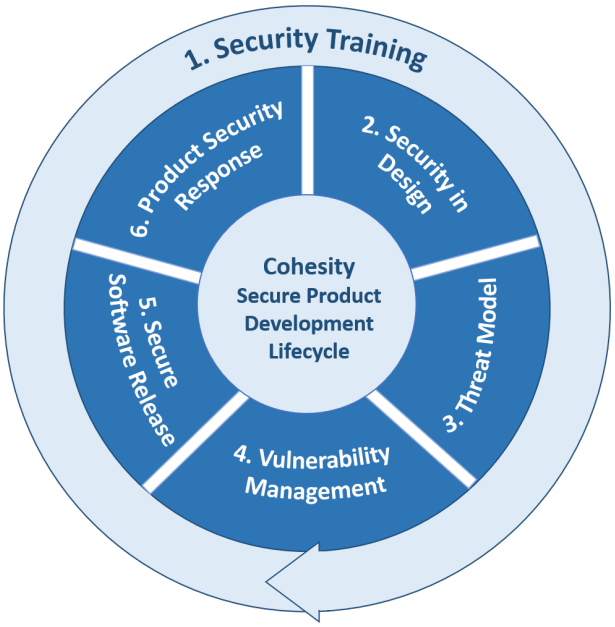

Security Training

Cohesity mandates secure coding training for all the developers to follow security best practices in product development. Teams in Cohesity such as developers, product managers, development managers, release management, and QA are provided with adequate training, awareness, and resources to influence product security during the product development lifecycle successfully.

Security in Design

In the design phase of product development, security is fundamental to develop highly secure systems and applications. To ensure the products' security posture and compliance or certification requirements are maintained with new features, design documents and design reviews are used to align the product design with the current security best practices. Cohesity follows security principles such as secure by default, fail securely, secure crypto algorithms, and so on.

Threat Model

The threat model is an engineering technique used to identify threats, attacks, vulnerabilities, and countermeasures that could affect an application or system. The threat model assists in design practice to meet security objectives in design and reduce risk. The threat model also helps in effective vulnerability management. Cohesity regularly updates the published product threat model based on new features and changes in threats.

Vulnerability Management

Cohesity adapts some of the best vulnerability management guidances from ISO 27001:2013 Control A.12.6.1 and NIST Cybersecurity framework Version 1.1 (Controls RA-1, RA-5, IP-12, CM-8, MI-3). For vulnerability management, Cohesity follows the following six practices.

Vulnerability Management Policy

Vulnerability management policy defines the severity classification of vulnerabilities based on the CVSS score and SLA for each severity class. It also defines the metrics for the backporting of security defects.

Penetration Testing

Penetration testing of products is conducted using third-party engagement and internal continuous mode using various automated techniques integrated into the release cycle. The potential vulnerabilities found in penetration testing are addressed as per the vulnerability management policy.

Static Code and Binary Analysis

Static analysis security testing (SAST) is an automated testing technique used to analyze source code to identify security flaws. Cohesity conducts static code and binary analysis using well known commercial off-the-shelf (COTS) tools.

Dynamic Scanning

COTS DAST scanners are integrated into the development pipeline to scan all significant development branches. Any vulnerabilities found are addressed meeting the SLA defined in the vulnerability management policy. Automated regression tests of previous vulnerabilities are run to avoid any security regression issues in future releases.

Third-Party Component Security

All imports of third-party components are tested for security risks. The third-party software inventory is maintained to assist in security risk mitigation and third-party component vulnerability patching. Third-party vulnerability patching follows a vulnerability management policy. Third-party component security also covers OEM firmware like BIOS and BMC software.

Product Infra and Tools Support

Engineering infrastructure tools are always maintained up to date, and security protections are configured. Security checks or options for compiler and linker are always enforced.

Secure Software Release

Cohesity follows a secure software release practice that brings predictability to the security posture of software releases. An acceptable level of security quality metrics is defined, and Cohesity ensures that each release meets the minimum goal.

JIRA workflows are defined to track the complete vulnerability life cycle using JIRA labels and fields and hence provide visibility of the release's security posture at any given time.

Product Incident Response

Cohesity uses a product incident response plan that supports analysis, mitigation, and remediation of vulnerabilities in the products. The plan also covers responsible disclosure from third-party researchers and customers.

Secure Acceptance Procedure

Cohesity uses the following procedures to ensure the integrity of its hardware and software products.

Hardware

Cohesity products are built by contract manufacturers (CM). CMs ensure the integrity of the products by:

1. Building to controlled specifications documented in Cohesity's Product Lifecycle Management (PLM) system and their shop floor control systems.
2. Performing tests jointly developed by the CM and Cohesity.

Software

The software image provided to customers is protected by using the SHA256 and MD5 hashing technique.

1. The customer logs on to the Cohesity Support Portal.

   In the downloads section, a link to a software package is provided, along with the corresponding MD5 checksum and SHA256 for the software package.
2. The customer downloads the software package and compares its checksum with the checksum provided on the Cohesity Support Portal. They might use commands such as:

   ```
   md5sum <software-package>.tar.gz
   sha265sum <software-package>.tar.gz
   ```
3. If the MD5 checksum or SHA256 values do not match, the integrity of the software package may have been compromised. The customer can try downloading the software package again to resolve the issue. If necessary, they can contact Cohesity Support for further assistance.

Cohesity Security Features

This topic explains how Cohesity builds security into the Cohesity cluster and is intended for:

- Administrators responsible for the configuration, maintenance and operation of the Cohesity cluster.
- Security engineers reviewing Cohesity for compliance with their company's security policies.

The Cohesity cluster provides the following security features.

Data Segmentation

The Cohesity cluster architecture provides virtual data isolation through Storage Domains. Each Storage Domain encrypts the data stored in it using its own independent keys.

During initial cluster setup, encryption can be enabled for the entire cluster so that the data in all Storage Domains will be encrypted. If encryption is not enabled during cluster creation, it can be enabled at the Storage Domain level. Note that once encryption is set during cluster or Storage Domain creation, it can not be disabled later.

Each Cohesity View is assigned a Quality of Service (QoS) Policy. A QoS Policy determines the priority of I/O (when contention occurs).

Data at Rest Security

The Cohesity file system provides full at-rest encryption based on the strong AES-256-GCM (Galois/Counter Mode) standard. Only 256-bit keys are used for encryption. The encryption architecture provides high security while maintaining the flexibility to optimally leverage the available hardware and software resources.

The built-in software-based encryption is hardware-accelerated through the latest Intel processors that support the AES-NI instruction set. Due to hardware acceleration, the software-based encryption route has minimal impact on performance.

Keys can be obtained in three ways:

- Internal Key Manager
- External Key Manager
- Cloud-based Key Manager

Keys can be rotated periodically. The default rotation is set to 90 days.

User data is encrypted before it is stored on any storage media used. The data is decrypted after it is read from the storage media in memory. The storage media can be SSD, HDD, tape, and cloud storage. This includes credentials for data sources protected by Cohesity.

Physical agents are designed to respond solely to cluster-issued commands. Thus, they do not have the ability to extract data, such as credentials or other sensitive information, from the cluster.

Data in Transit Security

The data being transmitted to a Cohesity cluster or being transmitted from a Cohesity cluster to External Targets is encrypted for security. Examples of data movement are ingesting data from a primary site, replicating between remote offices or transmitting data to public cloud.

The agent encryption cipher in use can be viewed in the Pulse logs or the agent report per the UI changes.

Primary Site to a Cohesity cluster

When the Cohesity backup software ingests data from the primary storage, it first establishes a secure connection with the primary storage over HTTPS. Cohesity backup software determines which data blocks have changed from the previous backup run by communicating with the primary storage over the secure connection, and then transfers all the changed blocks over the HTTPS connection.

Communication between the Cohesity cluster and Cohesity backup service (agent) is secured using SSL/TLS integrated transport. They perform mutual authentication and all data exchanged is encrypted.

The encryption is done using the Advanced Encryption Standard-256 protocol. The cipher suites used for backups through agents are:

1. `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`
2. `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA38`
3. `TLS_RSA_WITH_AES_256_GCM_SHA384.`

Replication from One Cohesity cluster to Another

A source Cluster starts a replication session by initiating a handshake with the target cluster over an HTTPS connection. Upon receiving the handshake request, the target Cluster obtains a session key from its local Key Management Server and sends it to the source cluster over the HTTPS connection.

The source Cluster stores the received session key in its local Key Management Server and completes the handshake. At this point, both the source and the target Clusters have the session key.

The source Cluster uses the session key to encrypt the data and transfer it securely for that replication session. The encryption uses Advanced Encryption Standard-256 protocol in Galois/Counter Mode (AES-256-GCM).

Data in Cloud Security

The cloud can be used:

- As another tier of storage (Cloud Tier)
- To store long-term retention copies (CloudArchive)

- Archive data directly to an external target without retaining a copy of the data on the cluster (CloudArchive Direct).

Cohesity supports all the major cloud vendors, including Amazon, Microsoft Azure, Google and Oracle, and any S3-compatible storage.

- You register the cloud storage as an External Target on the cluster and provide the cloud storage access keys. The access keys are encrypted and stored on the Cohesity cluster using AES-256 encryption in Galois/Counter Mode (AES-256-GCM). You can update the access keys from the Cohesity Dashboard or by using the Cohesity REST API. (Cohesity does not rotate the access keys.)
- Though the access keys are specific to the cloud storage provider, the Cohesity's key storage security is the same regardless of the cloud provider. Cohesity is FIPS certified and has an internal software-based key management system (KMS) that stores the key encryption key (KEK).
- Cohesity establishes a connection to the External Target *only* when making REST API calls to perform an archive or recovery operation.
- Cohesity does not include the External Target access keys in the API request. The request is signed using the External Target's access key and the signature is sent using the Authorization request header.
- Cohesity uses TLS 1.2 to encrypt in-transit data. Currently, Cohesity does not use mTLS (2-way SSL) in communication with External Targets.

CloudArchive and CloudArchive Direct use their own encryption, independent of data-at-rest encryption. The cipher suite used for encryption is `AES-256-GCM`.

- Cloud Archive encryption is on by default but can be turned off. Cohesity recommends keeping it on.
- Both the data encryption key (DEK) and KEK are encrypted, distributed for resiliency across the cluster and stored in SSD.
- The cluster creates a new DEK for each External Target. Cohesity provides two options. You can:
  - Store the data, metadata and encryption keys (DEK/KEK) in the External Target.
  - Store the data, metadata, and DEK only in the External Target. This requires the KEK to be downloaded and stored outside of the cluster. Without the downloaded keys, future CloudRetrieve operations will not be possible.

For Cloud Tier, encryption is always on and cannot be turned off.

Key Management

A Key Management Service (KMS) helps you manage the keys that encrypt data and secures it from unauthorized access. The Cohesity cluster has a built-in key manager or can

be configured to use external key managers - KMIP-based and Amazon Web Services (AWS) KMS - to manage encryption keys.

Ports Employed

For a list of firewall ports used by the Cohesity cluster,

Cluster Management Methods

The Cohesity cluster can be managed using the following methods:

- The Cohesity Dashboard web browser interface.
- The command line interface (CLI) available for Linux, Windows and iOS systems.
- The REST API endpoints on any platform using a REST client.

All access methods listed above require an SSL connection, and the Cohesity cluster prioritizes TLS 1.3. For backward compatibility, TLS 1.2 ciphers are also supported, but administrators can disable TLS 1.2 for enhanced security. The connection uses strong ciphers for the established session. To update the cipher list, see iris_cli cipher update. When the Cohesity cluster is created, a self-signed X.509 certificate (PEM format only) is generated for this purpose. The certificate must have the "Common Name" field set to the server's hostname (for example, a VIP's DNS name). Certificates are stored in the boot disk. They are replicated to other Nodes as part of Cluster formation using RPC calls. When the key expires, the web browser marks the session as invalid. Users can import new certificates to re-establish the session.

Users can overwrite this X.509 certificate according to their company policy.

Authentication

Authentication is required for all management methods (the Cohesity Dashboard, CLI and API). Only a user with an account on the Cluster can authenticate to the Cluster. Cohesity supports three types of accounts.

- **Local User Account**—A user account is created locally on the Cluster. This user can access only this Cluster. When the account is created, a password is assigned to the user and required to authenticate to the Cluster. This password is stored on the Cluster.
- **Active Directory**—A user account that exists on an Active Directory server. This user can access all the Clusters that have joined the Active Directory. The user uses the Active Directory password to authenticate to the Cluster. The Cluster authenticates the user with Active Directory. The password is not stored on the Cluster.
- **Single Sign-On** - A Single Sign-on (SSO) user account that uses the supported

Identity Provider (IdP) to authenticate to the Cohesity cluster. Cohesity supports major IdP vendors such as Okta, Duo, Ping, and Azure AD.

Protocols and Authentication

The Cohesity cluster can authenticate users defined locally in the Cluster and users defined in Active Directory. NFS uses allowlist-based access control. SMB uses Active Directory for authentication and authorization. Anonymous access, which doesn't require authentication, is not allowed for SMB.

Kerberos Authentication Protocol

Kerberos can be enabled for stronger authentication for SMB by joining the Cohesity cluster to Active Directory. Cohesity supports only Active Directory-based Kerberos authentication.

Client Side Protocols

Cohesity supports the SMB, NFS and S3 protocols.

SMB Protocol Security

Cohesity supports two types of authorization data on a file: Windows-style access control lists (ACLs) and POSIX mode bits (UNIX permissions). ACLs are used when a file is accessed from SMB. POSIX mode bits are used when a file is accessed via NFS.

Access Description	Supported	Notes
Access-control lists (ACLs)	Yes	Windows-style ACLs
Windows-style (NT) credentials for UNIX users	No	Not supported in NFS v3
SMB access of UNIX-created files	Yes	
NFS access of Windows-created files	Yes (with root access)	
Home directory permissions	Yes	Windows-style ACLs

NFS Protocol Security

- NFS access is controlled via IP addresses mentioned in the Allowlist. Client IP addresses that are mentioned in the Allowlist are allowed to mount an NFS share.
- The Allowlist is set at the Cluster level and View level override is available.
- Cohesity supports RPCSEC_GSS protocol for NFS.

Mixed SMB and NFS Environments

Beginning with version 6.0, support for mixed-permission environments is available. The supported scenarios are:

- A UNIX client can store and access a file over NFS.
- A Windows client can store and access a file over SMB.
- A UNIX client can connect by NFS to a file that was stored on the Cluster by a Windows client over SMB and access it based on the permissions.
- A Windows client can connect by SMB to a file that was stored on the Cluster by a UNIX client over NFS and access it based on the permissions.
- Data can be written at the same time to a file via NFS or SMB.

AWS S3

Cohesity supports:

- S3 server over HTTPs.
- Authentication with AWS Signature Version 2 and Version 4.
- AWS ACLs.
- Access controlled via allowed subnets.

Authorization

Authorization is implemented at two levels.

- The operation level, where operations are restricted to authorized users only.
- The object level, where objects are restricted to authorized users only.

The Cohesity cluster can also be added as an Active Directory domain.

Role Based Access Control (RBAC)

A role is a collection of operations, such as creating a View or deleting a Protection Group. A role grants authorization to perform one or more operations. Cohesity provides default roles such as Admin, Operator, Viewer, and Data Security. Customers can create custom roles. A user can have multiple roles but must have at least one role to log in to the Cluster.

A user can be assigned a set of objects on the Cluster. The user is authorized to access only those objects and perform operations on them allowed by the role assigned.

Accounting

Accounting is implemented at two levels.

- The operation level, where each operation is tracked.
- The file level, where each file access is tracked.

Audit Trail

Every user operation that changes the Cluster is audited and the audit log is saved. The log contains the original value and the value after the operation. For a list of logged operations, see Audit Logs. Cohesity stores audit logs on the Cohesity cluster. Users can read them using the standard NFS, SMB and S3 protocols and can export the logs to a syslog server.

Compliance

The following sections explain how Cohesity addresses various guidelines and specifications.

Federal Information Processing Standards (FIPS) Compliance

The Cohesity cluster is FIPS 140-2 certified. The encryption module has been validated on the Cohesity cluster with NIST for FIPS certification. The encryption module always operates in FIPS mode. You can find detailed information in the *Cohesity FIPS 140-2 Non-proprietary Security Policy* on the NIST website.

Common Criteria (EAL2+ Certification)

Common Criteria is an internationally recognized set of guidelines for evaluating IT security products. Common Criteria:

- is a catalog of criteria, and a framework for organizing a subset of the criteria, into security specifications
- is a methodology for evaluating security features
- can be applied to hardware, software, firmware or a combination thereof
- allows vendors to describe a product's security functionality with proof to support their claims

Cohesity realizes the value of being a Common Criteria participant. To view the CC certificate, see MyCC.

DoDIN APL Certification

Cohesity is DoDIN APL certified. The Department of Defense Information Network (DoDIN) Approved Products List (APL) is the single consolidated list of products that have completed Cybersecurity (CS) and Interoperability (IO) certification. DoD customers are allowed to select products for purchase from this list.

The DoDIN APL process is a rigorous security testing, validation, and certification process. This certification's security functional requirements come from a best practices guide called Security Technical Implementation Guides (STIG). It includes application security, OS hardening, web server security, and networking security as well. Cohesity adheres to these highest security standards.

As part of the DoDIN APL requirements, Cohesity provides the following security features:

- Login banner support
- Transfer audit logs to a Syslog server over a secure TLS connection.
- Certificate mapping-based authentication support
- Support to configure the classified or unclassified banner.
- Node-to-node encryption.

WORM (Write Once Read Many SEC 17a-4(f) Certification)

Long-term records retention is mandated by regulations and compliance rules, for example:

- In the financial services industry, SEC Rule 17a-4(f) specifies that "electronic records must be preserved exclusively in a non-rewriteable and non-erasable format."
- Other industries must meet similar requirements when they store mission-critical information.

WORM technology addresses these requirements by providing immutable locking and secure data retention capabilities.

Cohesity implements WORM as *DataLock* Views. DataLock Views are locked (read only) for a user-specified time period. Any user who has access to the View can clone a copy of the View as a DataLock View and set the lock retention time. During that time period, only users with Cohesity Data Security privileges can delete the cloned View, extend the lock period (it cannot be reduced) , add allowlisted IPs, change the QoS and change the description of the View. After the lock period expires, the View can be deleted by any user who has permission to delete it.

A replicated or archived DataLock View retains its DataLock properties. A DataLock View can be cloned to a regular or DataLock View.

General Data Protection Regulation (GDPR)

GDPR is an EU regulation designed to strengthen data protection for residents of the EU. It became effective on May 25th, 2018, and applies to any company controlling or processing Personally Identifiable Information (PII) of EU residents, regardless of the location of the company.

GDPR imposes a broad set of legal, governance and technical requirements on companies processing personal data. A subset of these requirements—those related to data protection and data management—are particularly relevant for storage platforms used to store personal data.

Cohesity has a strong vision to enable our customers to be at the top of their GDPR compliance. The tools and product functionality that Cohesity provides as part of its vision fall into the following categories:

Protect

- Mandate
  - GDPR mandates data protection by design and by default, making privacy by design an express legal requirement.
  - Organizations should also be able to demonstrate transparency when validating compliance.
- Features Offered
  - Encryption architecture, based on strong AES-256-GCM standard and has a FIPS-Certified mode, provides high end-to-end security while allowing optimal use of available resources.
  - Granular control through Role Based Access Control and strong Active Directory integration
  - File or view level WORM provides immutable locking and secure data retention capabilities

Minimize:

- Mandate
  - GDPR mandates data minimization as one of the main tenets to ensure that companies maintain reduced amounts of stored personal data.
  - Here, PII needs to be kept only for the time period relating directly to the original intent of capturing the data.
- Features offered
  - Cohesity architecture inherently minimizes data copies, reduces attack footprint, and tracks copies through centralized data management
  - Granular control over and automation of retention policies enables customers to keep PII only for intended periods
  - Effectively relocate files to on demand to minimize the spread of PII data

Locate:

- Mandate
  - GDPR mandates getting a 360-degree view of what PII is stored by an organization. In addition, the organization should be aware of:
    - What PII they store and why
    - Where does it come from
    - Where the data is stored, and the retention policy on the personal data
    - Who has access to the data
    - And, monitor the movement of the data

- Features Offered
  - Scheduled update to data maps that delineate the following:
    - Location and Movement Tracking (source & destination) of PII
    - Categories of PII stored
    - File containing PII
    - Retention policies of PII
    - Access rights to PII
  - Leverage Cohesity reports to complete Data Protection Impact Assessments (DPIAs), if needed.

Search:

- Mandate
  - GDPR mandates that any EU resident can request visibility of PII from a company.
  - They can also request for the PII to be deleted, modified/corrected, or exported out.
- Features offered
  - Search within unstructured data for multiple categories of PII
  - Input PII patterns and their variations, and file types (txt, doc, pdf, xls, zip, jpeg) to scan using templates
  - Inject custom code to run data processing jobs on stored data using in-place analytics tool
  - Report search results in txt file format or integrate with 3rd party data visualization or analytics tools

Monitor

- Mandate
  - GDPR mandates companies to monitor for and report security breaches.
  - Companies are obliged to communicate breaches to relevant supervisory authorities, and to affected individuals in certain cases.
  - Demonstrating appropriate reporting procedures post-identification of a breach can validate compliance.
- Features offered
  - Prevent portability of PII using Allowlists
  - Receive notifications or warnings when data is tiered, archived, or replicated to a non-Cohesity target
  - Receive a notification when data leaves EU while still on Cohesity
  - Ability to export cluster and system-level audit logs for additional analytics and breach detection

Payment Card Industry Data Security Standard (PCI DSS)

Cohesity adheres to the security benchmarks and requirements that are mandated by the PCI-DSS, including the following:

- authentication and access
- logging and auditing
- browser, plug-in and / or third-party software requirements
- management systems and software
- releases, patching and updates
- granular service controls
- exposed protocols
- monitoring, alerting and call home

Cohesity can provide details on demand.

Secure Government Clouds

Cohesity supports AWS GovCloud and Azure Government Cloud, and is working towards C2S support.

Trade Agreements Act (TAA)

All Cohesity appliances are compliant with the Trade Agreements Act (19 U.S.C. & 2501-2581) and ship from San Jose, CA.

Authorization to Operate (ATO)

Cohesity has the following:

- ATO for highly classified US Department of Defense (DoD) agency networks.
- DoD customers have created STIGs for Cohesity on DoD top-secret networks.
- ATO for highly classified Department of Energy networks.
- ATO for US Air Force networks that support mission-critical programs including the Air Force GPS ground systems.
- Support the US Intelligence Community on many programs.

Secure Data from Ransomware Attacks

Threat Defense Architecture

Cohesity's Threat Defense architecture is designed to help you combat threats against your data in the cluster by enabling you to achieve outcomes related to data protection, compliance, operational resiliency, and defense against sophisticated attacks.

For instance, these outcomes include:

- An ability to rapidly recover from a ransomware attack.
- Resilience of the data against aggressive attacks that actively target the data protection platform.
- An ability to recover from both operational and data center failures.
- Compliance through immutable and tamper-resistant data archives.

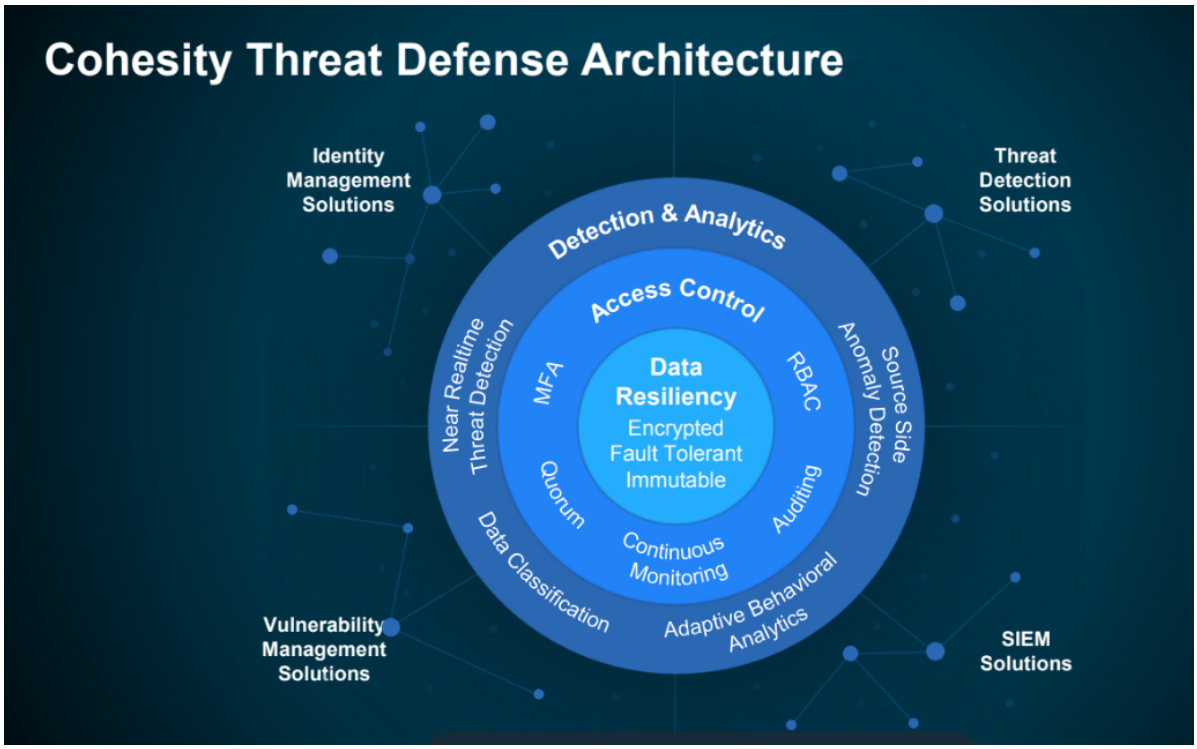

Cohesity Data Cloud

At its core, Threat Defense architecture is about providing a highly resilient platform that ensures confidentiality, availability, and integrity of the data with encryption, fault tolerance, and immutability built into the Cohesity Data Cloud.

The Cohesity Data Cloud is built around a hardened, self-contained Linux-based operating system (CohesitySpanOS) specifically designed to serve as a data management platform.

> **Note:** Cohesity's data protection application does not run on general-purpose operating systems such as Windows.

The SpanOS kernel operates with SELinux extensions enabled in enforcing mode by default. This helps restrict privilege escalation and file modification, and confine access to services to ensure zero-trust security. Cohesity SpanOS is not accessible to users.

Services

Services and subsystems on the Cohesity Data Cloud remain isolated with established APIs and other secure communication mechanisms used for service-to-service communications. Services that are essential to the operation of the cluster are installed and enabled, and other services included in general-purpose Linux distributions are not included on the Cohesity Data Cloud. This helps minimize the attack surface area of the Cohesity Data Cloud.

Cohesity Data Cloud includes an in-built firewall that minimizes the network surface area of the cluster, therefore, only exposing ports and services necessary for cluster operation based on the configured operation.

Cohesity SpanFS

Cohesity SpanFS is an immutable file system. It ensures strict consistency of data across nodes within a cluster to provide data resiliency. Immutability ensures that WORM-protected data cannot be deleted or modified by any process or user including cluster administrators. Only the backup service running on the Cohesity cluster can write to the file system through trusted APIs.

DataLock

DataLock enables your storage to be WORM-compliant (Write-once-read-many) and restricts any modification of backup snapshots. You can refer to Create or Edit a Standard Policy to learn more about creating or editing policy that includes DataLock.

Data Operations

In a Cohesity cluster, data integrity is maintained through multiple integrity checks whenever data is read or written. Data is securely written to the Cohesity Data Cloud. This means data is always written in empty space and never overwritten.

> **Note:** Every read or write operation to internal views is sandboxed.

Encryption

Data at rest and data in-flight encryption using FIPS-validated cryptography are employed with support for external key management.

Air Gap Protection

Cohesity Data Cloud can be configured to achieve air-gap protection that allows you to maintain a physically isolated network. Refer to Cohesity's Solutions for Air Data Protection to learn more about air-gap protection.

Zero Trust Access Control

Data resiliency is further augmented with comprehensive access control capabilities that ensure that entities accessing the data management platform conform to Zero Trust

principles.

Authentication

Cohesity supports a wide range of user authentication options, including a local user store (with in-built support for multi-factor authentication), integration with external identity providers, and certificate-based authentication (including support for smart cards). Besides native, in-built support for multi-factor authentication, Cohesity can also integrate with SAML-based authentication providers.

Access Control

Administration relies on role, object, and tenant-level access controls to ensure that Cohesity users are provided with the least privilege on the system. Access to the shell via SSH is disabled by default and can be controlled with organizational security policies. By default, you have read-only permission. Multi-man approval protocols are required for sensitive operations such as escalating privileges. Cohesity Support cannot access a cluster for troubleshooting unless explicitly granted limited-time permission by a cluster administrator. By default, read-only permission is granted.

Credentials Management

Each locally managed credential on the Cohesity Data Cloud is securely stored as a cryptographic representation of the credential and not in plaintext format. Therefore, you can reset it from factory defaults. In addition, security policy, which is enabled by default, can enforce password length and complexity rules and force password changes regularly to minimize the risk of brute force attacks.

Access to API

User and API access to the Cohesity Data Cloud are strictly controlled. Cohesity conforms with Zero Trust principles. This means that a user can access the Cohesity cluster only if they have the required permissions on the objects and files, and have been authenticated through any authentication protocol in the cluster. You can additionally control access to Cohesity APIs with rotatable API access tokens and user credentials.

Auditing

User activity is recorded in an audit trail that can be automatically shipped to an external log repository.

Analytics and Detection

The Threat Defense architecture incorporates additional layers of protection, including AI-driven detection and analytics. It combines a set of advanced capabilities ranging from threat detection to data anomaly detection.

Threat Detection

The Cohesity Data Cloud detects threats and vulnerabilities in several ways.

- The Cohesity Data Cloud detects anomalies in your cluster and provides a dashboard to monitor it. Refer to Detect Anomalies for more information.
- The Cohesity Data Cloud integrates with anti-malware and next-gen antivirus engines to detect when data managed by the Cohesity may contain or be infected with ransomware.
- The CyberScan app that runs on the Cohesity Data Cloud provides live vulnerability scanning of protected workload snapshots.

The final layer of Cohesity Threat Defense is about ensuring the data management platform can be operated securely. This means integrating and inter-operating with security infrastructure from vendors. The supported security infrastructure ranges from SOAR to identity management and from vulnerability management to key management.

Prepare Your Deployment Against Ransomware Attacks

In a ransomware attack, attackers focus on compromising privileged accounts with weak authentication mechanisms to gain complete access to the cluster and eventually, data.

You can eliminate these risks by preparing your cluster for ransomware attacks. The preparation includes:

- Ensure the Cohesity cluster is upgraded to the latest release and updated with latest patches. You can find the latest version at the Cohesity Downloads portal.
- Change default passwords for each user account such as UI Admin, Support, Intelligent Platform Management Interface (IPMI) admin, and cohesity_console and applying the required password policies. For more information, see Users and Groups and Configure Account Security Settings.
- Maintain separate Active Directory (AD) or SSO protection groups for primary and disaster recovery cluster credentials
- Store privileged credentials securely in credential vaults.
- Enable two-factor authentication (2FA) or multi-factor authentication.
- Disable sudo access for the support user account.
- Ensure users are assigned roles with the least privilege to perform their jobs through role-based access control.
- Enable DataLock on Protection Policies applied to critical data.
- Enable centralized logging and auditing (External Syslog Server).
- Enable offsite backups such as replication or archival in an operationally independent environment.
- Enable firewall policy to restrict network access and manage inbound traffic on Cohesity cluster. Refer to the topic Manage Application-Based Firewall Rules to learn more about firewall rules.

- Protect your Active Directory. Ransomware attacks often target Active Directories to gain control over identity management systems associated with your cluster. You can refer to Active Directory to learn more about prerequisites, best practices, and recovery workflows.

Detect Anomalies

You can connect your Cohesity cluster to Cohesity Data Cloud and keep track of anomalous snapshots in your cluster. Ransomware detection in Helios is based on machine learning and statistical anomaly detection techniques to identify sudden changes in the data across snapshots over time. When these changes are detected, it might be indicative of ransomware compromising the data.

Cohesity Data Cloud constantly monitors the metrics, such as how many modified snapshots, and detects any anomalies. In addition, it uses machine learning algorithms to monitor change rates in the backup data continuously.

Cohesity's ransomware detection techniques can analyze changes in:

1. The volume of data read or written over time – noticing whether the volume of data protected from one protection job run to the next changed dramatically.

2. Analyzing changes in data entropy – very good for detecting when the contents of files may have changed dramatically due to encryption or compression

3. Changes in the nature of the file system – identifying when a large number of files has been added, deleted, or modified

If the rate is out of the normal range — based on daily and historical rates — Cohesity flags it as a potential ransomware attack. Refer to Detect Ransomware Attacks to learn more about how Cohesity Data Cloud detects anomalies.

Respond to Ransomware Attacks

In response to ransomware attacks, it's important to immediately take the required measures and limit the potential damage to your organization's digital assets. You can integrate Helios Security Orchestration, Automation, and Response (SOAR) platforms to receive alerts when ransomware has attacked your cluster.

You can refer to the View the Alert section in the Anti-ransomware topic to learn about how to check the alerts in your cluster. As a first step to limit the exposure of your cluster to the ransomware attack, you should:

1. Reset Passwords. For more information, see Users and Groups.

   > **Note:** For 6.6.0d and earlier, and 6.5.1 releases, you have to change the password of Linux users.

2. Disable accounts that are not configured with multi-factor authentication.
3. Inform your Technical Account Manager and the Cohesity account team.

4. Enable Legal Hold for critical workloads. It ensures that snapshots of protection groups will not be deleted until the Legal Hold is removed or disabled. For more information, see Add or Remove a Legal Hold to a Snapshot.
5. Stop the cluster through a CLI command cluster stop.

After limiting exposure to the ransomware, you have to start recovery from a clean snapshot, recover active directory objects, and initiate mass restoration of your critical workloads such as VMware virtual machines and network-attached storage backups. For more information, see the Cohesity white paper: Cohesity Ransomware Protection — Prepare and Recover.

Key Management

A Key Management Service (KMS) helps you manage the keys that encrypt data and secures it from unauthorized access. A Cohesity cluster has a built-in key manager or can be configured to use external key managers - KMIP-compliant or Amazon Web Services (AWS) KMS - to manage encryption keys.

Cohesity uses two types of keys for encryption.

- Data Encryption Keys (DEKs) - these keys are used to encrypt the data.
- Key Encryption Keys (KEKs) - these keys are used to encrypt the DEKs.

Data on the Cohesity cluster is encrypted using a FIPS-validated implementation of the Advanced Encryption Standard (AES) in Galois/Counter Mode using a 256-bit DEK.

A unique DEK is created for the following entities in your Cohesity cluster.

- Encryption-enabled Storage Domains
- External targets, such as
    - CloudArchive
    - CloudArchive Direct

Regardless of the entity type, DEKs are encrypted using a KEK and persisted on the cluster in an encrypted format as Encrypted DEKs (EDEKs). DEKs are never stored or exposed in a plain text format.

If the external targets are CloudArchive and CloudArchive Direct, then the EDEKs are also stored alongside the data in the cloud.

The encryption and decryption process is performed within the Cohesity cluster and is transparent to all inbound or outbound protocols and applications such as backup, archiving, and file services.

Security Measures

Cohesity prioritizes security and access control to ensure that customer data remains protected at all times. A key concern is whether Cohesity administrators can access customer Key Management Systems (KMS) or encryption keys, potentially exposing

protected data. These Security measures clarify why such access is not a risk within the Cohesity Cloud Services (CCS) Backup-as-a-Service (BaaS) solution.

- **Restricted Production Environment Access**: Cohesity enforces strict access controls, access to the production environment is limited to authorized personnel.
- **Encryption and Decryption Controls**: Cohesity does not provide any mechanism to decrypt customer data outside of normal product operations. Encryption and decryption occur within the secure system framework, adhering to strict security protocols.
- **Key Management and Secure Storage**: Key Encryption Keys (KEKs) are securely stored in a Key Management System (KMS), purpose-built for secure key storage and management. This ensures that keys are not exposed or accessible beyond their intended use.
- **Customer-Controlled Key Management**: Customers can configure an external KMS, ensuring full control over their encryption keys. In this case, Cohesity does not retain direct access to KEKs but only invokes them to decrypt Data Encryption Keys (DEKs) when required for standard operations.
- **Risk Mitigation**: In the unlikely event of a security breach, Customers utilizing an external KMS can revoke Cohesity's access to the KEK, preventing further decryption operations.

Internal Key Management Service

Cohesity provides an internal key management service (KMS) to generate Data Encryption Keys (DEKs) and Key Encryption Keys (KEKs) on your Cohesity cluster.

- Cohesity stores data in Storage Domains.
- For each Storage Domain, the internal KMS generates a new KEK and DEK.
- At the end of this window, if additional data needs to be written to the Storage Domain, a new DEK is created.

When a new Storage Domain is added to the cluster, the internal KMS creates a new DEK and encrypts it with a KEK. Cohesity securely stores this DEK internally within the cluster's key store as an Encrypted DEK (EDEK).

To keep the encryption keys safe, Cohesity restricts the following actions:

- Exporting keys from your cluster
- Exposing keys in plain text format through any API or user interface

Key Rotation Policy

By default, the KEK rotation interval is 90 days, and you can configure this interval as needed. For more information, see Set Key Rotation Policy.

DEKs are created and managed as described in Internal Key Management Service.

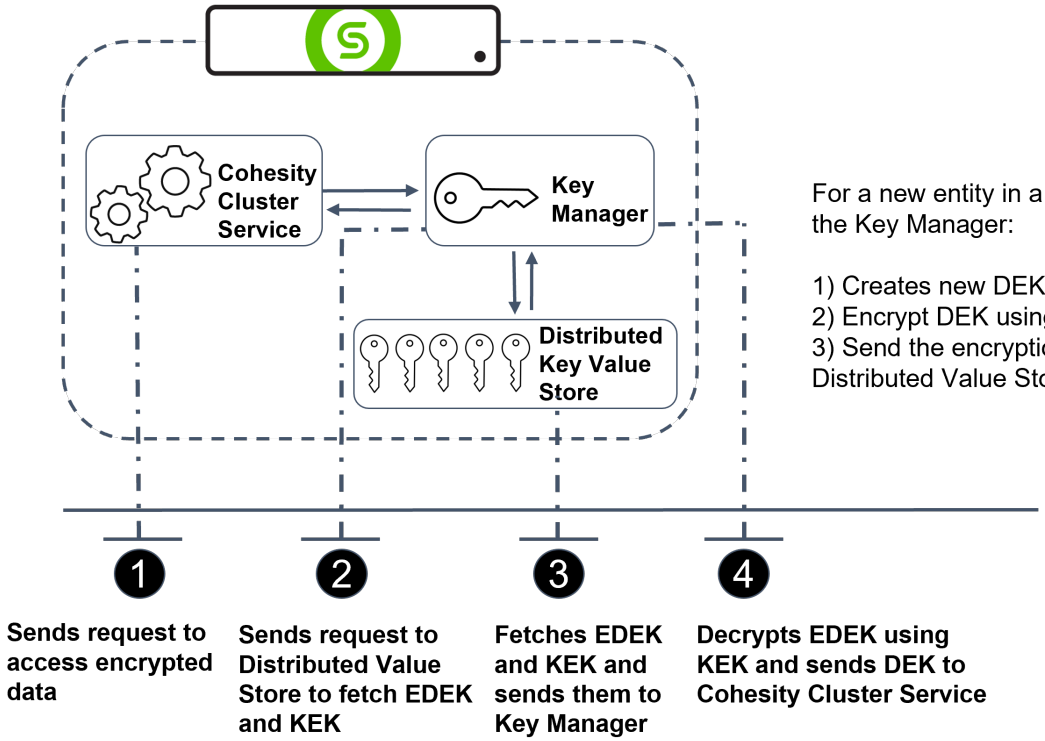

External Key Management Service

Cohesity supports Amazon Web Services-based KMS and Key Management Interoperability Protocol (KMIP) compliant KMS as external key management services. KMIP is used to facilitate communication between the Cohesity cluster and the key server of the external KMS. You can register multiple key management services.

When the Cohesity cluster is configured to use an external KMS, KEKs are generated and stored in the external KMS, though a KEK ID, a reference to the KEK, is maintained on the cluster. The cluster never persists a copy of the KEK, though the Cohesity cluster will persist the KEK ID as an identifier to fetch the KEK from the external KMS periodically.

> **Note:** The KEK ID does not contain any key material.

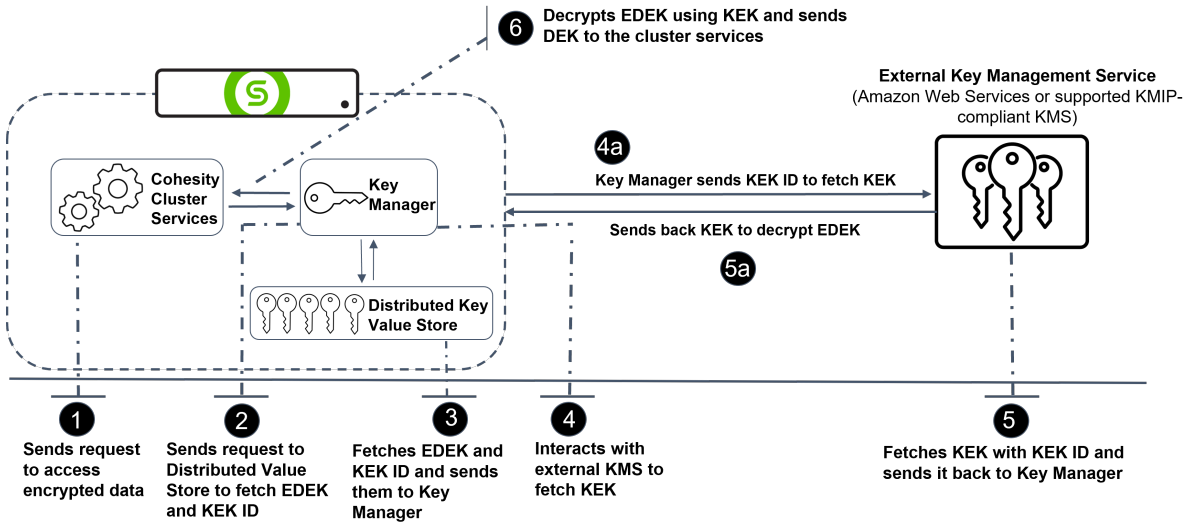

KMIP-based External KMS

When a KMIP-based external KMS is enabled, the Cohesity cluster will store both the EDEK and KEK ID locally. The external KMS will manage the KEK.

A Cohesity cluster interacts with an external KMS when a service on the cluster requests access to the encrypted data. After receiving this request, the Cohesity cluster sends a request to the external KMS for the KEK referenced using the KEK ID. When the external KMS responds with KEK, the Cohesity cluster uses it to decrypt the EDEK and derive the DEK. This DEK is sent to the service that requires the key to decrypt or encrypt the data for any read or write operations.

Cohesity allows caching of keys and it is implemented within various cluster services to ensure cluster throughput and minimize load on the external KMS. This allows the DEK to be cached in-memory on the cluster node for four hours before a call to external KMS must be made again to fetch the KEK and decrypt the EDEK.

> **Note:** If the external KMS is unavailable or unable to deliver the KEKs due to any reason, the data in the Storage Domains remains in an encrypted state and inaccessible.

By default, the key rotation interval is 90 days, though this is configurable.

Create a KMS Cluster

You can add multiple KMS and create a KMS cluster to ensure high availability of keys. A KMS cluster can include up to four key management services.

In a KMS cluster, every KMS node should share the same CA certificate and each Cohesity cluster node's certificate is shared internally. This means the additional KMS that is added must be connected to the same Cohesity cluster.

Follow the steps below to add a new KMS.

1. Go to **Settings** -> **Summary** and select **Key Management System** tab.
2. Select the edit icon against an existing KMS.

   The **Edit Key** dialog is displayed.
3. Enter a **Key Name**.
4. Under **Key Configuration**, select **KMIP Compliant** as the **Key Type**.

   KMIP Compliant: Select this option to add an external KMS. For more information, see External Key Management System.
5. Enter the following details.
   1. **Protocol Version**: Enter the KMIP version. If the version is KMIP 1.2, then enter KMIP1_2.
   2. **Server Address**: Enter the IP address of KMS. The IP address of each KMS node must belong to the same KMS cluster.

      > **Important:** The IP Address of the active KMS should be first on the list followed by the IP addresses of the secondary key management systems.

   3. **Port**: Enter the port number of KMS.
6. Under **Upload Certificates**, you have to upload the following keys and certificates.
   1. **Client Certificate**
   2. **Client Key**
   3. **CA Certificate**
7. Optionally, you can select the storage domain and external targets that must use this key for encryption. Ensure that encryption is enabled on the storage domain.
8. Select **Create**.

Validated External KMIP-based KMS

Cohesity supports KMIP versions 1.1 to 1.4, and the following KMIP-based KMS systems have been validated by the Cohesity cluster.

KMS	Documentation Reference
Entrust KeyControl	Cohesity and Entrust KeyControl Integration Guide
SafeNet KeySecure	Integration Guide - Cohesity and SafeNet AT KeySecure for Government

KMS	Documentation Reference
Vormetric Data Security Manager	Use Vormetric DSM to Manage Cohesity Encryption Keys
Fortanix KMS	Use Fortanix Self-Defending KMS to Manage Cohesity Encryption Keys
CipherTrust Manager	CipherTrust Manager Documentation Portal
IBM Security Guardium Key Lifecycle Manager	IBM SGKLM Documentation
Hashicorp Enterprise Vault	Manage Cohesity Encryption Keys Using HashiCorp Vault 1.4

AWS KMS

When a AWS KMS is enabled, the Cohesity cluster will store the KEK ID locally while the EDEK will be stored both on the cluster as well as alongside the data (in the case of CloudArchive and CloudArchive Direct).

A Cohesity cluster interacts with AWS KMS when a service on the cluster requests access to the encrypted data. After receiving this request, the Cohesity cluster will send a request to the AWS KMS for the Customer Managed Key (CMK). When you're using AWS KMS, the CMK acts as the KEK. When AWS KMS responds with the CMK, the Cohesity cluster uses it to decrypt EDEK and derive the DEK. This DEK is sent to the service that requires the key to decrypt or encrypt the data for any read or write operations.

If you are using AWS KMS as an external KMS, then

- DEKs are rotated every four hours
- KEKs, which is an AWS CMK, is rotated once every year by AWS

Caching of keys is implemented within various cluster services to ensure cluster throughput and minimize load on the external KMS. This allows the DEK to be cached in-memory on the cluster node for a limited period of time before a call to AWS KMS must be made again to fetch the KEK and decrypt the EDEK.

> **Note:** You can configure a HTTPS proxy on the Cohesity cluster and use it for communication with AWS KMS. This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Configure AWS KMS

Perform the following steps in the Cohesity cluster to configure AWS KMS.

1. Navigate to **Settings** > **Summary**.
2. Select **Key Management System**.
3. Select .

   The **New Key** pop-up dialog is displayed.
4. Enter the **Key Name** in the **New Key** dialog box.
5. Go to **Key Configuration** > **AWS**.
6. Enter the following details.
   1. **Customer Managed Key ARN**: You can choose to have an Amazon CMK ID or Amazon Resource Name (ARN). You should not enter key aliases in this field.
   2. **Access Key ID**: Amazon access key ID.
   3. **Secret Access Key**: AWS secret access key.
7. Enable or disable **SSL verification** of the connection to the AWS KMS server. By default, **SSL verification** is enabled.
8. Optionally, you can select the storage domain and external targets that must use this key for encryption. Ensure that encryption is enabled on the storage domain.
9. Select **Create**.

You have configured AWS KMS as your cluster's external key management service.

Azure KMS

When Azure KMS is enabled, the Cohesity cluster stores the Key Encryption Key (KEK) ID locally. The Encrypted DEK (EDEK) is stored both on the cluster and alongside the data (for services like CloudArchive and CloudArchive Direct).

The Cohesity cluster interacts with Azure KMS when a service requests access to encrypted data. Upon receiving the request, the cluster uses the Azure C++ SDK to authenticate via Managed Identity and fetch the private key from Azure Key Vault. This private key acts as the KEK. The cluster then uses it to decrypt the EDEK and derive the DEK. The DEK is passed to the requesting service to perform encryption or decryption operations.

If you are using Azure KMS as an external KMS:

- DEKs are generated locally on the cluster and rotated periodically (similar to internal KMS behavior).
- KEKs (RSA-HSM keys in Azure Key Vault) are rotated every 365 days for Cohesity-managed keys.
- Key caching is implemented across cluster services to optimize performance and reduce load on Azure KMS. DEKs are cached in memory for a limited time before the cluster re-authenticates and re-fetches the KEK to decrypt the EDEK.

You can now configure an HTTP or HTTPS proxy for both new and existing Azure KMS configurations on a Cohesity cluster. This allows Azure KMS operations to run through the

web proxy already set up on the cluster. Contact Cohesity Support to configure proxy for Azure KMS.

Manage Key Management Service

After you have configured a key management service (KMS) for your Cohesity cluster, you can do the following to manage KMS.

- Restart the Cohesity Keychain Service: The keychain service creates and manages encryption keys in a Cohesity cluster. Restarting it is required after you have updated the KMS configuration.
- Set Key Rotation Policy: You can set the lifetime of encryption keys on the Cohesity cluster. By default, it is set to 90 days.
- Delete an External Key Management System: You can delete an external KMS that is configured to manage encryption keys on the Cohesity cluster.
- Add a Secondary External KMS: If you're using an external key management system, then you can add more than one KMS to your Cohesity cluster to ensure high availability of keys.

Restart the Cohesity Keychain Service

If you update the KMS configuration on the Cohesity cluster, you must restart the Cohesity keychain service on the Cohesity cluster for the configuration changes to take effect. If you configure KMS for the first time, you do not need to restart the keychain service.

To restart the keychain service on the Cohesity cluster:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.
2. Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command.
3. Provide a Cohesity cluster username and password.

```
./iris_cli -server 172.16.3.101 -username=admin
```

4. Run the following command to restart the Cohesity keychain service:

```
cluster restart service-names="keychain"
```

Example:

```
admin> cluster restart service-names="keychain"

Success: Restarting the cluster services [keychain] ...
```

Set Key Rotation Policy

After the Cohesity cluster is created, the key rotation policy, by default, is set to 90 days. This means that the Cohesity cluster uses new encryption keys every 90 days. After the time period has been reached, all the old encryption keys used by the Cohesity cluster are

replaced by a new set of keys and the data on the Cohesity cluster remains as it was originally encrypted.

You can change the period of time that new keys are generated by the Cohesity cluster using the CLI command listed below. Specify the time period in seconds.

To change the default key rotation period:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Change the Key Rotation period. Specify the time period in seconds. For example to change the key rotation period to 30 days, enter the following CLI command:

   ```
   admin>cluster edit enc-key-rotation-period-secs=2592000
   ```

Delete an External Key Management System

You can delete an external KMS from the **Key Management System** page. Ensure that the external KMS is not used by a storage domain or external target.

> **Note:** If the external KMS is used by a storage domain or an external target, the delete icon does not appear.

To delete an external KMS, complete the following steps:

1. Go to **Settings** > **Summary**.
2. Click the **Key Management System** tab.
3. Hover over the key that you want to delete and click the delete icon.
4. Click **Delete**.

Add a Secondary External KMS

You can add a new KMS and create a KMS cluster to ensure high availability of keys. A KMS cluster can include up to four key management services.

In a KMS cluster, every KMS node should share the same CA certificate and each Cohesity Platform node's certificate is shared internally. This means the additional KMS that is added must be connected to the same Cohesity cluster.

Follow the steps below to add a new KMS.

1. Go to **Settings** -> **Summary** and select the **Key Management System** tab.
2. Select the edit icon against an existing KMS.

   The **Edit Key** dialog is displayed.
3. Enter a **Key Name**.
4. Under **Key Configuration**, select **KMIP Compliant** as the **Key Type**.

   KMIP Compliant: Select this option to add an external KMS. For more information, see External Key Management System.
5. Enter the following details.
   1. **Protocol Version**: Enter the KMIP version. For example, if the version is KMIP 1.2, then enter KMIP1_2.
   2. **Server Address**: Enter the IP address of KMS. The IP address of each KMS node must belong to the same KMS cluster.

      > **Important:** The IP Address of the active KMS should be first on the list followed by the IP addresses of the secondary key management systems.

   3. **Port**: Enter the port number of KMS.
6. Under **Upload Certificates**, you have to upload the following keys and certificates.
   1. **Client Certificate**
   2. **Client Key**
   3. **CA Certificate**
7. Optionally, you can select the storage domain and external targets that must use this key for encryption. Ensure that encryption is enabled on the storage domain.
8. Select **Create**.

You've successfully added a new KMS.

To migrate from Internal KMS to External KMS, see Migrate from Internal KMS to External KMS.

MegaFile

Storage of large volumes of data such as large file systems can be challenging because they are often stored in a single storage container, delaying backup and recovery. The MegaFile feature alleviates this delay by dividing large files into smaller parts and distributing the smaller parts across multiple nodes of a Cohesity cluster, allowing parallel and hence faster backup and recovery.

The following example helps to understand the impact of MegaFile. In a four-node cluster, the data ingest of a 2 TB VMDK is split into four parts, each of which is 512 GB, reducing the ingest time by a factor of four. In an eight-node cluster, the data ingest of a 2 TB VMDK is split into eight parts, each of which is 256 GB, reducing the ingest time by a factor of eight. As the number of nodes in a Cohesity Cluster grows, the time to ingest decreases significantly.

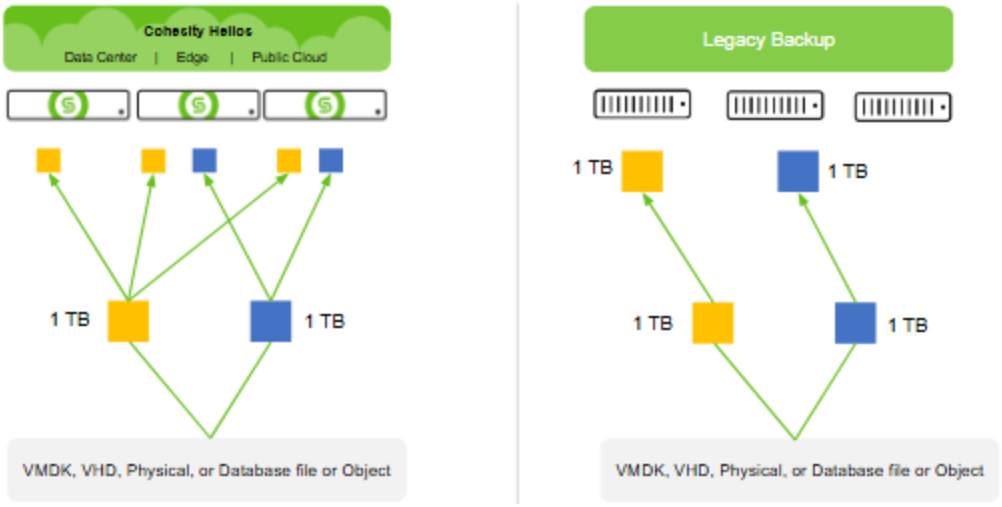

> **Note:** The size of the smaller parts is determined by MegaFile and varies depending on many factors such as the size of the larger file, the number of nodes in the Cohesity cluster, and so on.

> **Note:** The MegaFile feature is built in a Cohesity Cluster and is enabled by default.

Example using MegaFile

In the example below, two sub-tasks are created by MegaFile and it splits the VMDK to a sub-task for every 256 GB offset:

```
<date, time>    Number of sub-tasks: 2, total amount of data to be read:
260.00 GB
```

```
<date, time>    Setting up VM files in preparation to copy the snapshot
```

```
<date, time>    Setting up files completed
```

```
<date, time>    Querying disk [key: 2000, path: [HDD-3TB(0)] name-blank/-
name-blank.vmdk] from 0 with change id: *
```

```
<date, time>    Requesting permit for sub-task 0
```

```
<date, time>    Requesting permit for sub-task 1
```

```
<date, time>    Permit granted to sub-task 0
```

```
<date, time>    Permit granted to sub-task 1
```

Example without using MegaFile

In the example below, only a single sub-task is created:

```
<date, time>  Number of sub-tasks: 1, total amount of data to be read:
260.00 GB
```

```
<date, time>  Setting up VM files in preparation to copy the snapshot
```

```
<date, time>  Setting up files completed
```

```
<date, time>  Querying disk [key: 2000, path: [HDD-3TB(0)] name-blank/name-
blank.vmdk] from 0 with change id: *
```

```
<date, time>  Requesting permit for sub-task 0
```

```
<date, time>  Permit granted to sub-task 0
```

MegaFile in Backup and Recovery

The following sections discuss the various backup and recovery flows that employ MegaFile:

Source	Supported Backup Workflows	Supported Recovery Workflows
Microsoft Hyper-V	• Full backups use MegaFile by default. • Incremental backups only if the previous backup used MegaFile.	• Linux and Windows VM recoveries in Cohesity DMaaS portfolios. • File recovery to Linux VMs on On-Prem.
Nutanix AHV	• Full backups use MegaFile by default. • Incremental backups only if the previous backup used MegaFile.	N/A
VMware	• Full backups use MegaFile by default. • Incremental backups use MegaFile only if the previous backup used MegaFile.	File recoveries streamed over a Cohesity agent on VMware Cloud use MegaFile. **Note:** MegaFile is not supported if a file recovery on VMware Cloud is based on VMware tools.

Source	Supported Backup Workflows	Supported Recovery Workflows
Physical Server	All file-based AND block (volume) based backups.	All file and folder recoveries.
SQL Server	All SQL Server backups.	All SQL Server recoveries.
Pure Storage (SAN)	• File backups use MegaFile by default. • Incremental backups use MegaFile only if the previous backup used MegaFile.	All Pure Storage recoveries.
Active Directory	Volume and File backups.	N/A
Exchange	Volume and File backups.	N/A
MongoDB physical	Full backups use MegaFile by default.	N/A

Data Retention for Audits

In planning for audits, keep in mind that the Cohesity cluster retains metadata for Protection Groups and Objects (VMs, physical servers, databases, etc.) differently, as described below.

Data Retention for Backup

Backup auditors need to verify that your organization's assets are backed up at regular intervals and that backups are available for recovery. For example, you might be asked to show evidence of successful backups for the past three months or more.

Protection Group Metadata

Protection Group metadata requires minimal overhead in the Cluster and is retained forever for reporting and compliance purposes, unless the Protection Group is deleted. There is no set retention period for this information. The information includes:

- the Protection Group name
- the protection run start and end date and time
- the names of the backed up Objects
- the total number of bytes read and written
- whether the protection run succeeded or failed

Note that the protection run is considered successful only if *all* its Objects are backed up successfully. For example, if a Protection Group is configured to back up 100 VMs, but one of the VMs gets deleted in vCenter, the protection run status indicates failure. The more VMs included in a Protection Group, the greater the probability of the protection run reporting a failure. As a result, it may be desirable to capture more granular success/failure metadata at the Object level.

Object Metadata

Object metadata is **not** retained forever in the Cluster. A single protection run can back up hundreds or thousands of Objects and the metadata can grow quickly and require overhead. Therefore the Object metadata is retained only for the retention period specified in the Protection Group's Policy. After that period, the Object metadata is permanently deleted from the Cluster and is no longer available. You can guard against gaps in the Object backup history by collecting the Object information on a regular basis and storing it outside the Cluster.

Getting Audit Data

You can get Protection Group and Object information by running the following reports in the Cohesity cluster:

- Backup Summary report
- Protection Details Per Object report
- Protection Summary by Object Type report
- Failed Objects report

You can customize the reports with filters including a date range, Protection Group names and object names. You can run the reports as needed, on a recurring schedule, download them as CSV files or email them to backup administrators to review and keep for auditors.

Data Retention for Recovery

Recovery metadata requires minimal overhead in the cluster and is retained forever for reporting and compliance purposes. There is no set retention period for this information.

Recovery Metadata Retention

The following information is retained:

- Recovery Task
- Start Time
- Status
- Duration

Object Metadata Retention

The following information is retained:

- Object name
- Recovered From
- Recovery Point
- Options

For recovery tasks performed in the last month, the following information is retained:

- Object name
- Recovered From
- Recovery Point
- Recovery Task
- Start Time
- Duration
- Options

Getting Audit Data

You can get recovery task and object information for VMs by running the Clone and Recover VM Tasks Summary report in the Cohesity cluster.

Data Retention for Clone

Clone metadata requires minimal overhead in the cluster and is retained forever for reporting and compliance purposes, unless the clone is deleted. There is no set retention period for this information.

Clone Metadata Retention

The following information is retained:

- Task Name
- Objects Cloned
- Start Time
- Task Duration
- Status

Object Metadata Retention

The following information is retained:

- Clone Options
- Object
- Snapshot Cloned

- Teardown Tries
- Status

Getting Audit Data

You can get recovery task and object information for VMs by running the Clone and Recover VM Tasks Summary report in the Cohesity cluster.

Sending Audit Data Outside the Cluster

Consider the following options for retaining backup information outside the Cluster:

- Email Cohesity reports on a recurring schedule to an external repository such as a SharePoint site, an email inbox or a log collection server such as Splunk. For a list of reports and how to email and schedule them, see Reports. For information on how to set up your repository to receive and store emails, see the documentation from the repository vendor.
- Write a script that calls the Cohesity REST API operations to generate and save reports to a location outside of the Cluster.

Dynamic Disk Support

This topic describes what dynamic disks are and captures details of the supported and unsupported workflows of the Dynamic Disk use case in ESXi VMs.

What are Dynamic Disks?

Dynamic disks provide features such as the ability to create volumes that span multiple disks (spanned and striped volumes) and the ability to create fault-tolerant volumes (mirrored and RAID-5 volumes). Like basic disks, dynamic disks can use the MBR or GPT partition styles on systems that support both styles. All volumes on dynamic disks are known as dynamic volumes. Dynamic disks offer greater flexibility for volume management because they use a database to track information about dynamic volumes on the disk and about other dynamic disks in the computer. Because each dynamic disk in a computer stores a replica of the dynamic disk database, for example, a corrupted dynamic disk database can repair one dynamic disk by using the database on another dynamic disk. The location of the database is determined by the partition style of the disk. On MBR partitions, the database is contained in the last 1 megabyte (MB) of the disk. On GPT partitions, the database is contained in a 1-MB reserved (hidden) partition.

Dynamic disks are a separate form of volume management that allows volumes to have non-contiguous extents on one or more physical disks. Dynamic disks and volumes rely on the Logical Disk Manager (LDM) and Virtual Disk Service (VDS) and their associated features. These features enable you to perform tasks such as converting basic disks into dynamic disks and creating fault-tolerant volumes. To encourage the use of dynamic disks,

multi-partition volume support was removed from basic disks and is now exclusively supported on dynamic disks.

The following operations can be performed on dynamic disks:

- Create and delete simple, spanned, striped, mirrored, and RAID-5 volumes.
- Extend a simple or spanned volume.
- Remove a mirror from a mirrored volume or break the mirrored volume into two volumes.
- Repair mirrored or RAID-5 volumes.
- Reactivate a missing or offline disk.

Are Dynamic Disks Recommended?

Cohesity supports dynamic disks while recommending against them as they complicate the disaster recovery process. Cohesity also advises against using dynamic disks as a system disk.

Physical Windows workflows back up at the volume level and create one VHD / VHDX file with one volume. Thus, there is no issue with supporting dynamic disks if the VM is treated as a physical server.

Supported Workflows in ESXi VMs

The following table lists supported and unsupported workflows of Dynamic Disk use cases in ESXi VMs.

Workflow	Backup	Indexing	File Restore to Source	VM Restore to Source	Recommended Approach
VM with simple Dynamic Disks	Supported	Supported	Supported	Supported	Treat VM as a Physical Server
VM with spanned Dynamic Disks	Supported	Supported	Download file is supported • Search to download • Browse to download	Supported	Treat VM as a Physical Server

Workflow	Backup	Indexing	File Restore to Source	VM Restore to Source	Recommended Approach
VM with SQL Dynamic Disks	Supported	Supported	Supported	VM restore is supported DB restore to same VM not supported. DB restore to alternate VM supported.	Treat VM as a Physical Server

Limitations

- Recovery to an alternate location is supported with the following restrictions:
  - The volume should be a simple volume. Dynamic disks may be created unintentionally after some P2V workflows.
  - The alternate machine should ideally not have any other dynamic disks. If there are any other disks already, then manual intervention is required to clean up the missing dynamic disks that result from the disks attached during restores.
  - Quiesced VM snapshots fail with missing dynamic disks. Thus, if the alternate machine has dynamic disks and is part of a job that takes quiesced snapshots of the VM, then the cleanup becomes mandatory.
- Indexing (and hence search and download) will work for Simple, Spanned, Striped and Mirrored volumes. It will not work for RAID-5 volumes. So, File Level Recoveries are not supported for RAID-5 volumes.
- Instant Volume mounts on Hyper-V are not supported for Dynamic Disks.
- VM with SQL DB will have the same behaviors as above.

Chapter 4: Cluster Administration

From the (**Settings > Summary**) page of the Cohesity cluster, you can view overall statistics about the Cohesity cluster, configure or upgrade the Cohesity cluster and manage the Support Channel. See the following topics for details.

- Initial Cluster Setup 236
- Configure Settings 236
- Customize Preferences 268
- Cohesity Cluster Services 271
- Software Update 282
- Power Off and Power On a Cohesity Cluster 301
- Search the Cluster and Take Action 315
- Manage Storage Domains 317
- Manage Nodes 342
- Use Cohesity Data Cloud (Self-managed) in DoD Mode 386
- About Node-to-Node Encryption 395
- Quorum Protected Operations 398
- Antivirus Software to Protect Cluster Data 398
- Cohesity Licensing 403
- Manage Apps 414
- Networking Workflows 416
- Manage Firewall Rules and Ports 487
- Access Management 585
- Web Server Certificates 700
- Support Toolbox Overview 701
- Available Service Health Checks 705

Groot uptime check 710
Available Third-party Health Checks 719
Available Backup Health Check 720
Available Vendor Health Checks 720
Available Software Health Checks 721
Available Networking Health Checks 728
Available Archival Health Checks 737
Available Archival Health Checks 752
Available Performance Health Checks 753
Available Upgrade Health Checks 754
Available Vendor Health Checks 764
Environment Setup 774
Access Cohesity Host OS 803
Hardware Replacement 804

Initial Cluster Setup

The initial cluster Setup Guides, Quick Starts and Worksheet PDFs for all Cohesity platforms are available at Hardware Universe.

Configure Settings

In addition to the settings initially entered while setting up the Cohesity cluster, you must provide the configuration information listed below if it is not already specified.

To continue the configuration of Cohesity cluster settings:

1. Navigate to **Settings > Summary**. The Summary page displays the following details:

Field	Description
Cluster Name	Displays the name of your cluster.
Cluster ID	Displays the unique ID of your cluster.

Field	Description
Creation Date	Displays when you created the cluster.
Software	Displays the cluster's Cohesity release version.
Hardware	Displays which hardware the cluster is on. For example, Virtual ROBO.
Software Encryption	Indicates whether encryption is enabled at the cluster level. > **Note:** Encryption can be enabled at the cluster level only when creating the cluster and cannot be disabled or modified later. If Encryption is enabled at the Cluster level, all Storage Domains are encrypted. However, encryption can be enabled at the Storage Domain level after creating the cluster. For enhanced security, the 7.0.1 or later clusters use Galois/Counter Mode (GCM) encryption. Once you upgrade the Cohesity cluster to version 7.0.1, all data newly generated within the Cohesity cluster is automatically encrypted using GCM. Data that existed prior to the upgrade will remain encrypted using the previous method (CBC). You can apply GCM at the cluster level during cluster creation or at the storage domain level after the cluster is created.
Hardware Encryption	Indicates whether Encryption is enabled at the cluster hardware.
Storage Domain	Displays the total number of Storage Domains you have created and the encryption status. For details, see Manage Storage Domains.
Nodes	Displays the total number of nodes in the cluster. For details, see Manage Nodes.
Support Channel	Allows you to enable the Support Channel and copy the Support Channel token. For details, see Manage the Support Channel.
Storage Capacity for Metadata	Displays the total storage capacity available for the cluster's metadata.
Storage Used for Metadata	Displays the percentage of capacity utilized by the cluster.
Failure Domain	Displays the Failure Domain. For example, Node.

2. On the **Summary** tab, click **Configure**, and provide the following configuration information:

Required Settings

The following table lists the mandatory settings that are required for your Cohesity cluster to be up and running.

Settings	Description
Cluster Subnet Gateway	Specify the IP address of the Subnet Gateway for the Data (Primary) Network. You can either provide IPv4 or IPv6 addresses to configure the Cohesity cluster subnet gateway.
Domain Names	Enter the domain name(s) for the Cohesity cluster.

Optional settings

The following tables list the optional settings that are available for advanced Cohesity cluster configuration.

Web Proxy settings

Use an HTTP proxy server if the Cohesity cluster does not have direct internet connectivity to connect to Cohesity Data Cloud, the support channel, or to download upgrade software.

Settings	Description
Configure Web Proxy settings	Enable the **Configure Web Proxy settings** option, and enter the hostname or IP address and port number for the server. If the server requires a username and password, toggle on **Proxy Requires Password** and enter the username and password. You can either provide IPv4 or IPv6 addresses to configure the IP addresses or hostnames of the Web proxy server.

To configure Web Proxy using the CLI commands, see Use a Proxy Server.

> **Note:** To configure an HTTP proxy server for outbound internet connectivity from a Cohesity cluster to external targets, Office 365, AWS, and Azure Cloud Connector (for CloudSpin and native snapshot), see this KB article.

Security settings

Add an extra layer of security to your Cohesity cluster by changing the default password and enabling two-factor authentication.

Field	Description	Additional Notes
Enable Two-Factor Authentication	Enable this option for the users (belonging to the same AD domain to which Cohesity has joined) to sign in to the Cohesity cluster using a digital certificate instead of a password (Local users will continue to sign in with a password).	Changing this setting automatically restarts the web server.
Encryption	Encrypts metadata and data on the Storage Domain on the Cohesity cluster. **Note:** The metadata encryption impacts Cohesity cluster performance. For example, the read/write operations could delay the backup/archival/replication run times.	This option is displayed only if the **Encryption** setting is turned on during Cohesity cluster creation. **Note:** Once enabled during the Cohesity cluster creation, you cannot disable this option.

Field	Description	Additional Notes
Key Management Service Type	Select the required Key Management Service Type. The available options in the drop-down are: • Internal Key Management Service • KMIP Compliant • AWS > **Note:** By default, Cohesity uses Internal Key Management Type. For **KMIP Compliant**, you need to select the key from the **KMIP Compliant Key** drop-down. For **AWS**, you need to select the key from the **AWS Key** drop-down.	At the cluster-level encryption, DEK (Data Encryption Key) rotation is not supported for AWS KMS. If multitenancy is enabled on the Cohesity cluster, Cohesity uses the same key to encrypt all tenants' metadata.

Storage Domain Configuration

Activate cloud tiering and set default disk-level fault tolerance for Storage Domains.

Settings	Description	Additional Notes
Cloud Tier	Set **Local Storage Utilization Threshold** to activate cloud tiering for configured storage domains.	This is a global setting for all the Storage Domains created on the Cohesity cluster. Setting up thresholds for cloud tiering at individual Storage Domains will take precedence over the global setting (overrides the global settings).
Custom Failure Domain	Specify a failure domain at the global cluster level: **Node**–Sets node as the failure domain and protects the Cohesity cluster against node failures. This is the default setting.	This is a global setting that sets the failure domain to configure fault tolerance for all the storage domains created on the Cohesity cluster.

Settings	Description	Additional Notes
Failures Tolerated	Select the number of node failures that can be tolerated: • **1 Failure**–Can tolerate one random node failure at any given point in time. The usable metadata capacity will be 50%. A minimum of three nodes in a Cohesity cluster is required to tolerate one node failure. • **2 Failures**–Can tolerate two random node failures in a Cohesity cluster simultaneously. The usable metadata capacity will be 33%. A minimum of five nodes in a Cohesity cluster is required to tolerate two node failures. For more details, see About Fault Tolerance.	• Applied to metadata and used as the default setting for Storage Domains. • This setting is not supported for single-node Virtual Edition. • A 2-node failure has a higher fault tolerance than a 1-node failure.

Email Configuration

Configure the connection to the SMTP Server and verify the settings when you save the Cohesity cluster configuration.

Settings	Description	Additional Notes
Enable SMTP Server	Enable the **Enable SMTP Server** option to configure the connection to the SMTP Server. Specify the following SMTP settings: • **SMTP Server**: Specify the name of an SMTP server that is used to send emails when Warning or Critical Alerts are generated by the Cohesity cluster. Cohesity recommends specifying an SMTP mail server. • **Port**: Specify the port number used to access the SMTP Server. • **SMTP Server uses SSL/TLS without STARTTLS (typically for port 465)**: Select this option if the SMTP Server uses SSL/TLS without STARTTLS. Typically, SSL/TLS without STARTTLS uses port 465. • **SMTP Username**: Specify the name of the account used to authenticate with the SMTP Server. • **Change SMTP Password**: Specify the password of the account used to authenticate with the SMTP Server. • **Sender Email Address**: The SMTP server uses the email address you enter in this field to send alert notification emails. **Note:** From release 6.6.0d_u3 CPD patch onwards, the email address you enter in this field will take precedence over the email address specified for the local System Admin account (admin) of the Cohesity cluster.	The SMTP Server is used to send emails when Warning or Critical Alerts are generated by the Cohesity cluster. Cohesity recommends that you configure the SMTP Server.
Test Email on Save	Enable the **Test Email on Save** option to verify the SMTP settings when you save the Cohesity cluster configuration. In the **Test Email Recipient** field, enter an email address to receive the test email.	

Multitenancy Settings

Configure multitenancy for the Cohesity cluster.

Settings	Description	Additional Notes
Enable Organizations	Enable the **Enable Organizations** option to configure multitenancy for the Cohesity cluster. Optionally, toggle on **Allow multiple organizations to use one Storage Domain** to reduce storage consumption in the Cohesity cluster by providing Cohesity deduplication across all the organizations using the Storage Domain. Organizations will be isolated from other organizations, whether this option is on or off.	Enabling multitenancy for a Cohesity cluster cannot be undone. You cannot revert the Cohesity cluster to a single tenancy state.

Apps Management

Manage applications from the Cohesity App MarketPlace in the Cohesity cluster.

Settings	Description

Enable Apps Management	Enable the **Enable Apps Management** option to manage applications from the Cohesity App MarketPlace in the Cohesity cluster. Provide the following details:
	• **Subnet**: Provide an unused subnet in a private address space between /12 (maximum size) and /24 (minimum size). The subnet will be used to create an overlay network on which to run the apps. Make sure the subnet is unused to avoid routing conflicts. The Internet Assigned Numbers Authority (IANA) reserves the following IP address blocks for use as private IP addresses:
	  • 10.0.0.0 to 10.255.255.255
	  • 172.16.0.0 to 172.31.255.255
	  • 192.168.0.0 to 192.168.255.255
	• **Allow external connectivity from apps**: Toggle on to allow the apps network access beyond the Cohesity cluster nodes. Not all apps require external connectivity. For example, a backup app that needs to access a database will require external connectivity, but the Cohesity Insight app, which searches Views on the Cluster, does not require external connectivity.

Statistics

Send anonymized data about product usage.

Setting	Description
Share Product Usage Statistics	Enable this option to send anonymized data about product usage. For more information, see this Google Analytics topic.

Product Documentation

Use the documentation help that is included directly on the Cohesity cluster without internet access.

Settings	Description	Additional Notes
Documentation	Select the **No internet access? Access Help directly from the Cluster** option to use the documentation help that is included directly on the Cohesity cluster without internet access.	The Cohesity cluster pages have context-sensitive links to the online help available from the Cohesity Documentation site or offline help from the Cohesity cluster. Click the icon in the upper right corner of the Cohesity cluster to open the help topic for the page you are currently viewing. Then, if prompted, enter your MyCohesity portal credentials.

Login Banner

Displays a login banner message on the Login screen.

Settings	Description
Enable Login Banner	Toggle on to see the configured login banner message on the Login screen. In the displayed text field, configure the login banner message.

Time Settings

Change the time zone for reports and audit logs.

Settings	Description
Time Zone	Click **Change Time Zone** and select a time zone from the drop-down to change the time zone for reports and audit logs.

Throttling Background System Activities

The activities, such as garbage collection, system checks, post-process storage optimization, and other system activities, run in the background on the Cohesity cluster. These system activities utilize the Cohesity cluster resources and impact the workloads, such as file shares and backups.

Schedule the throttling of the system activities to reduce the utilization of the Cohesity cluster resources.

Settings	Description
Background System Activity	Enable this option to schedule the throttling of the system activities running in the background. **Time Zone Settings** To change the time zone for throttling schedules, click **Change Time Zone** and select a time zone from the drop-down menu. By default, your current time zone is displayed. **Add Throttling Schedules** > **Important:** Cohesity recommends that you schedule the throttling in such a way that, on average, the background system activities utilize at least 50% of Cohesity resources in a week. To add schedules: 1. In the **Days** column, select one or more weekdays from the drop-down menu. You can also select **Everyday** to schedule throttling to run daily. 2. In the **Time** column, enter the time duration for throttling. 3. In the **Activity Level** column, select a level of throttling from the drop-down menu. • **No restriction**: The background system activities will run without any restriction. • **Moderate**: The background activities will utilize 66% of the Cohesity resources. This is the default option. • **Low**: The background activities will utilize 33% of the Cohesity resources. • **No activity**: Other than critical processes, all other background system activities will be completely stopped. The background activities will not utilize any of the Cohesity resources. 4. Click **Add**. 5. Optional. To configure multiple schedules for throttling the background system activities, repeat steps a to d.

Settings	Description
	Note: • If there is an overlap between the schedules added, then the latest added schedule will take precedence over other schedules. For example, you have added a schedule (say S1) for running the background activities on **every Monday** between **9.00 am to 5.00 pm** with **Activity Level** set to **Low**. Later, you added another schedule (say S2) for running the background activities every day between **1.00 am to 3.00** pm with **Activity Level** set to **Moderate**. In this scenario, on Monday from 1.00 am to 3.00 pm, there is an overlap between these two schedules. In such cases, the recently added schedule, S2, will take precedence during the overlap time. • Cohesity will perform throttling of the system activities running in the background only for the days and duration for which you have configured the throttling. Saving the Cohesity cluster configuration with this option disabled will clear all the existing throttling schedules. When you save the Cohesity cluster configuration, the following error message appears if the Cohesity resources utilization by the background system activities is less than 50%: `Error occurred while configuring background activity schedule. Background activity has been excessively throttled. Please relax the throttling conditions.` **Delete an Existing Throttling Schedule** To delete a throttling schedule, click the delete icon against the schedule.

3. Click **Save**.

Manage the Support Channel

The Cohesity Support Channel is a secure, simple, and effective way for Cohesity employees to provide on-demand assistance to customers. When the channel is enabled, a qualified Cohesity employee can log in to the Cohesity cluster to troubleshoot issues in a remote session. The following diagram illustrates how the Support Channel works.

The Support Channel server is located in the public cloud and has a secure login mechanism. The server is configured with a public key and accepts connections from the Cohesity cluster. The Cohesity cluster selects a master Node, which then establishes an SSH reverse tunnel with the Support Channel server on a TCP port. (For information on ports,

see Firewall Ports for Support Channel.) This setup ensures that all communication is encrypted between the Support Team and the Cohesity cluster.

The server configuration is limited to only certain Cohesity employees, and there are strict rules for user registration that are checked regularly.

Only a limited set of Cohesity employees can configure the Support Channel server, and a strict user registration policy is enforced and audited regularly.

Firewall Ports for Support Channel

The primary node in the Cohesity cluster initiates an SSH reverse tunnel with the Support Channel server on TCP port 443. To use the Support Channel server server, you need to configure the ports listed in the Cohesity Support Channel section in the Manage Firewall Ports topic before you enable the Support Channel server.

> **Note:** Make sure your network team has unblocked and is not monitoring the Cohesity support channel sites or ports for seamless connectivity and effective communication.

Enable Support Channels

Cohesity recommends that you enable the Support Channel when you need assistance from Cohesity employees. You can enable the Support Channel using the Cohesity Dashboard or the Cohesity CLI. When you enable the Support Channel, you must specify how long you want the Support Channel enabled. When the period you specified expires, the Support Channel turns off. You can also manually disable the Support Channel before the specified period expires.

> **Note:** Support Channels are disabled by default in Cohesity clusters.

Enable or Disable the Support Channel

To enable or disable the Support Channel:

1. Navigate to **Settings > Summary**.
2. Next to the **Support Channel** at the bottom of the page, click the edit icon ( ).
3. Select one of the following from the list:
   - **Off**—Disables the channel.
   - **Temporarily On**—Lets you enable the channel temporarily. Enter a number (1 to 30) and then select hours or days from the drop-down. If the support channel is in use and about to expire within 7 days, the support channel duration will be automatically extended for the next 7 days.

> **Note:** You can toggle off the automatic extension if it is not required.

4. Click **Save**.

You can also manage the Support Channel with the Cohesity CLI. See the `iris_cli support support-channel` command.

Support Channel Tokens

Once the Support Channel is enabled, a Support Channel token is autogenerated. Cohesity employees require the Support Channel token to remotely log in to your Cohesity cluster using SSH. Copy the Support Channel Token and share it with the Cohesity employee. Support Channel tokens provide a more secure and hassle-free method than sharing the password with Cohesity employees.

Support Channel tokens are encrypted and are longer than the tokens in the previous Cohesity versions. Because encryption appends a random number, an encrypted token may appear different every time it is fetched. However, the token is still valid until the Support Channel expires or is disabled.

Support Channel tokens expire when the set time duration for the Support Channel expires, as a security enhancement. A new token is generated when you re-enable the Support Channel. Cohesity employees logging into Cohesity clusters require the new Support Channel token every time the Support Channel is enabled.

> **Note:** If you extend the Support Channel duration before it expires, the existing Support Channel token is retained. However, if you manually disable the Support Channel before its session expires, the Support Channel token also expires.

Copy the Support Channel Token

To copy the Support Channel token:

1. Navigate to **Settings > Summary**.
2. From the **Support Channel token** field, copy the Support Channel token using the copy icon.

Or, from the Cohesity CLI, run the following `iris_cli` command to copy the Support Channel token:

```
iris_cli user linux-bash-access
```

The following is a sample output of the command:

```
admin@127.0.0.1> user linux-bash-access

supportUserToken             :
```

```
JDEkTUhjQ0FRRUVJQ2VOLzR4QTBwUllJNWJEK1VyVEZGR2l0NXRkODRmZWwzTVB0WlgzaEZvS29
Bb0dDQ3FHU000OQpBd0VIb1VRRFFnQUVDWWJESDFoRVU2U0gvdkZ6a3pQcndaZU41aGNLZWdrdD
AvNzNpZktxL2NXODloWkJqbXgzCkJaZlFF
SEZmTzhjTEt1MllkbXRjNWxVbkh2cE0raUNBbGc9PQ==
```

Support UI User

When you enable a Support Channel, a Support UI user role (cohesity_ui_support_<clusterid>) is created by default and displayed along with the other user roles under the **Users & Groups** tab on the **Access Management** page. Cohesity employees remotely logging into the Cohesity cluster use this user role. By default, this user role will have viewer permissions and a secure password. The user role is automatically disabled and deleted if the Support Channel expires or if the Support Channel is manually disabled.

Migrate from Internal KMS to External KMS

This section describes the procedure to migrate the storage domain or metadata encryption from Internal KMS to External KMS.

> **Note:** Once you migrate from Internal KMS to External KMS, the change is irreversible. Migration from External back to Internal KMS is not supported.

Migrate Storage Domain Encryption from Internal KMS to External KMS

If encryption is already enabled and the storage domain uses **Internal KMS**, you can migrate from Internal KMS to AWS KMS or any KMIP-compliant KMS.

To migrate the storage domain encryption from Internal KMS, complete the following steps:

1. Go to **Settings > Summary**.
2. Click the **Storage Domains** tab.
3. In the **Storage Domains** page, hover over the storage domain, click ⋮ and select **Edit**.
4. In the **Encryption** section, from the **Key Management Service Type** drop-down list, select **AWS** or **KMIP Compliant**.
5. Click **Update Storage Domain**.

Migrate Metadata Encryption from Internal KMS to External KMS

If you enable encryption when creating a Cohesity cluster, by default, Cohesity uses **Internal KMS** for cluster-level encryption. After the Cohesity cluster is created, you can migrate from Internal KMS to AWS KMS or any KMIP-compliant KMS.

> **Note:** If encryption was not enabled when you created the Cohesity cluster, the **Encryption** option does not appear.

To migrate metadata encryption from Internal KMS, complete the following steps:

1. Go to **Settings > Summary**.
2. Click **Configure**.
3. In the **Encryption** drop-down list, select **AWS** or **KMIP Compliant**.
4. Click **Save**.

Set Key Rotation Policy for Encryption

After the Cohesity cluster is created, the key rotation policy, by default, is set 90 days. This means that the Cohesity cluster automatically creates and uses new encryption keys every 90 days. After the time period has been reached, all the old encryption keys used by the Cohesity cluster are replaced by a new set of keys and the data on the Cohesity cluster remains as it was originally encrypted. You can change the period of time that new keys are generated by the Cohesity cluster using the CLI command listed below. Specify the time period in seconds.

To change the default key rotation period:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Change the Key Rotation period. Specify the time period in seconds. For example to change the key rotation period to 30 days, enter the following CLI command:

   ```
   admin>cluster edit enc-key-rotation-period-secs=2592000
   ```

Update SSL Certificates Using Cohesity CA or Customer CA certificate for Agent to Cluster Communication

Cohesity software packages ship with an auto-generated self-signed certificate. You must accept the self-signed certificate to access the Cohesity clusters. This topic describes the following ways to update or manage the cluster certificate from the Cohesity cluster (User Interface):

- Load a Certificate Authority (CA) signed certificate.
- Refresh the factory shipped (Cohesity signed) certificate if it expires.
- Configure CA integration with Venafi.

Important Notes on SSL Certificates

- Cohesity does not recommend using self-signed certificates.
- Cohesity cluster supports X.509 certificates in PEM format only.
- The certificate must have the "Common Name" field set to the server's hostname (for example, a VIP's DNS name). The Common Name must be the same as the name you plan to use to access the Cohesity cluster.
- Private Key Files must not be encrypted.
- If the certificate admin provides a single file, ensure to separate the file to obtain the certificate and key files.
- By default, OCSP validation is disabled. Contact your Cohesity account team to enable the feature.
- Cohesity does not support ECDSA certificates.

Examples

The following is a sample cert.pem file.

-----BEGIN CERTIFICATE-----
JHFDRTCCAi2gAwIBAgIRAMnovKoL9P3hveY2r9Y8wXIwDQYJKoZIhvcNAQELBQAw
NTEVMBMGA1UEChMMQ29oZXNpdHkgSW5jMRwwGgYDVQQDExNDb2hlc2l0eSBJbmMg
U2VydmVyMB4XDTE4MTIwMzE2MzM0MFoXDTE5MTIwMzE2MzM0MFowNTEVMBMGA1UE
ChMMQ29oZXNpdHkgSW5jMRwwGgYDVQQDExNDb2hlc2l0eSBJbmMgU2VydmVyMIIB
IjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAvrE4iQ4BROiVytcMvaBXu2Nm
vBpjjrVdolQ83pdufZhD4vnL22Dnu0U2EWx0YrdaiQX3c8iEC2u951PInFs0NXAW
Y+KGqtRrg17aEJyHVqqEy5NZ0mxPMpmDUs0kNUzjpj6jX521PytjVx3y+/qchuu1
GM+qDWrpLO2s6O5lnkE1sxkR/MK14f4mXUYCmywaevQEr0/0lIjyqDmqEZg1+NQn
XBUgRw5c9Zvod/chpyZQN/seuM7oXRRcgfnEtd8Fm1ztCDy3Ui3eZs1X2d31D4y4
zzBBNdawEv25yTAOcHMvBeLLWwtKnm0RFLDF3SxEByKyKhziT1imOJYv2RaTpQID
AQABo1AwTjAOBgNVHQ8BAf8EBAMCBaAwEwYDVR0lBAwwCgYIKwYBBQUHAwEwDAYD
VR0TAQH/BAIwADAZBgNVHREEEjAQgg4qLnJ0cHNyZS5sb2NhbDANBgkqhkiG9w0B
AQsFAAOCAQEAMqZ8rYzDXoZErRSL/xXhqcMwM+/Mm7TYymrF2ieYIpSWVhi0NRzf
2VlPvlHA+iNkScbV9lB+0hWblz5khK3nTbXmjnL65DgpJrms91LkQbGAh8Tp4ILf
fLdE9/j01uXs5uv1gZxss/lw+a3NW/4jQHtMcuWdYtJqcuGBJxktWGnjcsJWe9Hq
GrLlsuIYsHHPuqNin2HvxdqHCZbk2wumAmv47WGC/rzoCQ/83Sp74GuyK3W1z76x
fkpI9dVKUaox8L81NV8cyneWQe81ma9uoROJuBs6Wj/1qb5seVrbEooC87xtN8ns
EJ8A22yecDHtEJZVfowllVENqqa6mJLl1A==
-----END CERTIFICATE-----

The following is a sample key.pem file.

-----BEGIN RSA PRIVATE KEY-----
JHFEpQIBAAKCAQEAvrE4iQ4BROiVytcMvaBXu2NmvBpjjrVdolQ83pdufZhD4vnL
22Dnu0U2EWx0YrdaiQX3c8iEC2u951PInFs0NXAWY+KGqtRrg17aEJyHVqqEy5NZ
0mxPMpmDUs0kNUzjpj6jX521PytjVx3y+/qchuu1GM+qDWrpLO2s6O5lnkE1sxkR
/MK14f4mXUYCmywaevQEr0/0lIjyqDmqEZg1+NQnXBUgRw5c9Zvod/chpyZQN/se
uM7oXRRcgfnEtd8Fm1ztCDy3Ui3eZs1X2d31D4y4zzBBNdawEv25yTAOcHMvBeLL
WwtKnm0RFLDF3SxEByKyKhziT1imOJYv2RaTpQIDAQABAoIBAQCsL8g18WByAXZo
jyBhBazSduIoTcfB78CZuGiyMZjCbcA27ooGEI2OiuNUflLlZFyqx6l2cgPRcDss
63mEvdH5wrlk/gHNtTtTh01unHiUaCq6y9XabJRiwz1SggwE2OXikkLhLr2fSaVM
360tnFDEf4ep3LHXpcbhG8f7GHDHoNZrs0ZyVXqbafsQ15CY0qukekOUWF4i/MSF
xCRmTTgn3X09gQ1uCrSqT27s6G1NBt5rqeOIS+KjSJfPwdeij9h/fnzazv/NRuIt
/fA1E/x+/fmxaWn/XyLvZbtpAa/WzRyknDs7DWBCA5PYNIzBHIFbW/SPviOlhBSU
QPVnQHMtAoGBAPjmlYPrL5BHSfPXAQlIk7UIZNC8dWmtHZWXKhEHLffhjKBQemTU
8ERBD0LX9Cj/XpIjCgYNPR6SzSP40UY4GSZiVx/g5rRiXXanuZvY1Va8y1ZvbIi+
dO3ctmzIstM41IBjDWEXJu9yhAgTcpNp2B8Q+Xx17X4MMRvxTBrFi1T/AoGBAMQh
nK8gFr+gu1eDObupk75qXKzIyJQ1KrEfeOjMHRCL4BwN+mbv1cM8f32WHdOE3URL
Num0ztq65NjaB68Z0qV6TnSDNT5hFDNPuaczQOArYwA4PMPGgzgN7mpKQ7NQTv7H
5SQMcEv2xBrMdJEEuGONAMrRdv522fUIwBWRj6NbAoGAbkNwTx2cV5VNTKmd+ufd
D6DlUJd8iFE5vvBVoW+FzL9t1gNxBURUIFiWEy2qK5Nd6O1+3kl1Hxs3p8ztoObv
NVLxd9zhN14wO4V4M3Nup1tB1UvNUBLXqBxT8INEvU/6y9mVrsJb7E3V8eHzqxyD
steJPrrbeDCcfN17sox6+i8CgYEAjjrpyz30AUdoVFA/f8277KxpRIvemUvtB8GR
BxFuSl+FSBMC8b8fP5QRyFsiig+h/jKlQdny4EAEURXb2KlbTtFhL7Eu3Wu9IHDC
a3QJqgcMpk4iEr9UzUOdD8uHSPQGNVF05C15Cj2LfPuj1L6dkOPiHTLFnkdGTCHq
WjnS9l0CgYEAzbutugqBCIXxflR27uWg7mb74qhEPEM0SgNN+hMLFVBg1d8/7xtI
KS4XmRqB0JQP4/wicWNYCSiEBk+2tPwUr+7Jed/Oiz/U2TJXWdq/j+yqE1oZHeOU
B4csWtfc7rD1N2RfQax1iojN87He8nIL/YrT6jNFeRMW0awLHMYDQHI=
-----END RSA PRIVATE KEY-----

Configuring Cohesity Cluster with a CA Signed Certificate

A Certificate Authority (CA) signed certificate is a digital certificate issued by a trusted organization called Certificate Authority. It authenticates the identity of the entity (certificate holder) before digitally signing the certificate, thereby assuring users that they are interacting with a legitimate and secure entity.

You can replace the Cohesity cluster's self-signed certificate with a Certificate Authority (CA) signed certificate. Optionally you can use the Subject Alternative Name (SAN) field to specify additional host names (such as sites, IP addresses, and common names) to be protected by a single SSL certificate, such as a Multi-Domain or Extend Validation Multi-Domain Certificate.

This procedure applies to configuring the Cohesity cluster using SHA256 and SHA384 cryptographic hash algorithms. When configuring, update the default_md parameter in step 3 accordingly. To configure a Cohesity cluster with a CA certificate using SHA256 or SHA384:

1. Copy the following template to a text editor, replace the sample values in red color with actual values, and save the file in your local directory.
   Sample

```
#REQ FILE: cert_req.conf
#---------------------------
[ req ]
default_bits    = 2048
default_md      = sha384 (or sha256)
prompt          = no
encrypt_key     = no
distinguished_name = req_distinguished_name
req_extensions     = req_ext

[ req_distinguished_name ]
countryName         = "US"
stateOrProvinceName = "CA"
localityName        = "San Jose"
organizationName    = "Data Protection"
organizationalUnitName = "example Organization Name"
commonName          = "cluster1.example-domain.com"

[ req_ext ]
subjectAltName = @alt_names

[alt_names]
DNS.1 = cluster1.example-domain.com
DNS.2 = cluster1-vip.example-domain.com
IP.1 = 10.1.1.1
IP.2 = 10.1.1.3
IP.3 = 10.1.1.5
#---------------------------
```

2. Run the following command to the SCP file saved in the local directory to support user:

```
scp "/path_to_file/filename" user@ip_address:"/Path to destination/"
```

   Example

```
scp cert_req.conf support@111.1.1.1:/var/tmp/
```

3. Get the IP address for any Cohesity Cluster node (You can get the list of IP addresses by selecting Settings > Summary and then opening the Nodes tab). Use the IP

address to log in to that node using a support user account that has the admin role.

Command

```
ssh support@<node IP address>
```

Example

```
ssh support@111.1.1.1
```

When prompted, enter the password of the support user account.

```
[support@111.1.1.1's password:
```

Once your password is successfully authenticated, the prompt changes to the Secure Shell prompt.

```
[support@restricted-8888222f444 ~]\>
```

4. Generate the private key and certificate signing request files using cert_req.conf file:

```
openssl req -out cert_req.csr -newkey rsa:2048 -nodes -keyout cert_
key.key -config <path to file>/cert_req.conf
```

Example

```
openssl req -out cert_req.csr -newkey rsa:2048 -nodes -keyout cert_
key.key -config /var/tmp/cert_req.conf
```

> **Note:** Private key must not be encrypted.

5. Convert the cert_key.key to a PEM private key using the following command:

```
openssl rsa -in cert_key.key -out /var/tmp/key.pem
```

6. Make sure all users (owners, groups, others) have read permissions for the private key (key.pem) file:

To check the current permissions settings, run the following command:

```
ls -ltr <key.pem>
```

Run the following command to set the required permissions for all users:

```
chmod 600 <key.pem>
```

Sample Output:

```
[support@restricted-a121 ~]\> chmod 600 key.pem
```

```
[support@restricted-a121 ~]\> ls -ltr key.pem
```

```
-rw------. 1 support support 1704 Jun 13 08:23 key.pem
```

7. Run the following command to verify that the CSR file includes the Subject Alternative Name.

```
openssl req -text -noout -verify -in cert_req.csr
```

Sample Output

```
verify OK
Certificate Request:
Data:
Version: 0 (0x0)
Subject: C=XX, ST=PROVINCE, L=CITY, O=EXAMPLE, CN=HOST-A.EXAMPLE.COM
Subject Public Key Info:
Public Key Algorithm: rsaEncryption
Public-Key: (2048 bit)
Modulus:
       00:e0:76:f7:7d:d7:42:b2:46:35:6d:67:a6:f0:10:
       b6:36:54:44:1f:04:60:23:91:7c:4e:0e:c4:34:e8:
       b8:a0:6e:1e:3a:09:cf:3e:26:c8:22:06:2a:f8:6a:
       22:fb:c0:9d:cd:03:ab:6a:5d:7e:34:4a:69:db:fb:
       07:55:b5:c1:df:a1:e4:2a:c2:7e:e4:cc:5f:28:04:
       fe:79:1c:c8:0f:e3:a4:fa:f6:b2:7a:89:6f:84:9a:
       c5:e9:9e:d9:8b:6a:3f:39:ed:7a:d2:40:c4:85:ed:
       ec:45:95:a9:d1:bf:51:3e:ed:f7:2b:fa:36:ed:77:
       19:fd:84:ba:24:33:92:87:1e:4d:18:11:61:ad:de:
       88:7a:77:e9:f8:9b:99:4e:9f:a1:52:4a:77:20:a7:
       70:0d:7d:89:8b:e5:d1:f9:f3:b0:0a:f4:d9:dc:12:
       87:57:8d:84:0b:34:bc:4b:93:30:9a:4f:e8:4e:3b:
       3c:44:95:e8:bb:63:90:b4:e3:a9:d3:d8:46:49:02:
       98:f1:63:88:23:9d:bb:36:02:2a:0b:7c:3e:18:cb:
       84:bb:8d:59:5e:da:87:31:ab:01:e6:61:29:fe:51:
       9c:9c:62:17:8b:e8:9c:53:09:75:03:68:7c:e9:94:
       86:7c:61:8b:c2:53:63:97:5e:78:c9:93:c3:d8:15:
       8d:f3
Exponent: 65537 (0x10001)
Attributes:
Requested Extensions:
X509v3 Subject Alternative Name:
DNS:HOST-A.EXAMPLE.COM, DNS:HOST-A, IP Address:10.0.10.1
Signature Algorithm: sha256WithRSAEncryption
   ac:fb:6c:1c:4c:52:94:7c:6d:99:85:21:a1:c4:c7:eb:97:8e:
   e0:bf:ee:d6:c8:31:b2:6d:ba:38:c5:15:48:97:8d:e0:2d:4d:
   da:52:d1:44:8f:86:3e:73:91:9b:85:b9:b9:2b:77:f2:98:cc:
   d2:a5:b6:6f:65:7f:fe:dd:d8:f9:76:3e:66:5c:01:af:d7:13:
   31:a6:51:44:ca:46:a4:c9:fe:c1:b8:c8:fb:9c:fe:fb:8d:e3:
```

```
a8:53:70:38:b2:63:10:ad:5d:83:93:38:d4:38:eb:e3:d0:bd:
5d:a9:06:47:d1:e6:87:1c:29:ac:33:fb:4e:75:99:70:df:71:
de:f2:84:7b:15:ce:a0:c3:bc:3e:06:9b:82:4b:07:ca:f0:53:
fd:10:82:b5:66:1b:af:ac:eb:bd:4b:51:f1:c1:6f:1e:f4:8d:
87:4d:d6:b0:c1:90:0f:4d:56:52:0f:41:5e:ac:8f:f4:cd:95:
db:d4:3c:14:b2:e3:7e:1e:95:94:27:96:42:af:4b:3a:e6:ae:
3d:40:27:63:bc:7e:cf:ba:7f:74:9d:72:30:ac:c8:77:c4:ee:
45:0f:a9:eb:57:ab:3a:5e:a5:3c:34:ce:1f:e1:4f:6c:63:88:
dd:e3:23:00:85:35:5e:c3:af:14:d0:26:2e:21:2e:7f:bc:7b:
68:10:36:bf
```

8. Use WinSCP or SCP command to download the `cert_req.csr` file from the Cohesity node to your local computer and upload it to the Certificate Authority (CA) for signature. Obtain the certificate in PEM format `(cert.pem)`.

   > **Important:** Submit a certificate request to the CA using the contents of the CSR file, following the CA's enrollment process.

Note:

- Cohesity cluster only supports X.509 certificates in PEM format only with Base64 encoded.
- When you are getting the cert_req.csr certificate file signed to create the cert.pem file, ensure not to choose chain certificates.
- If the certificate is in DER format, run the following command to convert it into PEM format:

```
openssl x509 -in cert.cer -outform PEM -out cert.pem
```

- If the server certificate is signed by an intermediate CA then the `cert.pem` file must have the server certificate first and then the intermediate certificate.

  Example

```
-----BEGIN CERTIFICATE-----
Server Certificate
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
Intermediate CA Certificate
-----END CERTIFICATE-----
```

9. After you obtain the certificate (cert.pem), copy it to the Cohesity cluster so you can update the cluster with the new certificate. Copy the file to a node on the cluster using the scp command on Linux systems as follows:

```
scp cert.pem support@<Node IP Address>:/var/tmp
```

Note: If you created the key pair on your enterprise certificate management tool, you must import both the files (signed certificate and key) on the Cohesity cluster. Copy the file to a node on the cluster using the scp command on Linux systems as follows:

```
scp cert.pem key.pem support@<Node IP Address>:/var/tmp
```

10. Make sure all users (owners, groups, others) have read permissions for the self-signed certificate (cert.pem) file:

    To check the current permissions settings, run the following command:

    ```
    ls -ltr <cert.pem>
    ```

    Run the following command to set the required permissions for all users:

    ```
    chmod 600 <cert.pem>
    ```

11. Replace the Cohesity cluster's current SSL certificate with the new certificate you created in the previous steps:

    a. Start the Cohesity CLI using the following command:

    ```
    support@restricted-test-cluster-005056ba314a-node-1 ~]\> iris_cli
    ```

    When prompted, enter the Username and password you use to log into the Cohesity Cluster's User Interface. Once the password is successfully authenticated, the Cohesity CLI console opens.

    b. Replace the certificate by running the following command:

    ```
    cluster update-ssl-certificate ssl-certificate=<absolute path of the cert.pem file> ssl-cert-private-key=<absolute path of the key.pem file>
    ```

    Example

    ```
    admin@127.0.0.1>cluster update-ssl-certificate ssl-certificate=/var/tmp/cert.pem ssl-cert-private-key=/var/tmp/key.pem
    ```

    c. Restart the UI and REST API Services using the following command:

    ```
    admin@198.51.100.12>cluster restart service-names=iris
    ```

    > **Note:** It is recommended to wait a minute before proceeding with the next command.

    d. Restart the I/O Operations service using the following command:

    ```
    admin@198.51.100.12>cluster restart service-names=bridge
    ```

Configuring Cohesity Cluster with a Cohesity CA Issued Certificate

Use this procedure if you do not have an internal CA and do not use SSL certificates issued by a public CA. In such cases, you can use the Cohesity cluster CA to generate and manage SSL certificates.

Before you Begin

You need Host Shell Access for the Cohesity cluster to run the last four commands in this procedure. For details, see Using the Secure Shell.

1. Get the IP address for any Cohesity Cluster node (You can get the list of IP addresses by selecting Settings > Summary and then opening the Nodes tab). Use the IP address to log in to that node using a support user account that has the admin role.

   Command

   ```
   ssh support@<node IP address>
   ```

   Example

   ```
   ssh support@111.1.1.1
   ```

   When prompted, enter the password of the support user account.

   ```
   [support@111.1.1.1's password:
   ```

   Once your password is successfully authenticated, the prompt changes to the Secure Shell prompt.

   ```
   [support@restricted-8888222f444 ~]\>
   ```

2. Run the following command to start Cohesity CLI.

   Command

   ```
   $ iris_cli
   ```

   When prompted, enter the Username and password you use to log into the Cohesity Cluster's User Interface. Once the password is successfully authenticated, the Cohesity CLI console opens.

   Example

   The following is a example of Cohesity CLI.

   iris_cli, Version: 6.8.2_release-20240317_97f56d9a

   Host: cork-clu-lts-005056ba9163-node-2, 127.0.0.1

   Cluster Server: 127.0.0.1:443

   admin@127.0.0.1>

3. Run the following command to generate Cohesity CA certificate.

   Command

   ```
   cert-manager new-cert city=<city-where-the-company-located-or-cluster-installed> country-code=<two-letter-code-of-the-country-defined-by-the-ISO-standard> organization-unit=<organization-unit-that-owns-the-node> organization=<name-of-the-company> state=<state-where-the-city-is-located> san-list=<fqdn-ipaddresses-of-all-interfaces-and-VLANs-separated-by-comma>" common-name="<name-of-cluster-to-which-certificate-is-to-be-assigned>" output-dir=/home/support/
   ```

> **Note:** Include all FQDNs and VIP addresses from all interface groups and VLANs, and ensure that at least one FQDN or wildcard domain name is specified for the san-list parameter.

Example

If cluster FQDN is xyz.cohesity.com and if it resolves to IP addresses 10.1.1.112 and 10.1.1.115, then:

```
cert-manager new-cert city="San Jose" country-code="US" organization-unit="IT" organization="Cohesity Inc" state="California" san-list="xyz.cohesity.com,10.1.1.112,10.1.1.115" common-name="xyz.cohesity.com" output-dir=/home/support/
```

Certificate and Private Key will be written to following path:

- /home/support/new-cert/privateKey.pem
- /home/support/new-cert/certificate.pem

4. Run the following command to replace the certificate:

   Command

   ```
   cluster update-ssl-certificate ssl-certificate=<absolute path of the cert.pem file> ssl-cert-private-key=<absolute path of the key.pem file
   ```

   Example

   ```
   admin@127.0.0.1>cluster update-ssl-certificate ssl-certificate=/home/support/new-cert/certificate.pem ssl-cert-private-key=/home/support/new-cert/privateKey.pem
   ```

   Sample Output

   ```
   CLUSTER SSL CERTIFICATE : -----BEGIN CERTIFICATE---
   MIIFpDCCA4ygAwIBAgIQDOxhTmAVWGsxVbypDcdj2TANBgkqhkiG9w0BAQsFADBs
   MQswCQYDVQQGEwJ1czERMA8GA1UEBxMIU2FuIEpvc2UxFTATBgNVBAoTDENvaGVz
   aXR5IEluYzEVMBMGA1UECxMMQ29oZXNpdHkgSW5jMRwwGgYDVQQDExNDb2hlc2l0a
   eSBJbmMgU2VydmVyMB4XDTI0MDYwMTExNTAwNloXDTM0MDUzMDExNTAwNlowbDEL
   MAkGA1UEBhMCdXMxETAPBgNVBAcTCFNhbiBKb3NlMRUwEwYDVQQKEwxDb2hlc2l0
   eSBJbmMxFTATBgNVBAsTDENvaGVzaXR5IEluYzEcMBoGA1UEAxMTQ29oZXNpdHkg
   SW5jIFNlcnZlcjCCAiIwDQYJKoZIhvcNAQEBBQADggIPADCCAgoCggIBALLcHuJX
   gx3bDw/a6Rag1wIwzQZdYaGIVjTe+B03EcTPLPfVMBG6fXVk2XApPPVu54lkxoGI
   ```

```
YJKj44hOGAOMfo+ecpPyWeDfp9VTguPxbwHL9UiqESE36JGMgVUt7yKXWW6qoLf/
mbsV0ublJhu9R19pfSuLt8dbMo6WCxUx2JAvFw0ZFBfOTLCOGIFeQlLOKcj77+XW
IyqSwCBAn+VLg7EaJ+u+cWegab48v5KZwrEvNcncd5iiMC5p9Em0QmqjzHXGC6Mq
IC5HiTYL4xrNj+XWjZYs7ZckbW5SBtzG0PSFW/tRW6fuuDl8EwyQVpxG0BcxkMM4
HbMyUiVXmovqBlJ/6qascQQBnmgV2xYBlFbaRVjgIgpFeQw7qML8675pckQqrsB1
oKbVtN0qDcEVaT1Jy811JRi+IIeKTQpfOH4EgXD3+fyvzY/qjE2wr/0q5bXpZ4yP
vepj6lXYue0UDQ/l0X8szn7P/97ppuv/j98SOK5CLTukRQNJ2tSQF/aPkj6r2bT9
yr0tF6NBJayX2Wv02/10nhRaFR00+y935oZSPGreYlX73rmRSCVjTIab1VCjm43y
o7Ds9yvWNuqeJy+zci0xL3ttZxaWiu6c4IMkGD+saqw5Qy50+XsPEIjnC1nY832f
T3QBcxIhE/4UMxK0mhwtq845EleaOZb/rmvTAgMBAAGjQjBAMA4GA1UdDwEB/wQE
AwIBBjAPBgNVHRMBAf8EBTADAQH/MB0GA1UdDgQWBBS8y1X1dn5EXIkZs9q2lVWv
CEVoOjANBgkqhkiG9w0BAQsFAAOCAgEACwInp6yPGJ0gQ99OSmpXmRmRIMG2Md3u
u03Z/wOUf7kCCpdwsfXsJQPz65bw6UtEqDzwfQfTBFuQUU4vNYEZOcb8PlSV9H0c
Jtr/jwi0VanUU2BmhqG6AlpfUa33sQR+c1VFhPwLDFm0mCKx+2T7UyGLCGCCH9KG
oGCuH8hksR+UTIFX3t7J4nSq5A3Oqm2ygxTfEeYOWS50jtcHXZAO+xWCLS9waUDS
Qoo0GhDjTT7XJ7bqVWiQ7RQQxwtJPxmYId9Ez2ZhLjKtXioX9R3YlXwuTaLYsfwB
MI2v+TbRSuUC8Y4/ql04835DaQ8QnlYvbGjptkVDr/Y/gA4rNclleOG9hswpHRZk
kUHEqkguV+DW169IJgnL/695tl2ites7bkrIUmj3kqnBeiIWM91Fvw7smXEgtTgZ
qA4vWt6ew+Xk8SaeOvL0e5OcZnFjTnUsbldXYZ2rToO65G18JOIcCW1HCrou/wmc
U072MP3GA8LS5iP+qrlLDfIAdoRzN42QmLoJ1gFuFfXIYUQOaqHczjMLi+aO5ljn
ZvBUyjDSocuN9ITFg59g5Ma25TlCW0ebyoulPqKfHb3JKeFjE3DVeBQOifJ1j1PM
bCekJuDB9EtKUzoWUOkGw0YRDDhtu9T9QoeDPiLWI4Te9
```

5. Run the following command to restart bridge and iris services:

   Command

   ```
   iris_cli cluster restart service-names=bridge,iris
   ```

   Sample Output

   ```
   [support@stor1-525488a33cal-node-1 ~]$ iris cli cluster restart
   service-names-bridge, iris
   ```

   ```
   Username: admin
   ```

```
'admin' Password:
```

```
Success: Restarting the cluster services [bridge iris] ...
```

Configure Certificate Authority Integration with Venafi

In addition to using the Cohesity CA or manually uploading CA-signed certificates, you can configure the cluster to obtain certificates from an external Certificate Authority using Venafi Trust Protection Platform (TPP).

With this integration, the cluster retrieves certificates from Venafi based on your organization's policy and automatically distributes them to all registered Cohesity Agents. This removes the need to manually install Cohesity Agent certificates.

Prerequisites

Before integrating with Venafi, ensure that:

- The Venafi Trust Protection Platform (TPP) server is deployed and reachable from all cluster nodes, with network connectivity, DNS resolution, and firewall access configured.
- A policy directory (Policy DN) exists in Venafi TPP. The policy must define certificate attributes, be associated with a certificate authority template, and you must have the Policy DN value.
- An OAuth application (Client ID) is configured in Venafi TPP with the required scope to request and retrieve certificates. You must also have a service account username and password with access to the application and policy.
- The certificate values you provide follow the rules defined in the Venafi policy.
- No other external CA is already configured on the cluster. If one is already configured, remove or update the existing configuration before proceeding. Contact Cohesity Support to remove an existing CA configuration.
- You have administrative access to run certificate management commands using the CLI on the cluster.
- You have the following details from your Venafi setup. These values are required when running the configuration command.

Parameter	Description	Required
ca-type	Type of customer managed CA provider. Set this value to VENAFI.	Yes
ca-url	Base URL of the customer managed CA provider API endpoint.	Yes

Parameter	Description	Required
venafi-ca-tpp-url	Base URL of the Venafi Trust Protection Platform (TPP) server used for authentication and certificate operations.	Yes
venafi-ca-username	Username of the Venafi service account used for authentication.	Yes
venafi-ca-password	Password of the Venafi service account.	Yes
venafi-ca-application-id	Venafi TPP application (OAuth client) ID used for certificate requests.	Yes
venafi-ca-oauth-scope	OAuth2 scope used when authenticating to Venafi TPP. This scope must allow certificate request and retrieval operations.	Yes
venafi-ca-policy-dn	Policy Distinguished Name (DN) in Venafi TPP that defines how certificates are issued.	Yes
venafi-ca-certificate-authority	Certificate Authority template (issuer) in Venafi TPP used to issue certificates based on the policy.	Yes
venafi-ca-zone	Descriptive zone or folder label in Venafi used to identify the CA configuration.	No

To configure the Venafi integration:

1. Run the following command to start Cohesity CLI:

   ```
   $ iris_cli
   ```

   When prompted, enter the username and password you use to log in to the Cohesity Cluster's user interface. After authentication, the CLI console opens, as shown below:

   ```
   iris_cli, Version: 7.4.1_release-20260506_8236697a

   Host: cork-clu-lts-005056ba9163-node-2, 127.0.0.1

   Cluster Server: 127.0.0.1:443

   admin@127.0.0.1>
   ```

2. Run the following command to configure the external CA integration:

```
cert-manager configure-generic-ca ca-type="VENAFI" ca-url="<customer_
ca_api_url>" venafi-ca-tpp-url="<tpp_server_url>" venafi-ca-
username="<username>" venafi-ca-password="<password>" venafi-ca-
application-id="<application_id>" venafi-ca-oauth-scope="<oauth_scope>"
venafi-ca-policy-dn="<policy_dn>" venafi-ca-certificate-authority="<ca_
template_name>" venafi-ca-zone="<optional_zone>"
```

This command requires details from your Venafi configuration. These values allow the cluster to connect to Venafi, authenticate, and request certificates based on the configured policy.

3. Run the following command to verify that the customer-managed CA configuration is saved. Confirm that the displayed values match your Venafi configuration, including the TPP URL, Policy DN, certificate template, application ID, OAuth scope, and username.

```
get-generic-ca-config
```

4. After the configuration is complete, generate a new certificate to confirm that Venafi is issuing certificates correctly. To generate a new certificate, run the following command:

```
cert-manager new-cert city="<city>" common-name="<common_name>"
country-code="<country_code>" organization="<organization>"
organization-unit="<org_unit>" state="<state>"
```

This command generates a new cluster certificate using the configured customer managed CA. After the command runs successfully, the following files are created:

- CA Certificate
- Private key
- Certificate

Sample Output

```
Wrote CA Cert to /home/cohesity/data/cert-manager/new-cert/ca_cert.pem
```

```
Wrote PrivateKey to /home/cohesity/data/cert-manager/new-
cert/privateKey.pem
```

```
Wrote Certificate to /home/cohesity/data/cert-manager/new-
cert/certificate.pem
```

5. To retrieve and inspect the generated certificate, run the following command:

```
openssl x509 -in /home/cohesity/data/cert-manager/new-cert/ca_cert.pem
-text -noout
```

Verify that the following details are correct:

- The issuer matches the Venafi-configured Certificate Authority common name and not the Cohesity internal CA.
- The subject details, validity period, and other paramaters match the rules defined in the Venafi Policy DN.

> **Note:** This configuration applies only to newly generated certificates. Existing certificates and agents continue to use their current certificates.

Related Topic

Generate SSL Certificates with a Subject Alternate Name

Disable AES-128 Ciphers Using the CLI

You can disable AES-128 cipher suites on the cluster's web services and S3-compatible object store endpoints.

Disable AES-128 Ciphers on Port 443

Perform the following steps to disable AES-128 Ciphers on Port 443:

1. Open the Cohesity CLI.
2. To view the ciphers currently enabled on the cluster, run the following command:

   ```
   iris_cli -username=<username> -password=<password> -skip_password_
   prompt=true cipher list
   ```

   The command output displays the list of **ENABLED CIPHERS** and **DISABLED CIPHERS**.
3. Run the following command to disable the AES-128 ciphers on port 443:

   ```
   cipher update enable=false ciphers=<AES_128 cipher1>,<AES_128 cipher2>
   ```
4. Restart the Iris service to apply the change:

   ```
   cluster restart-services service-names=iris
   ```
5. After the service restarts, run `cipher list` once again to confirm that AES-128 cipher suites are disabled.

Disable AES-128 Ciphers on Ports 3000 and 3001

Perform the following steps to disable AES-128 and other legacy Ciphers on ports 3000 and 3001:

1. Open the Cohesity CLI.
2. Run the following command to disable the AES-128 TLS 1.3 cipher suites on the S3-compatible object store on ports 3000 and 3001:

   > **Note:** The naming scheme for the ciphers in the command below is different from the scheme used for port 443.

   ```
   iris_cli -username=username -password=password -skip_password_
   prompt=true cluster update-gflag service-name=bridge reason=disable_
   aes_128 gflag-name=bridge_s3_disabled_cipher gflag-value= TLS_AES_128_
   GCM_SHA256,AES128-GCM-SHA256,AES256-GCM-SHA384,AES128-SHA,AES256-SHA
   effective-now=true
   ```

   Following is the list of ciphers used by port 3000 and 3001:

   - TLS 1.2:
     - "ECDHE-ECDSA-AES128-GCM-SHA256"
     - "ECDHE-ECDSA-AES256-GCM-SHA384"
     - "ECDHE-ECDSA-AES128-SHA"
     - "ECDHE-ECDSA-AES256-SHA"
     - "ECDHE-RSA-AES128-GCM-SHA256"
     - "ECDHE-RSA-AES256-GCM-SHA384"
     - "ECDHE-RSA-AES128-SHA"
     - "ECDHE-RSA-AES256-SHA"
     - "AES128-GCM-SHA256"
     - "AES256-GCM-SHA384"
     - "AES128-SHA"
     - "AES256-SHA"
   - TLS 1.3:
     - "TLS_AES_128_GCM_SHA256"
     - "TLS_AES_256_GCM_SHA384"
     - "TLS_CHACHA20_POLY1305_SHA256"
3. Restart the Bridge service to apply the changes:

   ```
   cluster restart-services service-names=bridge
   ```

Disabling AES-128 ciphers lets organizations that require AES-256-only encryption policies restrict the cluster to stronger cipher suites.

> **Note:** Before disabling any ciphers, verify that your clients and backup agents support AES-256 cipher suites. Disabling AES-128 ciphers on a cluster that has clients or agents that only support AES-128 will prevent those clients from connecting.

Enable File Modification Detection

The **Enable File Modification Detection** feature sends notifications to the user if a shared file or folder is modified. To enable this feature, contact Cohesity Support.

Update the Change Notification Events Limit

This feature enables you to update the number of events generated per file object within a given time interval. The default limit is 5 events in one minute (60000000 usecs).

To update the events notification, contact Cohesity Support.

Customize Preferences

The **Customization** page in the Cohesity cluster allows you to customize various settings and options to tailor your experience according to your personal preferences. You can modify settings related to your account, user interface, and interactions with the Cohesity cluster.

For example, you can hide the options on the left navigation that are not relevant to your environment, hide the options from the **Data Protection** menu, or select a favorite dashboard. You can customize the preferences both at the cluster level and user level.

> **Note:** The Customization feature is not available for tenant users.

Customize Preferences at the User Level

User-level customization of the Cohesity cluster applies to the logged-in user.

Perform the following steps for user-level customization:

1. Log in to the Cohesity cluster.
2. Navigate to **Settings > Customization**.
3. Mouse hover on the option and click the edit icon to customize the following:

User Interface Components	Description
Language	Set the default language. Cohesity supports the following languages: • English (US) • Japanese (JP)
Byte Scaling	Adjust the scale or size of data in terms of bytes. Cohesity offers the following byte scaling options: • Base 1024 (1 KiB = 1024) • Base 1000 (1 KB = 1000)
Recovery Location	Set the recovery location. Cohesity supports the following recovery locations: • Original Location • New Location
Default Dashboard	Set the default dashboard to view when you log in to the Cohesity cluster. The available options are: • Summary • Data Protection • SmartFiles • Cloud • MS SQL • Continuous Data Protection
Simple Create Workflows	Set the simple create workflows to display the quick protect or quick create pop-ups when creating a Protection Group, Policy, or View. Cohesity supports the following: • Protection • Policy • View

Time Format	Select how time must be represented in the Cohesity cluster: • 12-hour clock • 24-hour clock
Persist Snack Bars	Snack bars are brief notifications that appear at the bottom of the Cohesity cluster after an action is performed. Choose whether to keep messages, alerts, or notifications visible until you interact with them or let them disappear automatically: • **Persist Snackbar Messages** - Messages, alerts, or notifications stay visible until you acknowledge the message or dismiss it manually. • **Disappearing Snackbar Messages** - Messages, alerts, or notifications disappear automatically after a few seconds.

4. Click **Save**.

Customize Preferences at the Cluster Level

Cluster-level customization of the Cohesity cluster allows you to hide dashboards, navigational items, and protection adapters that do not apply to your environment.

> **Note:** The **Cluster** section is available to users only with the Manage Cluster privilege.

Perform the following steps for cluster-level customization:

1. Log in to the Cohesity cluster.
2. Navigate to **Settings > Customization**.

3. Mouse hover on the option and click the edit icon to customize:

User Interface Components	Description
Dashboards	Select the dashboards to be hidden from the drop-down menu on the **Dashboards** page. **Note:** A hidden dashboard at the cluster level is not available for the user to select at the user level.
Navigation Items	Select the options to be hidden from the left navigation pane. **Note:** You cannot hide **Settings** and **Customization**.
Protection Adapters	Select the adapters to be hidden from the **Protection**, **Recovery**, and **Sources** pages. **Note:** You cannot hide an adapter if it includes registered sources.

4. Click **Save**.

Cohesity Cluster Services

A Cohesity cluster consists of multiple internal services that work together to deliver core cluster operations like data protection, storage management, monitoring, security, and overall cluster operations. Each service performs a specific function, such as orchestrating backup and recovery workflows, managing cluster health, handling I/O operations, or enabling indexing and search.

This topic provides a consolidated overview of the cluster services that support core cluster operations. It includes:

- **Internal service names**: The service identifier as it appears in logs, alerts, support bundles, and diagnostic outputs. You may encounter this name during troubleshooting or when working with Cohesity Support.
- **External service names**: A simplified, descriptive name that represents the service's function. Use this name when discussing system behavior at a high level.
- **Description**: A brief explanation of what the service does and the responsibilities it handles within the cluster.

- **Primary function**: A concise summary of the service's main role, such as backup orchestration, monitoring, indexing, or cluster management.
- **Functional category**: A high-level grouping that indicates the functional area to which the service belongs, such as Data Protection, Base Infrastructure, Monitoring, Networking, or Agent services. Categories help you quickly identify related services and understand how they fit into overall cluster operations.

The following table lists the Cohesity cluster services and their corresponding details:

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
alerts	Alerts	Continuously monitors the cluster for system events, anomalies, and threshold breaches, generating real-time notifications to administrators.	Provides proactive alerting and notification for critical system events, enabling timely intervention.	Alerts
apollo	Healer and Space Reclamation	Responsible for background healing, garbage collection, and space reclamation to optimize storage and maintain data integrity.	Optimizes storage utilization and ensures data health through automated space reclamation and healing.	Disk
ageis	Fault Isolation	It is related to alert management and metadata handling	Manage alert metadata and handle RPC timeout errors	Monitoring / Alert Management
athena	Marketplace Apps	Manages Cohesity Marketplace, hosting and orchestrating third-party and Cohesity-developed applications for extensibility.	Provides a secure platform for deploying and managing Marketplace applications.	Marketplace

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
athena_proxy	Marketplace Apps Helper	Acts as an auxiliary service supporting Marketplace app operations, handling orchestration and integration tasks.	Facilitates deployment and integration of Marketplace applications.	Marketplace
athena_ watchdog	Marketplace Apps Monitor	Monitors health and status of Marketplace applications, ensuring optimal performance and alerting administrators to issues.	Provides health monitoring and operational assurance for Marketplace apps.	Marketplace
atom	Continuous Data Protection	Delivers continuous data protection for VMware and other workloads, capturing changes in near real-time.	Provides continuous backup and granular recovery for mission-critical workloads.	DataProtection
baseos	BaseOS	It is a component within the software bundle that includes critical CVE fixes, kernel updates, driver updates, and optionally bug fixes for other user-mode packages	It is to hold critical CVE fixes, kernel updates, driver updates, and optionally bug fixes for other user-mode packages	Patch
bifrost	Hybrid Extender Agent	It is a framework used within Cohesity's infrastructure to manage multi-tenancy and network isolation challenges	It is to serve as a capabilities framework that abstracts individual Hybrid Extender (HyX) capabilities for tenants, providing a global view of these capabilities	Networking

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
bifrost_broker	Hybrid Extender	Enables secure, multi-tenant networking and hybrid cloud connectivity, supporting seamless data movement.	Facilitates secure networking and hybrid cloud extension.	Disk
bridge	I/O Operations	Manages all input/output operations within the cluster, ensuring high-throughput and low-latency data access.	Optimizes cluster I/O operations for efficient data handling.	Disk
bridge_proxy	I/O Operations Helper	Supports the main I/O service by handling auxiliary data operations for performance optimization.	Assists in efficient data movement and I/O processing.	Base Infrastructure
ca	Cohesity Certificate Authority	A Certificate Authority (CA) is a critical component in the Public Key Infrastructure (PKI) framework, responsible for issuing digital certificates	It's role in enabling secure communication between Cohesity clusters and host agents by replacing self-signed certificates with those generated in the customer environment	Certificates
compass	Cluster DNS Server	Provides DNS resolution and load balancing for cluster VIPs, ensuring seamless node communication.	Handles internal DNS lookups and distributes traffic for high availability.	Networking

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
core_helper	Core Dump Handler	Manages collection and handling of core dumps generated during system failures for diagnostics.	Provides core dump management for troubleshooting and analysis.	Support
eagle_agent	Support Automation	Automates support tasks and integrates with Helios for proactive analytics and issue detection.	Enables proactive support and automated issue resolution.	Disk
elrond	Prediction Model	Predicts and optimizes I/O operations using analytics to anticipate workload patterns.	Enhances performance by predicting and optimizing I/O workloads.	Base Infrastructure
gaia	Gaia dataplane service	Gaia is a generative AI application designed to empower users with secure and efficient data management capabilities within the Cohesity Data Cloud	It is to serve as an advanced generative AI application that utilizes retrieval-augmented generation (RAG) technology to manage and process backup data securely within the Cohesity Data Cloud	AI
gandalf	Cluster Configuration Manager	Manages cluster configuration, consensus, and distributed coordination for consistent operations.	Maintains configuration integrity and manages distributed consensus.	Reporting

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
groot	Reporting	Generates and delivers comprehensive reports on cluster operations, usage, and health.	Provides actionable insights through detailed reporting and analytics.	(Cloud) Disk
heimdall	heimdall	Heimdall is responsible for managing the lifecycle of fleet instances launched for supporting various workflows across our platform.	Manages as a component within the Cohesity infrastructure responsible for managing cloud infrastructure operations, particularly in the context of fleet management and gflag settings	Infrastructure / Cloud Operations
icebox	Archive	Handles data archiving for long-term retention and retrieval to external or cloud storage targets.	Manages efficient data archiving and retrieval for compliance and cost optimization.	Web Server
iris	UI and REST API	Provides web-based user interface and REST API endpoints for cluster management.	Delivers unified management interface and API access for cluster operations.	Web Server
iris_proxy	UI and REST API Helper	Supports UI and API services by handling auxiliary tasks for responsiveness.	Enhances performance and reliability of UI and API operations.	Base Infrastructure
keychain	Key Management	Manages encryption keys and secure storage to ensure data security and compliance.	Provides robust key management and secure storage for encrypted data.	Indexing

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
librarian	Index Catalog	Maintains index catalog for fast search and file-level recovery across the cluster.	Enables efficient indexing, search, and recovery operations.	Base Infrastructure
logwatcher	Log Monitor	Monitors, collects, and rotates system logs for diagnostics and compliance.	Ensures comprehensive log management for troubleshooting and auditing.	DataProtection
madrox	Replication	Manages data replication across clusters or sites for disaster recovery and high availability.	Orchestrates reliable data replication for business continuity.	DataProtection
magnus_server	Disaster recovery Helper	The magnus_server is involved in several operations, primarily related to patch management and execution within the Cohesity infrastructure	Manages to run tasks related to Disaster Recovery within the Cohesity infrastructure	DataProtection / Disaster Recovery
magneto	Backup and Recovery Manager	Orchestrates all backup and recovery operations, ensuring data protection and recoverability.	Manages backup scheduling, execution, and recovery workflows.	Base Infrastructure
newscribe	Distributed Key Value Store	Provides distributed configuration and metadata storage for high availability.	Stores and replicates configuration and metadata across the cluster.	Base Infrastructure
nexus	Cluster Manager	Oversees cluster management, orchestration, and lifecycle operations including upgrades.	Coordinates cluster-wide management and orchestration tasks.	Base Infrastructure

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
nexus_proxy	Cluster Manager Helper	Assists the main cluster manager with auxiliary operations for smooth orchestration.	Supports and enhances cluster management operations.	Base Infrastructure
nfs_proxy	NFS proxy	The nfs_proxy is responsible for handling NFS v4.1 backups and restores	It manages within the Cohesity infrastructure is to facilitate the handling and management of NFS (Network File System) operations	DataProtection / Networking
rtclient	Support Channel	Facilitates secure support communication channels for direct interaction with Cohesity support.	Provides secure, reliable channel for support communications.	DataProtection
siren_server	Service Monitor and Logs Collector	Monitors cluster health and collects system logs for visibility and troubleshooting.	Aggregates metrics and logs to enable real-time monitoring and diagnostics.	Monitoring
smb2_proxy	SMB Backup Helper v2	Assists with backup operations for SMB v2 workloads, ensuring compatibility and performance.	Supports and optimizes SMB v2 backup workflows.	DataProtection
smb_proxy	SMB Backup Helper v1	Assists with backup operations for SMB v1 workloads, maintaining legacy compatibility.	Supports and optimizes SMB v1 backup workflows.	Base Infrastructure

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
stargate	Stargate	The stargate service is involved in handling task updates and managing connections within the Cohesity infrastructure	Manage and update adapter flag files	Infrastructure / Task Management
stats	Cluster Metrics	Collects, aggregates, and reports on cluster metrics for visibility into performance and resource utilization.	Enables real-time monitoring and analytics of cluster health and usage.	Base Infrastructure
statscollector	Cluster Metrics Store	Stores and manages historical cluster metrics for trend analysis and capacity planning.	Provides persistent storage and analysis of cluster metrics.	DataProtection
storage_proxy	Data Handler Storage Service	Assists with NAS backup operations for efficient data protection of network-attached storage.	Optimizes and supports NAS backup workflows.	DataProtection/Disk
throttler	Adaptive Throttler	Manages resource allocation and rate limiting for data protection and disk operations.	Controls and optimizes resource usage across data protection and disk operations.	Base Infrastructure
tricorder	Cluster Health Monitor	Continuously monitors cluster health, performs system checks, and validates upgrade readiness.	Ensures cluster health and pre-upgrade validation for seamless operations.	Web Server

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
ui	User Interface	Provides the main graphical interface for cluster management with intuitive navigation.	Delivers a user-friendly interface for administrators and operators.	DataProtection
vault_proxy	Archive Helper	Assists with data archiving to cloud or external targets for secure long-term storage.	Supports and optimizes data archiving workflows.	Indexing
yoda	Indexing	Manages indexing for search and file-level recovery, enabling fast data retrieval.	Provides high-performance indexing and search capabilities.	Indexing
yoda_agent	Indexing Helper	Assists the main indexing service by handling distributed indexing tasks for scalability.	Supports distributed indexing and enhances search performance.	Agent
magneto_agent_windows	Secure Connector for Windows	Agent for secure backup and recovery on Windows systems.	Enables secure, reliable backup and recovery for Windows workloads.	Agent
magneto_agent_aix	Secure Connector for AIX (Go)	Agent for secure backup and recovery on AIX (Go) systems.	Enables secure, reliable backup and recovery for AIX workloads.	Agent
magneto_agent_aix_7_1	Secure Connector for AIX (7.1)	Agent for secure backup and recovery on AIX 7.1 systems.	Enables secure, reliable backup and recovery for AIX 7.1 workloads.	Agent
magneto_agent_aix_java	Secure Connector for AIX (Java)	Agent for secure backup and recovery on AIX (Java) systems.	Enables secure, reliable backup and recovery for AIX Java workloads.	Agent

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
magneto_agent_solaris	Secure Connector for Solaris	Agent for secure backup and recovery on Solaris systems.	Enables secure, reliable backup and recovery for Solaris workloads.	Agent
magneto_agent_solaris_10	Secure Connector for Solaris (10.0)	Agent for secure backup and recovery on Solaris 10.0 systems.	Enables secure, reliable backup and recovery for Solaris 10 workloads.	Agent
debian_agent	Agent for Debian	Agent for secure backup and recovery on Debian systems.	Enables secure, reliable backup and recovery for Debian workloads.	Agent
redhat_agent	Agent for Redhat	Agent for secure backup and recovery on Redhat systems.	Enables secure, reliable backup and recovery for Redhat workloads.	Agent
suse_agent	Agent for SUSE	Agent for secure backup and recovery on SUSE systems.	Enables secure, reliable backup and recovery for SUSE workloads.	Agent
linux_ppc_agent	Secure Connector for Linux PPC Agent	Agent for secure backup and recovery on Linux PPC Agent systems.	Enables secure, reliable backup and recovery for Linux PPC workloads.	Agent
hpux_agent	Secure Connector for HPUX Agent	Agent for secure backup and recovery on HPUX systems.	Enables secure, reliable backup and recovery for HPUX workloads.	Agent
magneto_agent_linux	Secure Connector for Linux	Agent for secure backup and recovery on Linux systems.	Enables secure, reliable backup and recovery for Linux workloads.	Disk / DataProtection

Internal Service Name	External Service Name	Description	Primary Function	Functional Category
Hydra	Write Cache service	Acts as the high-speed write cache for the Cohesity cluster, capturing incoming data and enabling rapid, reliable data ingestion.	Buffers and optimizes all incoming writes for immediate access and high-performance data protection.	Support / Base Infrastructure
Genie	Support Channel & Health Monitor service	Provides secure remote support access and continuous health monitoring for the Cohesity cluster, ensuring proactive maintenance and rapid troubleshooting.	Enables encrypted remote support and health monitoring, facilitating secure diagnostics, telemetry collection, and system checks.	Support
mongo_db_ physical	MongoDB Physical	Provides MongoDB DataProtect solution, which involves physical or virtual MongoDB instances as part of its core components	It is to provide consistent backup and recovery for MongoDB deployments. Cohesity's MongoDB DataProtect solution integrates natively with MongoDB to facilitate this process	DataProtection
aws_dynamo_ db	Aws DynamoDB	Provides protection of AWS DynamoDB tables by requiring specific permissions for backup and recovery operations	It is supported as a service that can be registered and protected using Cohesity's DataProtect as a Service	DataProtection / Cloud Services

Software Update

Cohesity provides a comprehensive solution for managing software upgrades and patch installations within your cluster. The new design streamlines the update process, enabling you to upload and apply both upgrade and patch packages, perform necessary pre-checks, and initiate the upgrade or patch on all cluster nodes. When you download an upgrade package, you also have the option to download a corresponding patch, ensuring your cluster is fully updated. Additionally, Cohesity provides real-time tracking for both upgrades and

patches, allowing you to monitor the status of each node, track completion times, and resolve any issues that may occur during the process.

Cohesity ensures a seamless and efficient way to maintain your cluster with the following key functions:

- Download Package
- Upgrade Cluster
- Apply Patch
- Upgrade with Patch

Before you begin, it is important to familiarize yourself with the following upgrade paths, prerequisites, and considerations to ensure a smooth upgrade and patching process.

Upgrading to 7.4

You can upgrade your Cohesity cluster from previous releases to 7.4.1. The following table provides details on supported upgrade paths.

Your Current Release	Upgrade Path to 7.4.1
• 7.4 • 7.3.2 • 7.3.1 • 7.3 • 7.2.2_u2 • 7.2.2_u1 • 7.2.2 • 7.2.1 • 7.2 • 7.1.2_u6 • 7.1.2_u4 • 7.1.2_u3 • 7.1.2_u2 • 7.1.2_u1 • 7.1.2 • 7.1.1 • 7.1	7.4.1 directly

> **Note:** Direct upgrades from 6.8.x to 7.4.1 are not supported. If your cluster is running any 6.8.x release (such as 6.8.1, 6.8.1_u1, or 6.8.2), a step upgrade is necessary. First, upgrade to 7.1.2_u6, and then proceed to 7.4.1.

Release Upgrade Policy

Policy	Example
Cohesity will support upgrades from the latest release of the prior LTS release branch, which includes all LTS designated releases within the branch, to the most recent release of the current LTS branch.	6.6.0d+ (LTS designated releases: 6.6.0d_u3, 6.6.0d_u4, 6.6.0d_u5, 6.6.0d_u6) to 6.8.2 LTS designated release will be supported.
Cohesity will not allow upgrades by default to any release that is older in time irrespective of the release branch. Exceptions are to be managed on a case-by-case basis.	6.5.1f_release-20210825_596bb917 is released after 6.6.0c_release-20210822_0d731348. Therefore, an upgrade from the 6.5.1f version to the 6.6.0c version is not supported. This policy is also applicable to patches. If you have upgraded your Cohesity cluster to a patch released after the LTS release, upgrading to that LTS release is not supported. However, you can upgrade to any LTS version released after the patch. For example, your Cohesity clusters were upgraded to 6.6.0d_u5 in July 2022. Cohesity released 6.8.1_u1 on Nov 2022 and the 6.6.0d-p32 patch on March 2023. If you've applied 6.6.0d-p32, you cannot upgrade to 6.8.1_u1. However, you can upgrade to the upcoming 6.8.1_u2 release.
Cohesity will support the release N-1 upgrade without an intermediate step. (N is defined as the current release branch).	7.2.x to 7.4.1 is supported. 7.4.1 is the current release branch.
Cohesity will support the release N-2 upgrade without an intermediate step. (N is defined as the current release branch).	7.1.x to 7.4.1 is supported. 7.4.1 is the current release branch.

Policy	Example
When a specific release is declared LTS, Cohesity will support upgrading from the open LTS releases to the new LTS release. This will include the three most recent releases on the LTS branch to the new LTS release.	6.8.1, 6.8.1_u1, 6.8.1_u2, 6.8.1_u3, 6.8.1_u4, 6.8.1_u5, 6.8.1_u6, 6.8.1_u7 to 6.8.2.

Upgrade Considerations

Note the following about upgrading the Cohesity cluster to 7.4.1:

- Cohesity does not support rolling back to older versions.
- To upgrade the Cohesity cluster from a version that is no longer supported, Cohesity recommends you to upgrade to any of the supported versions mentioned in the Upgrade Paths, and then perform an upgrade to the latest release version. For information on Cohesity Products that have reached the end of support, see Cohesity Products End of Support.
- See Cohesity Feature Deprecation to review the list of features marked for deprecation for Cohesity 7.4.1 and later releases.
- Before performing the upgrade, ensure that the cluster data space and metadata space utilized is less than 85%. After the cluster upgrade, the Garbage Collection algorithms take 3 to 4 days to trigger. Hence, ensure that the cluster has enough space during this period. Space constraints may lead to backup and replication failures on the Cohesity cluster.
- If you are running remote adapter jobs and the cluster is upgraded, the jobs will be disrupted during the upgrade process. The jobs will be killed and restarted multiple times during the upgrade.
- Starting 6.8.2 and 7.1.2, the Cohesity indexing service is optimized to automatically identify and delete stale directories at regular intervals, which were created for indexing. After upgrading from a version without this optimization, the cluster indexing service will remove any stale directories identified, which may result in cluster-free space increase.
- Cohesity recommends upgrading the Cohesity Agent on Physical Servers and the Cohesity installed Agent on VMs to the latest release version of the Cohesity cluster.
- Cohesity recommends upgrading the Cohesity cluster first, followed by the Cohesity Agent. Upgrading an agent before the cluster is likely to impact the existing functionality and disruptions may be observed due to agent being on a higher version

than the cluster. Cohesity also recommends the agents be on the same, latest major version as the Cohesity cluster to get the latest security fixes and benefit from newer features.

- After upgrading to the latest version, if there is an IP subnet conflict, the **Enable Apps Management** toggle in **Marketplace** > **My Apps** is turned off. Navigate to **Settings** > **Summary** > click **Configure** and specify a different IP address in the **Configure Apps management network** field and then turn on the **Enable Apps Management** option.
- If you are on a Cohesity Cloud Edition cluster and using Marketplace Apps, then when you upgrade the Cohesity Cloud Edition cluster to 7.4.1, connectivity among the Marketplace Apps could be impacted* due to Flannel moving to etcd v3 APIs. It is recommended to pause any Marketplace Apps before the upgrade and resume them once the upgrade is complete.

  ***Impact**: Running workloads, Protection Groups, or scans related to the Marketplace App might see network disruption during the upgrade.
- For pure PXG clusters, before upgrading from version 6.6 to 6.8.1 or above, make sure that your cluster usage is below 95%. After the upgrade, there is a known issue where the available data space may decrease. Even clusters that are using only 88% of disk space before the upgrade have experienced out-of-space errors afterward, which can lead to backup failures. To avoid disruptions, Cohesity strongly recommends reducing disk usage well below 95% before starting the upgrade process.
- Cohesity is working to qualify the Windows agent for Japanese Windows version that uses Shift-JIS encoding. Wait to upgrade agent to version 7.4.1 until Cohesity releases a version that supports Japanese language (Shift_JIS) based Windows OS.
- Manual disk-space verification on nodes is not required. During the upgrade, the system performs pre-upgrade checks and automatically cleans up disk space when possible. If sufficient space is still unavailable, the upgrade stops immediately and provides a clear error message indicating the specific node and path that require additional space.
- If you are leveraging storage snapshots for VMware VM backups in releases prior to 7.2 and plan to upgrade the Cohesity cluster to release 7.2 or later, ensure that SAN connectivity between Cohesity nodes and the Storage Array (Pure, HPE Alletra, Nimble, and IBM) is in place before the upgrade to prevent VM backup failures.
- Starting with Cohesity version 7.2.2, the Cohesity BaseOS transitioned from RHEL 7.x to RHEL 9.x, and the time synchronization service changed from ntpd to chronyd. Due to this architectural change, Windows-based NTP sources (Windows Time Service / W32Time) are not supported for Cohesity clusters running version 7.2.2 or later. Configure Linux-based NTP sources for time synchronization before upgrading to ensure cluster stability and supportability.

Databases

- The addition of the new Postgres database could cause UI slowness until the ETL process completes. The bootstrap run of the ETL process pulls the entire data set to populate the database. The initial run has a slight performance impact. In the case of upgrades, data population happens in the post-upgrade step. Subsequent upgrades will not be affected.
- After upgrading to the latest version, to display SAP HANA log backups in the Cohesity cluster, you need to modify the existing registered source and set the `et-log-backup` source registration parameter to `true`. Only the log backups triggered after enabling `et-log-backup` will be shown on the Cohesity cluster.

  > **Note:** After modifying the source configuration (with`-et-log-backup=true`), a full backup is mandatory. The initial full backup must be completed before any log backups appear on the Cohesity user interface.

- If you are upgrading to version 7.4.1 or later, you need to update the existing SAP HANA source configuration to enable auto-discovery and entity hierarchy. Set the `--entity-hierarchy` source registration parameter to "`true`." After updating the source configuration (with `--entity-hierarchy=true`), a full backup is mandatory.
- For **SQL log backups**, if you are upgrading the cluster to version 7.2 and above, ensure that the Cohesity cluster bridge node VIPs on port 11117 are reachable from the SQL source. In case of the multitenant environment, ensure that the Hybrid Extender IPs on port 11117 are reachable from the SQL Source.

  > **Note:** The SQL log backup will fail post-upgrade if the above-mentioned port requirement is not satisfied.

- If you have upgraded your cluster to Cohesity 7.4.1 from an earlier version and already have an AWS account registered, you must add new permissions to protect the newly supported AWS workloads:
  - Amazon RDS (Ingest-based protection)
  - Amazon Redshift
  - Amazon DocumentDB
  - Amazon DynamoDB

Administration

- To generate a new SSH key after upgrading the Cluster, contact Cohesity Support.
- Cohesity Support Engineers require a Support Channel token to remotely log into the Cluster using SSH for on-demand assistance. From your Cohesity cluster, you need to

copy the Support Channel token and provide it while raising a request for on-demand assistance.

- The Secure Shell restricts access to the host commands or scripts. After you upgrade to 6.7 or later version, the secure shell might have the following impact on your existing Cohesity Data Cloud deployments:
    - Access to the bash shell using SSH will be no longer available to the support user account without authorization from Cohesity.
    - If you run custom scripts using SSH on your Cohesity cluster, the scripts may fail. In this case, Cohesity recommends the following:
        - Verify if there is an alternate method to use Cohesity CLI commands or REST API and update your scripts accordingly.
        - Verify if a corresponding Cohesity CLI command is available in the supported list of CLI commands; if so, use the supported CLI command. If the CLI command is not available in the supported list of commands, contact Cohesity Support to enable the CLI command.
    - The private binary or tools running on the Cohesity nodes might fail. Contact Cohesity Support for options to install private binaries or tools.
    - Sudo access is disabled by default. For support channel access, enable the sudo access. For more information, see Enable or Disable Linux Sudo Access.
- If there is a source that is registered before the upgrade and assigned to an organization, then unassigning its root entity is not allowed. You can unassign the source if it is not assigned to an organization, and it will get assigned after the upgrade.

NoSQL and Hadoop

- To continue using Cohesity NoSQL & Hadoop services on the Cohesity cluster version 7.2.2, you must upgrade the NoSQL & Hadoop service to the 7.0.0 version available on Helios.
- If you are running NoSQL and Hadoop app, Cohesity recommends the following before upgrading the Cohesity cluster:
    - Pause the protection runs by navigating to **Data Protection > Protection** . From the Action Menu (⋮) of the required protection run, select **Pause Future Runs**.
    - Pause the NoSQL and Hadoop app by navigating to **Marketplace > My Apps** . From the Action Menu (⋮) of the app instance, select **Pause**.

  After upgrading the Cohesity cluster to the latest version, contact your Cohesity account team to check if the upgraded Cohesity cluster requires a new NoSQL and Hadoop app. If it requires a new version of the app, you must upgrade to the latest version of the NoSQL and Hadoop app. Once the cluster upgrade is complete, resume the app, and then the protection runs.

Microsoft 365

- If you upgrade your Cohesity cluster to 6.8.2 or later versions and currently backing up Microsoft 365, ensure that you add the required Microsoft Graph Permissions related to MS Groups to your custom application to continue using your existing Protection Groups and protect your Microsoft 365 data.
- After upgrading to the 7.4.1 version, if you are replicating the Mailbox data to a remote Cohesity cluster, then ensure that you upgrade the remote Cohesity cluster to the 7.4.1 version.

Single Node Cluster Upgrades

Single node cluster upgrades must be run when the upgrade will have the least impact. During the upgrade of a single node cluster, the node is rebooted and during the reboot, the cluster is unable to process Protection Groups, recover tasks, or any other workflow.

Virtual Edition Deployment

The following are the requirements for the Virtual Edition deployment for 6.8 and later versions:

- small (8 TB) configuration supports Virtual Machines with 12 vCPUs, 32 GB of memory, and 64 GB virtual disk to store the operating system.
- large (16 TB) configuration supports Virtual Machines with 24 vCPUs, 64 GB of memory, and 64 GB virtual disk to store the operating system.

For more information, see Virtual Edition for VMware Setup Guide and Virtual Edition for Clustered VMware Setup Guide.

Replication Environments

- If the cluster replication is configured, verify that the network connectivity is functioning properly during the upgrade to ensure the cluster replication relationship is successfully upgraded to use AES-256-GCM for encryption.
- In a replication setup, when you upgrade your Cohesity cluster to Cohesity 6.6 or later and you use the default System Admin password, you will be prompted to change the password. After changing the password, you must update the new password on the replication partner cluster.
- For information about using replication between Cohesity clusters running different versions, see Replication Compatibility.

  For the latest notes about replication compatibility, prerequisites and limitations, see the *Release Notes*.

Upgrade and Patch Workflows

If the Cohesity cluster is not running the newest software version, upgrade the Cohesity cluster. Only upgrade packages that are newer than the current version can be installed.

You can upgrade the software using the Cohesity cluster as described in the procedures below. You can use the same procedure and upgrade package to upgrade any Cohesity cluster, regardless of whether the Cohesity cluster is hosted on on-premises hardware, Cloud Edition or Virtual Edition.

Download Package

To download an upgrade or patch package, you can select the desired package type from either a URL or your local machine. Once the package is chosen, you can proceed with downloading, or opt to apply the upgrade or patch immediately.

Perform the following steps to download the package:

1. Go to **Settings** > **Software Update** and click **Get New Package**.

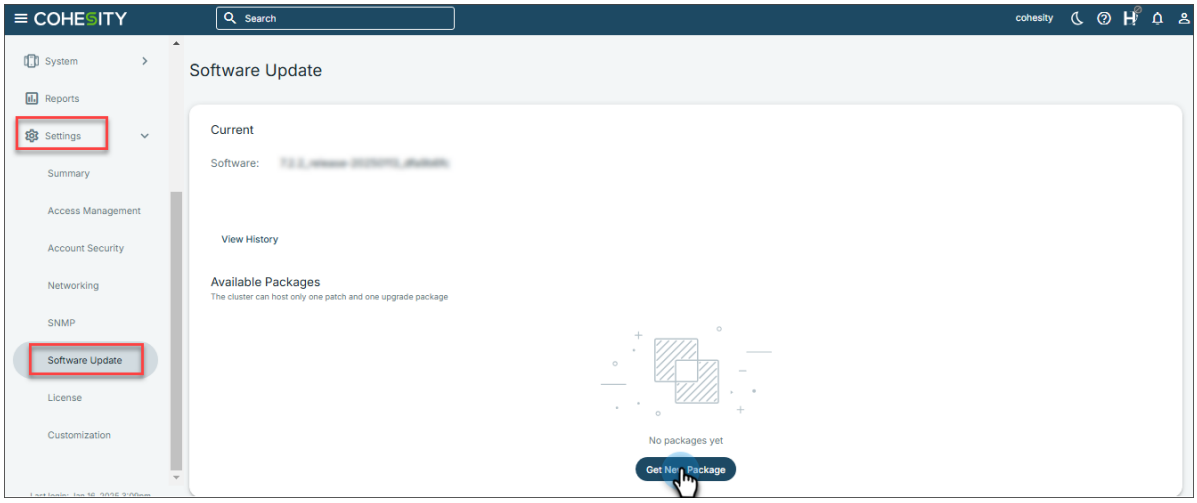

A dialog appears with the following package options:

1. **Upgrade Package**
2. **Patch Package**
3. **Upgrade and Patch Packages**

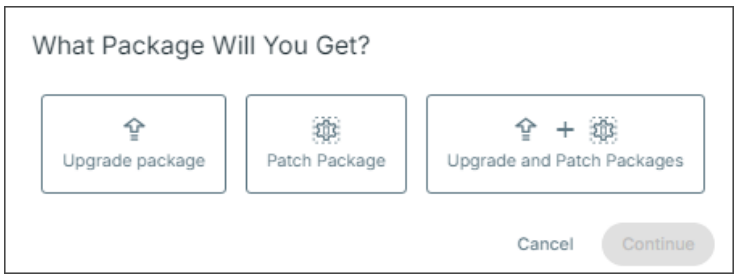

2. Select the required package and click **Continue**.
   A **Get New Upgrade Package** dialog appears.

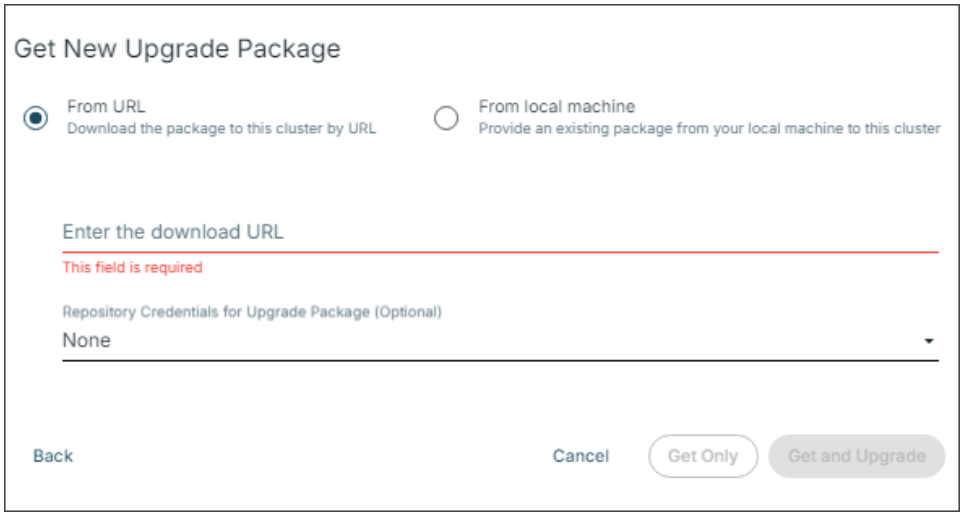

3. Choose one of the following options to download the package.

Option	Field	Description
From URL	**Enter the download URL**	Provide the package's URL
	Repository credentials for package (Optional)	Repository credentials can be used to download packages from secure repositories that require authentication to access the files. Choose one of the following options from the drop-down: **Basic** - Select this if your repository requires HTTP basic authentication. Enter the **Username** and **Password**. **Custom** - Select this if your repository requires custom authentication schemes and requires you to pass credentials as HTTP headers. Enter **Key** and **Value**. Key is the HTTP header name and Value is the HTTP header value.
From Local Machine	**Select File**	Browse and select the downloaded package from your local machine.

> **Note:**
> When you select the **Upgrade and Patch Package** option, you can either specify the download URLs of both upgrade and patch packages or you can upload the downloaded packages from your local machine to cluster.

4. Depending on your package selection, you can click:
    1. **Get Only** to download the package.
    2. **Get and Upgrade** to download and immediately upgrade the package.
    3. **Get and Apply** to download and apply the patch.

Once the package is downloaded, you can perform the following actions by clicking ⋮ on the package:

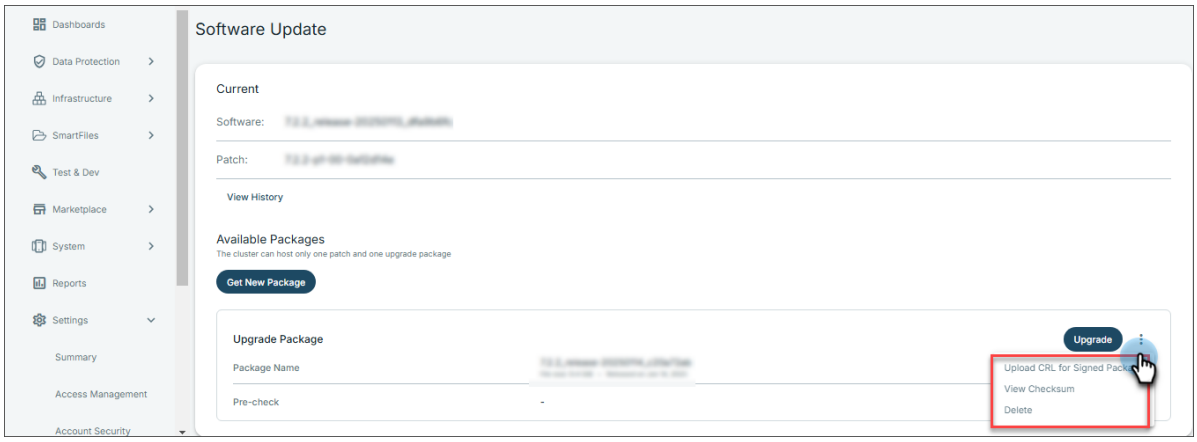

- **Upload CRL for Signed Package** - To upload Certificate Revocation List (CRL) for signed packages. CRLs are used to validate whether the certificate used for signing the package has been revoked after the package was downloaded to the Cohesity cluster.
- **View Checksum** - To view and verify the checksum of the package.
- **Delete** - To delete the package.

Run Pre-check

The pre-check option is available after the package download. It is important to run the pre-checks once the packages have been downloaded. Doing so in advance allows you to identify and resolve any issues, ensuring a smooth upgrade and patching process without any interruptions.

To run the pre-check on the downloaded package, click **Run**.

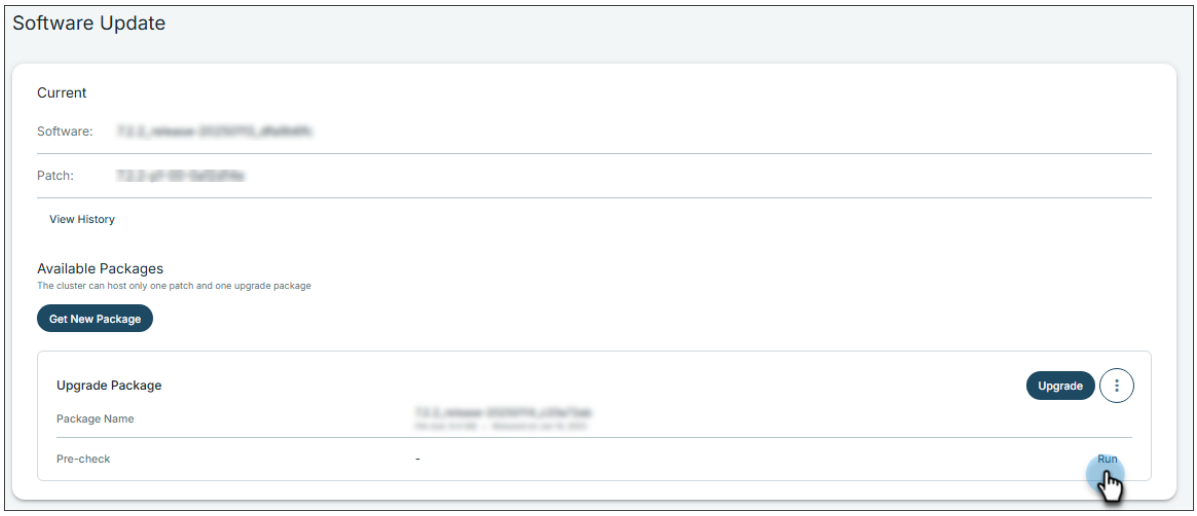

The pre-check will start and you can see the pre-check details by clicking **View Details**.

> **Note:** The cluster can host only one patch package and one upgrade package.

Upgrade Cluster

To upgrade the cluster, you can select the upgrade package and initiate the upgrade process after performing the pre-check. You can track the upgrade progress for all nodes until the operation is completed.

Perform the following steps to upgrade the cluster:

1. On the **Software Update** page, click **Upgrade**.

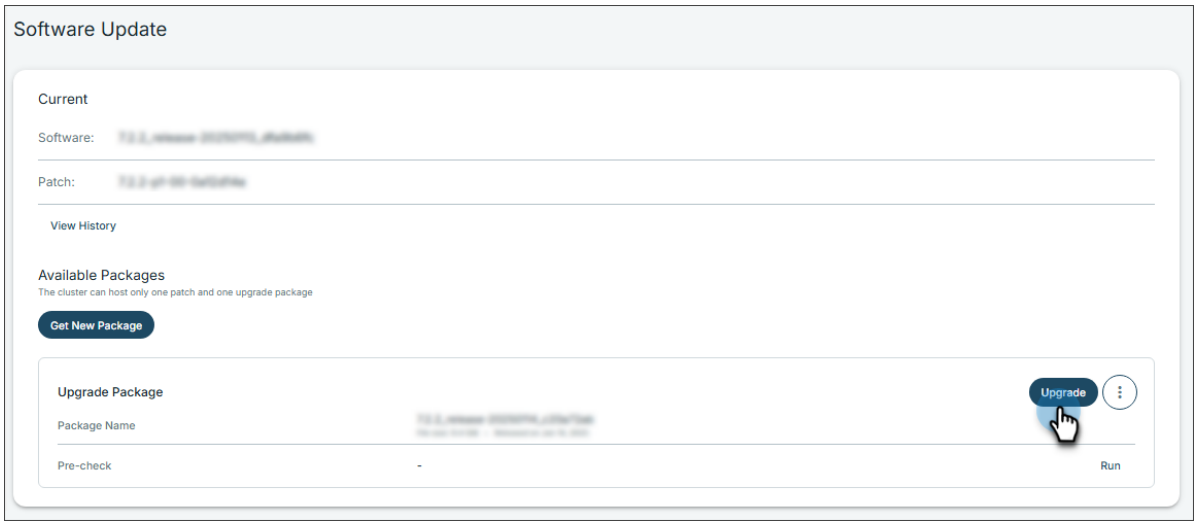

   The **Upgrade Review** dialog appears.

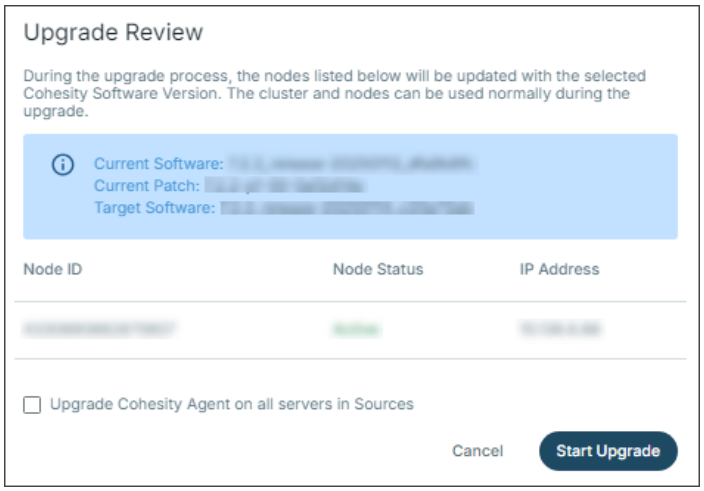

2. After reviewing the upgrade information, click **Start Upgrade**.

   The **Ok to Upgrade?** dialog appears.

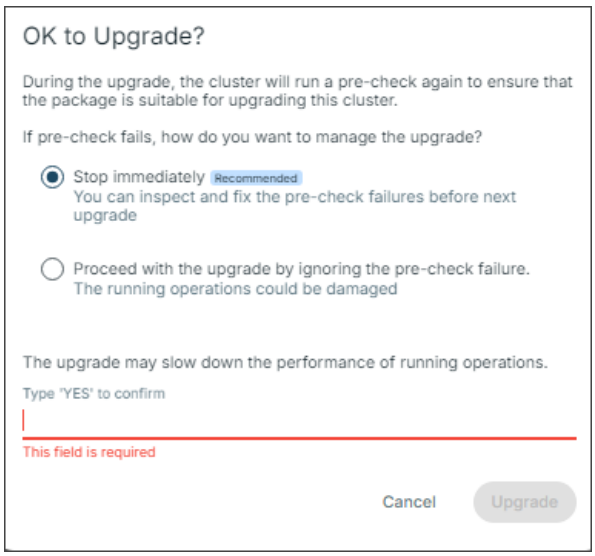

3. If the pre-check fails, you can either:
   1. **Stop immediately** to stop the upgrade process (recommended). This allows you to review and address the pre-check failures before retrying the upgrade.
   2. **Proceed with the upgrade by ignoring pre-check failure** to continue the upgrade process despite the pre-check failure.
4. Enter **Yes** in the text field and click **Upgrade** to confirm and start upgrading.

The upgrading process will begin. You can observe the upgrading progress for all the available nodes in the cluster.

Apply Patch

To apply a patch, select the patch package and begin the process after performing the pre-check. The patch will be applied across all nodes in the cluster, with progress tracked until

completion.

Perform the following steps to apply a patch:

1. On the **Software Update** page, click **Apply**.

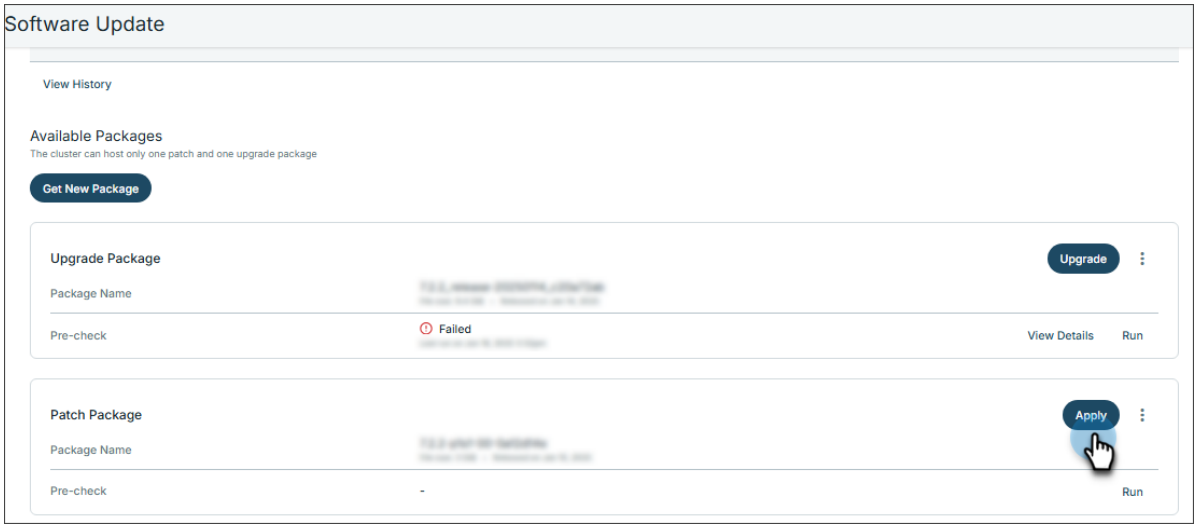

The **OK to apply?** dialog appears.

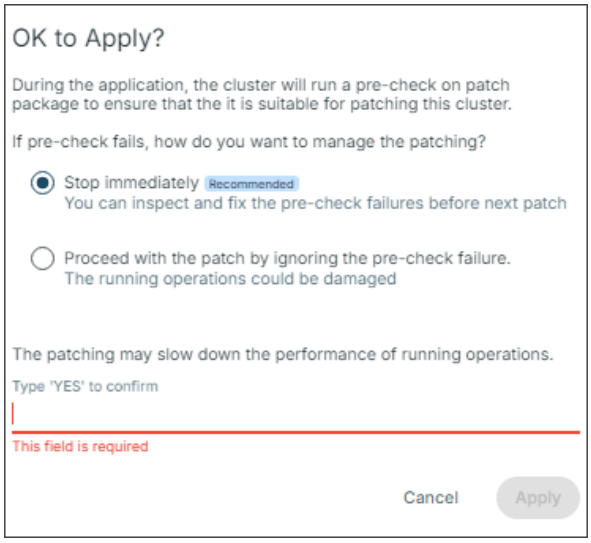

2. If the pre-check fails, you can either:
    1. **Stop immediately** to stop the patching process (recommended). This allows you to review and address the pre-check failures before retrying the patch.
    2. **Proceed with the patch by ignoring pre-check failure** to continue applying the patch despite the pre-check failure.
3. Enter **Yes** in the text field and click **Apply** to confirm and start applying the patch.

You can observe the patching progress for all the available nodes in the cluster. The patching process will begin after the pre-checks are completed on all nodes, starting with the node having the lowest ID.

To cancel the patching, click **Cancel**.

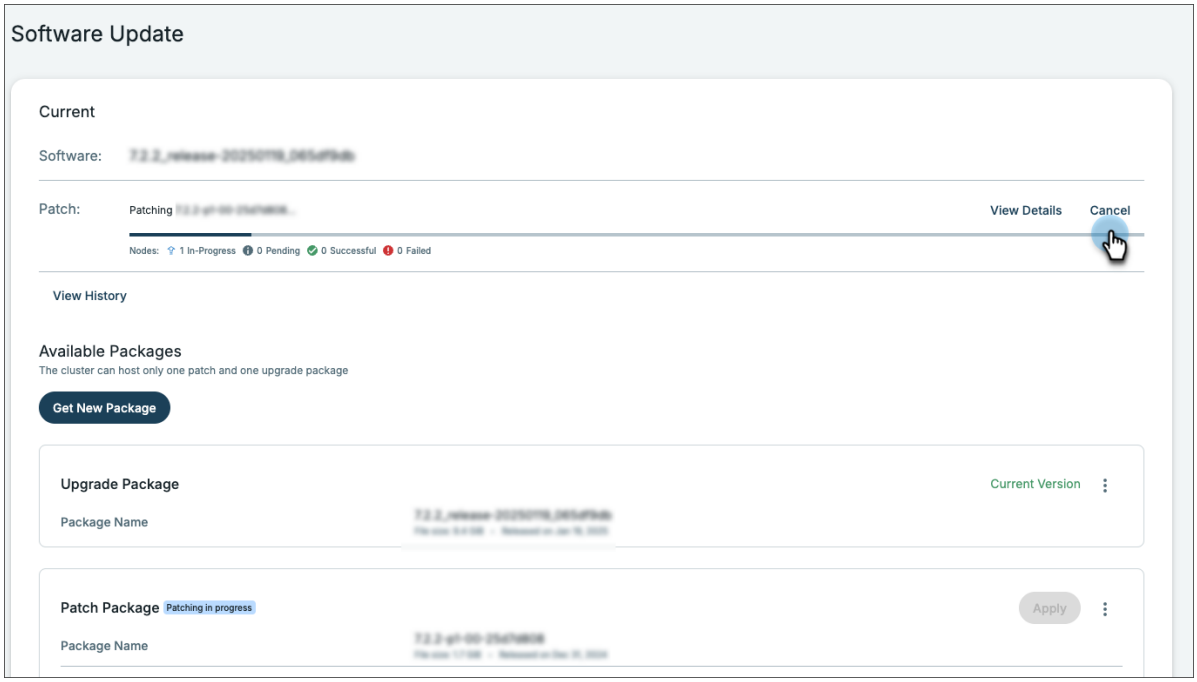

The **Cancel Patching?** dialog appears.

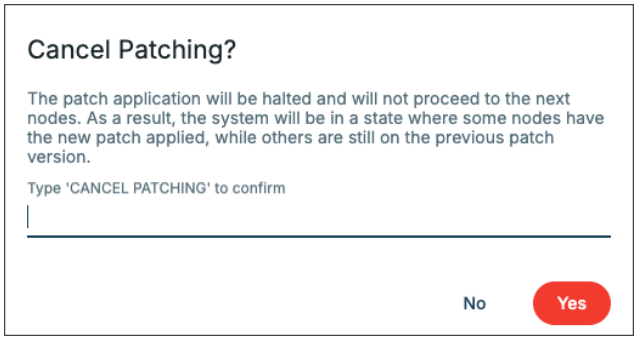

Enter **Cancel Patching** in the text field and click **Yes** to cancel the patching or **No** to continue the patching.

Upgrade with Patch

Upgrading with a patch allows you to apply both an upgrade and patch in one action. After downloading the upgrade and patch packages, you can proceed with the upgrade, including the patch, following a pre-check process.

Perform the following steps to upgrade with the available patch:

1. On the **Software Update** page, click **Upgrade**.

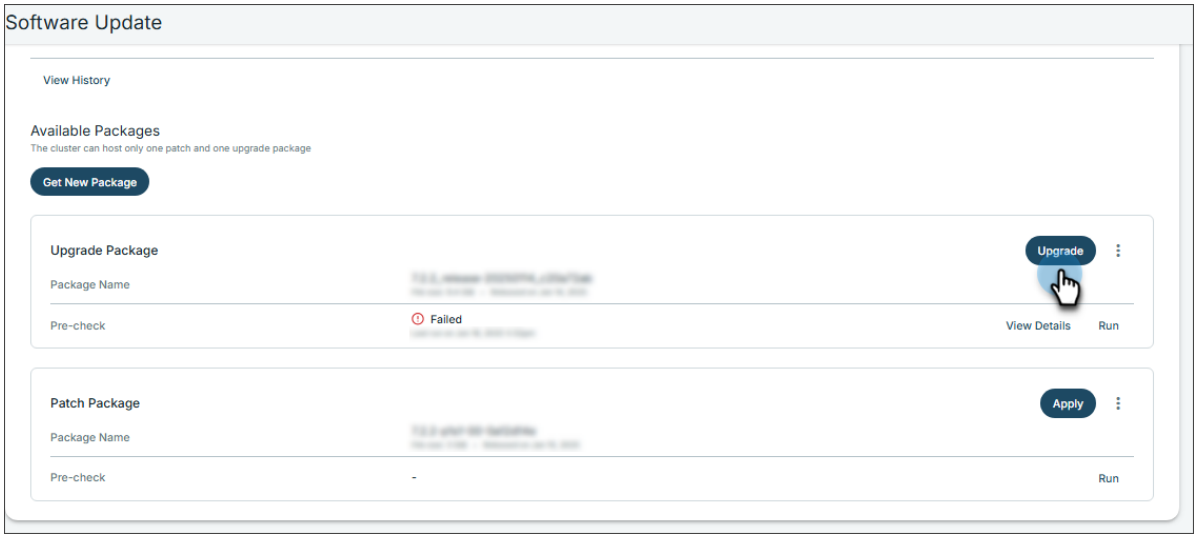

The **Upgrade with the Available Patch?** dialog appears.

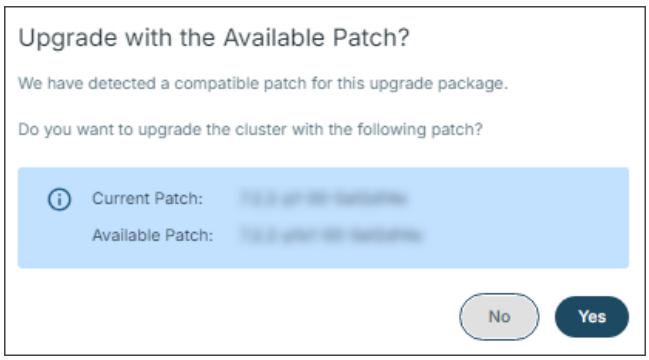

2. Click **Yes** to upgrade the cluster with the available patch.

   The **Upgrade Review** dialog appears.

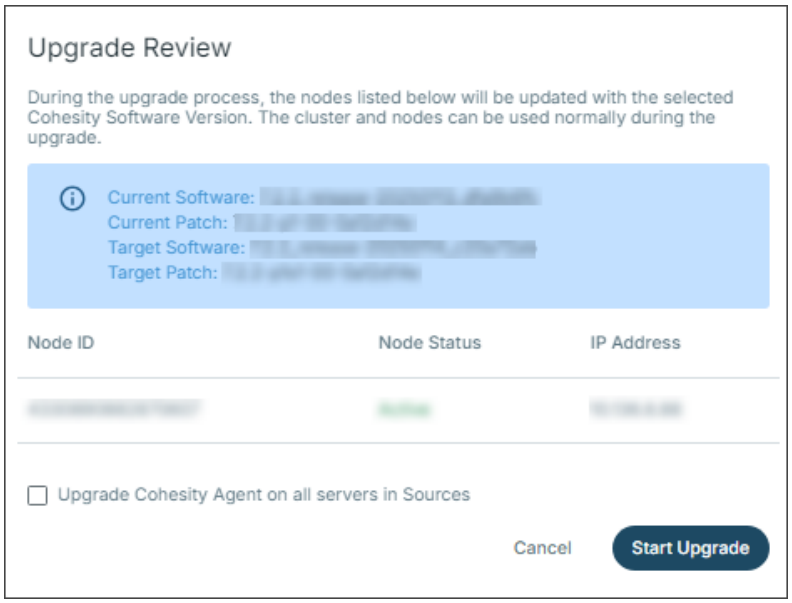

3. After reviewing the upgrade information, click **Start Upgrade**.

   The **Ok to Upgrade?** dialog appears.

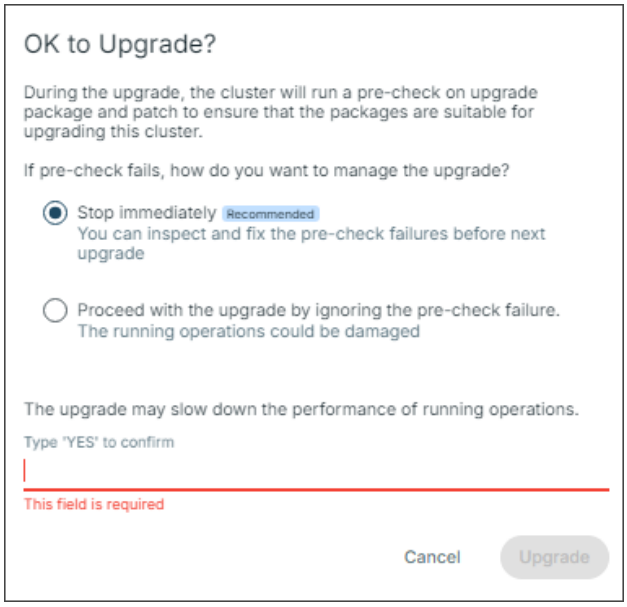

4. If the pre-check fails, you can either:
   1. **Stop immediately** to stop the upgrade process (recommended). This allows you to review and address the pre-check failures before retrying the upgrade.
   2. **Proceed with the upgrade by ignoring pre-check failure** to continue the upgrade process despite the pre-check failure.
5. Enter **Yes** in the text field and click **Upgrade** to confirm and start upgrading.

The upgrading process will begin. You can observe the upgrading progress for all the available nodes in the cluster.

To monitor upgrading progress, click **View Details**.

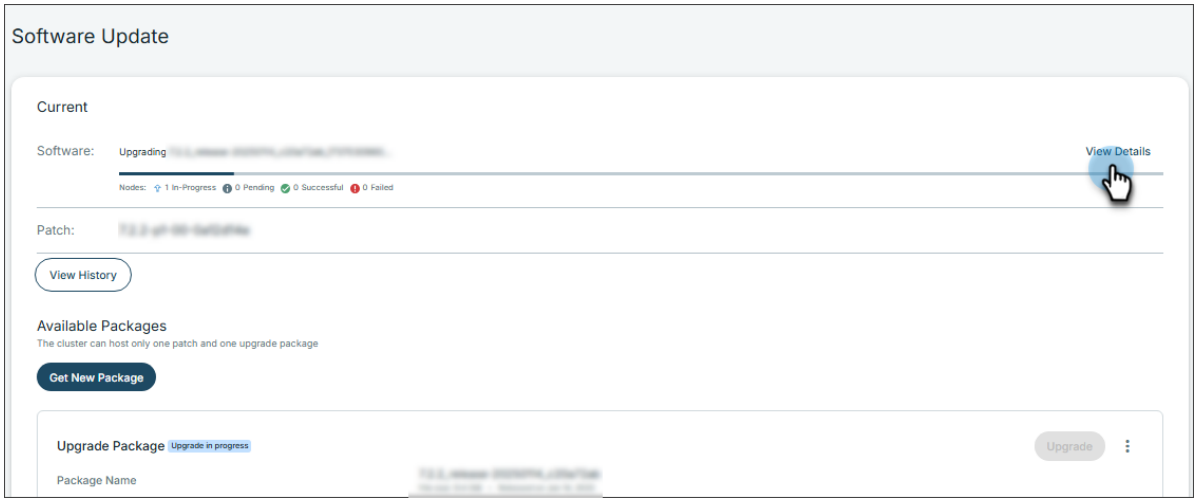

The **Upgrading the Cluster** page appears, showing the following information:

- Target version for upgrading
- Node Status
    - In-progress
    - Pending
    - Successful
    - Failed
- Node IP Address
- Upgrading Duration
- Upgrading Completion Time

To view the upgrade and patch history, click **View History**.

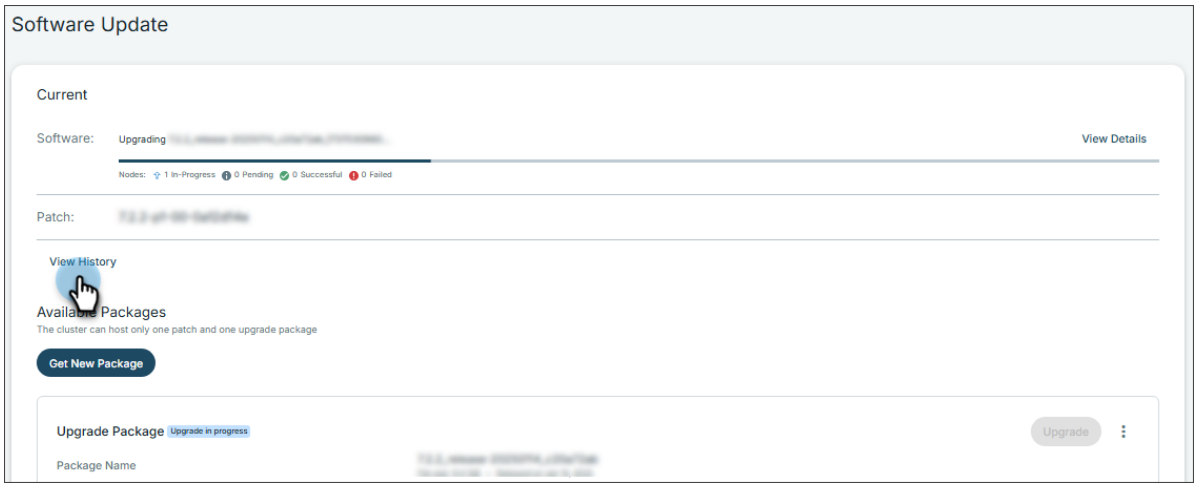

The **View History** dialog appears, where you can see the following information:

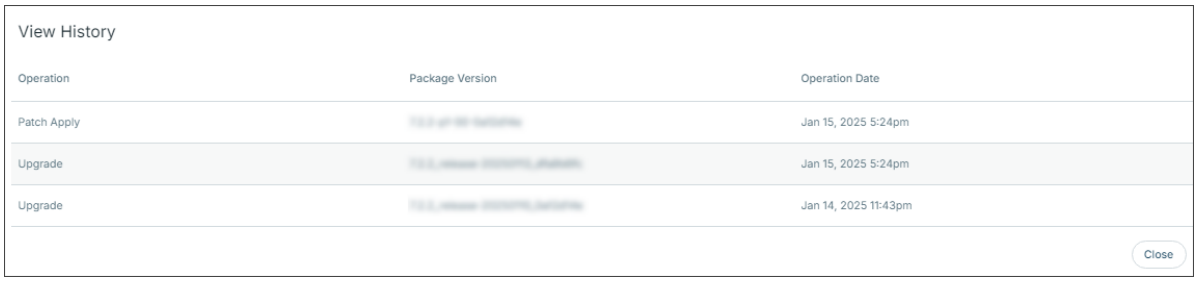

- **Operation** - Displays the type of operation you performed.
- **Package Version** - Displays the package version used during the operation.
- **Operation Date** - Displays the date and time when the operation was completed.

Cancel Upgrade

Cancel Upgrade in a Single Node Cluster

Run the following command to cancel the upgrade in a single node Cohesity cluster:

```
iris_cli cluster upgrade cancel=true
```

Sample Output

```
Username: admin
```

```
'admin' Password:password123
```

```
Success: Aborted upgrade.
```

When running the CLI command to cancel the upgrade of a Single Node Cohesity Cluster, all upgrades, including ongoing upgrades on its nodes, are canceled.

Cancel Upgrade in a Multinode Cluster

Run the following command to cancel the upgrade in a multinode Cohesity cluster:

```
iris_cli cluster upgrade cancel=true
```

Sample Output

```
Username: admin
```

```
'admin' Password:password123
```

```
Success: upgrade cancellation is requested.
```

If upgrade cancellation is requested when an upgrade is in progress on one of its nodes, the Cohesity cluster allows the upgrade to complete on that node. Once the ongoing upgrade is completed, subsequent upgrades on its remaining nodes are aborted. While the approach improves cluster stability, it may lead to nodes on different versions. You can resolve this by upgrading the remaining nodes to ensure all the nodes are of the same version.

Verify an Upgrade is Canceled

To verify the upgrade is canceled in a single or multinode Cohesity cluster, Cohesity recommends that you run the `cluster status` command:

- To check the current version of the nodes in the output.
- To verify there are no ongoing upgrades on the cluster by ensuring that `CLUSTER ACTIVE OPERATION` is empty.

```
admin> cluster status
```

Sample Output

```
Username: admin
'admin' Password:Password123
CLUSTER ID                    : 1111111
CLUSTER NAME                  : LocalCluster
SERVICE STATE SYNC            : DONE
CLUSTER ACTIVE OPERATION      :
CLUSTER HEAL STATUS           : NORMAL
```

If a node (irrespective of whether it's the first or last node of that cluster) becomes inoperable during an upgrade, cancel the upgrade, delete the node, and restart the upgrade.

Power Off and Power On a Cohesity Cluster

This topic describes the following:

- How to power off a Cohesity cluster or Node.
- How to power on a Cohesity cluster or Node.

Power off a Cohesity cluster or a Cohesity Node

1. Login to the Cohesity CLI.

   If your default username and password is not admin, then login to Cohesity CLI using the `iris_cli -username=XXXX` command. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

2. Stop all IOs from the Clients, and wait for 2-3 minutes.

   Skip this step when powering off a single node on a Cohesity cluster.

3. Run the following commands:

   To stop the Cohesity cluster, use the command:

```
cohesity>cluster stop
```

To stop a single node on a Cohesity cluster, on the Linux CLI, run the:

`node_stop.sh script`

The `node_stop.sh script` is located at: `/home/cohesity/bin/node_stop.sh`

This command will disconnect the secure shell and come back to normal in a few minutes.

> **Note:** Enable Host Shell to run the node_stop.sh command as Cohesity restricts custom scripts to the host operating system shell. For more information, see Enable Host Shell Access.

4. Wait for 5 minutes for the services to stop and run the status command to the view service status.

   In case of aCohesity cluster, use command:

   ```
   cohesity>cluster status
   ```

   In case of a Cohesity node, use command:

   ```
   cohesity>node status
   ```

   For Cohesity cluster, output similar to the following is expected:

   ```
   NODE ID                        : 101954678100000
   NODE IPS                       : <Node IP>
   SOFTWARE VERSION               : <Software Version>
   ACTIVE OPERATION               :
   SERVICE NAME                   :
   alerts                      :
   apollo                      :
   athena                      :
   atom                        :
   bifrost                     :
   bifrost_broker              :
   bridge                      :
   bridge_proxy                :
   compass                     :
   eagle_agent                 :
   elrond                      :
   etl_server                  :
   gandalf                     : 25981, 25988
   groot                       :
   ```

```
heimdall                    :
icebox                      :
iris                        : 4600, 13237
iris_proxy                  : 982, 14524
janus                       :
keychain                    :
librarian                   :
logwatcher                  : 5403
magneto                     :
newscribe                   :
nexus                       : 2757
nexus_proxy                 : 27979, 27982
nfs_proxy                   :
node_exporter               : 5398
patch                       :
pushclient                  : 5399
rtclient                    : 24050, 24066
smb2_proxy                  :
smb_proxy                   :
stats                       :
statscollector              : 5396
storage_proxy               :
throttler                   :
tricorder                   :
vault_proxy                 :
yoda                        :
```

For a Cohesity node, output similar to the following is expected:

```
[support@hpe-apollo-4510-2m282800yf-node-1 ~]$ iris_cli node status
```

```
E0726 10:47:16.470086 32855 handle_rest.go:676] Client.Do Error:Post
https://127.0.0.1:443/login: dial tcp 127.0.0.1:443: connect:
connection refused after making /login request.
```

```
Unable to connect to server 127.0.0.1:443.
```

> **Note:** The numbers for each Cohesity cluster service may differ for each node or Cohesity cluster.

5. Run the following command to stop the cluster services and shut down the Linux kernel.

```
shutdown.sh
```

6. Based on the hardware or hypervisor, perform the following:

C2000 Series

Log in to the IPMI and power off each node by performing the steps below:

1. Click the **Remote Control** tab.
2. In the left-hand pane, select **Server Power Control** > **Power OFF Server** or **Graceful Shutdown**> **Perform Action**.

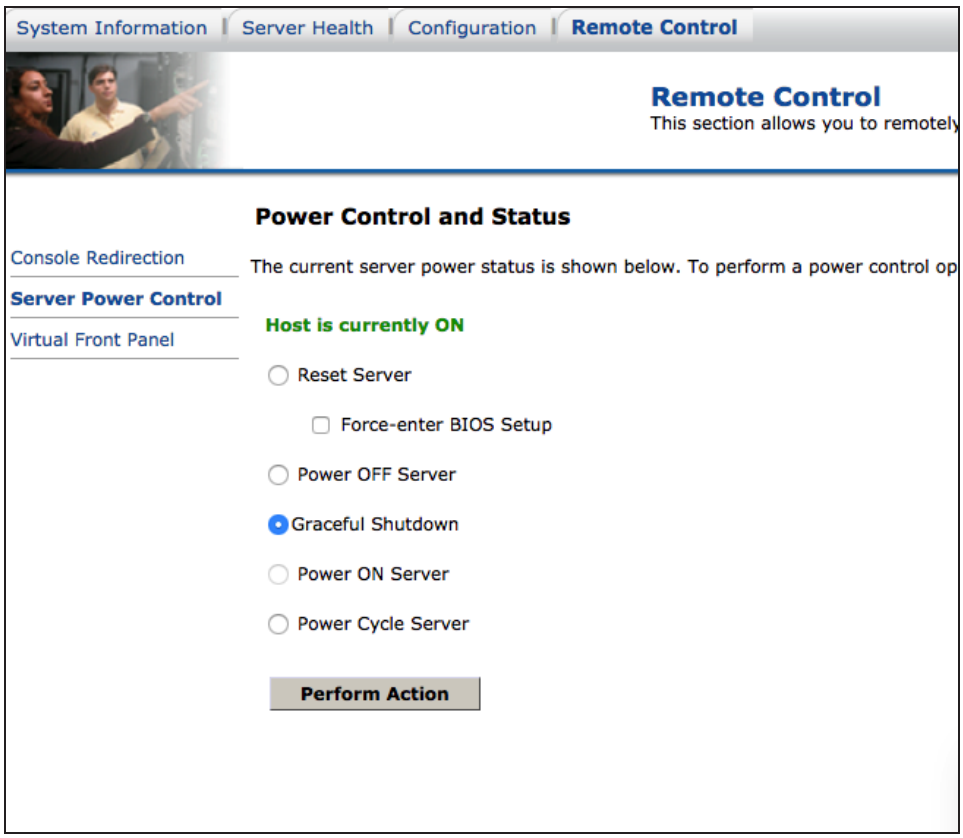

C3500

Log in to the IPMI and power off each node by performing the steps below:

1. Click the **Remote Control** tab.
2. In the left-hand pane, select **Power Control** > **Power Off Server -**

Immediate or **Power Off Server - Orderly Shutdown** > **Perform Action**.

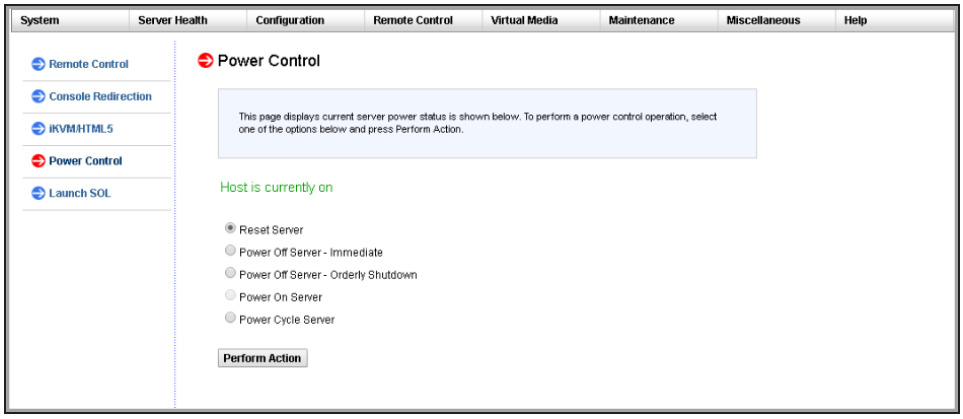

C4000 Series

Log in to the IPMI and power off each node by performing the steps below:

1. In the left-hand pane, select **System Power Control**.
2. In the **Remote Power Control** section, from the **Operation** drop-down, select **Power Off**.
3. Click **Perform**.

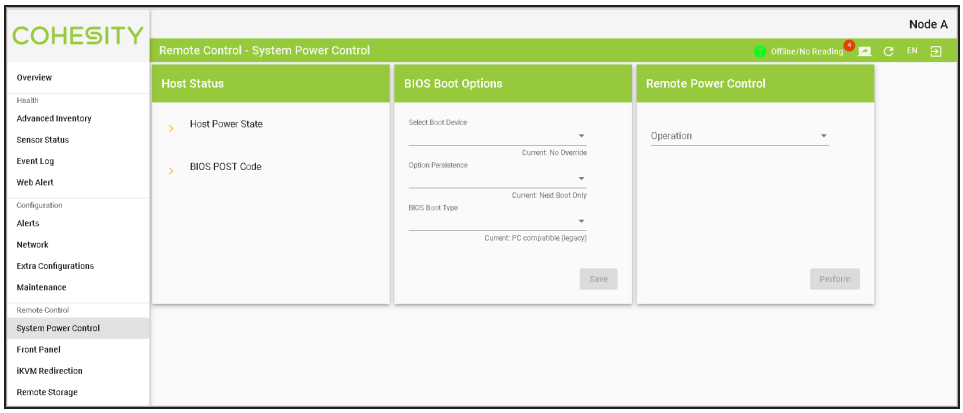

C5000

Log in to the IPMI and power off each node by performing the steps below:

1. In the left-hand pane, select **System Power Control**.
2. In the **Remote Power Control** section, from the **Operation** drop-down, select **Power Off**.

3. Click **Perform**.

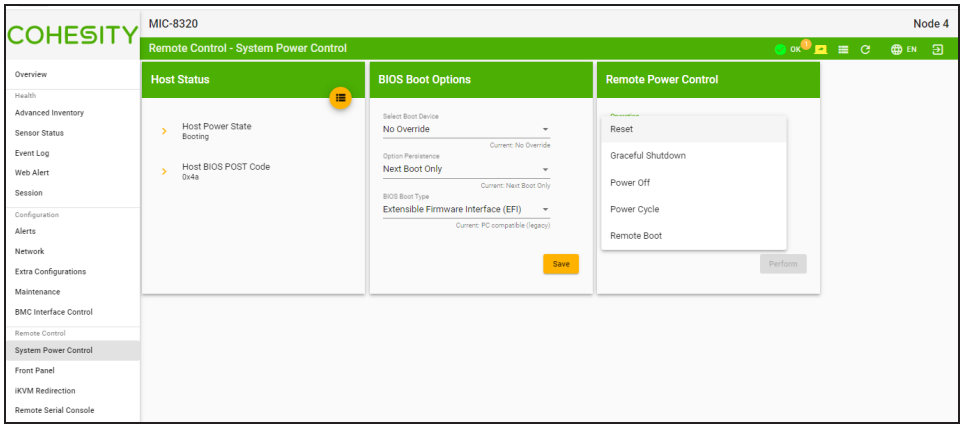

C6000 Series

Log in to the IPMI and power off each node by performing the steps below:

1. In the left-hand pane, select **Power Control**.
2. Under **Power Control** > **Power Actions**, select either **Power Off Server - Immediate** or **Power Off Server - Orderly Shutdown.**
3. Click **Perform Action**.

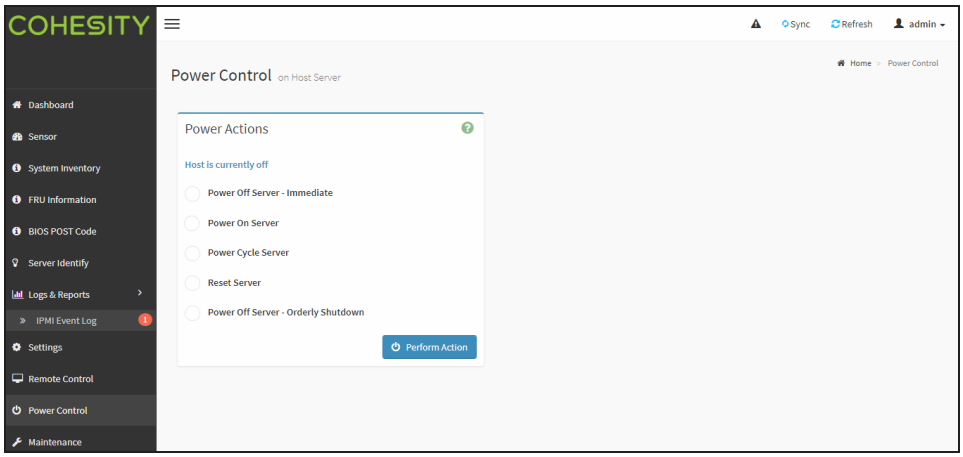

CX8000 Series

Log in to the IPMI and power off each node by performing the steps below:

1. In the left-hand pane, select **Power Control**.
2. Under **Power Control** > **Power Actions**, select **Power Off**.

3. Click **Perform Action**.

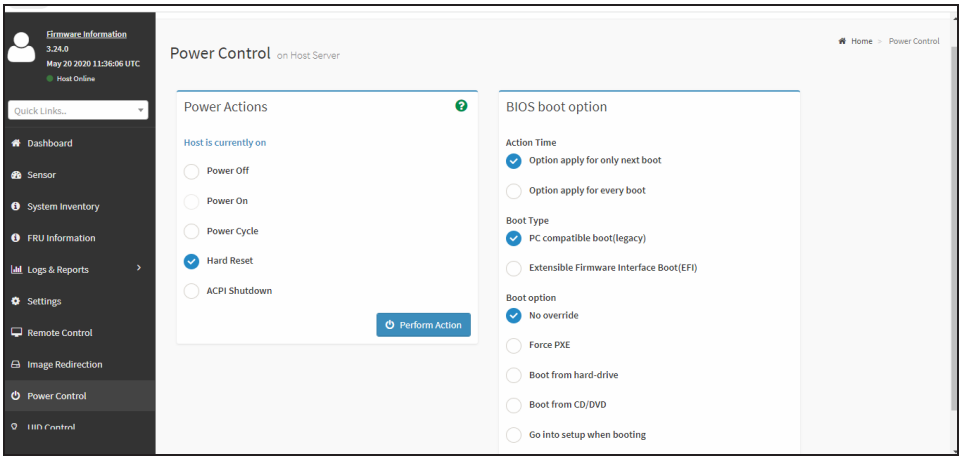

Virtual Edition Cluster

1. Log into your Hypervisor's Management Console to power on the node. For details on how to do this, see the documentation of your Hypervisor.

   For example, if you are using a VMware VM:

   1. Log into the vCenter Server using the vSphere Web Client.
   2. Follow these steps to power on each VM in the Cohesity cluster:

      - From the **Hosts and Clusters** tab, right-click the VM and select **Power** > **Power Off**.

        (or)

      - Click the **Summary** tab of the VM, and from the **ACTIONS** menu, select **Power** >**Power Off**.

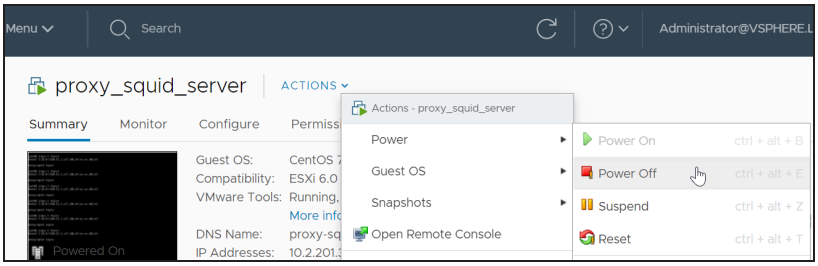

Cloud Edition Cluster

You can stop your cloud instance using the Cloud Management Console of your Cloud Service Provider. For details on how to stop a cloud instance, using the:

- **AWS**: Amazon EC2 console or the AWS CLI.
- **Azure**: Azure portal.
- **GCP**: Google Cloud console or the gcloud CLI.

Power on a Cohesity cluster or a Cohesity Node

1. Based on the hardware, perform the following:

   C2000 Series

   Log in to IPMI and power on each node by performing the steps below:

   1. Click the **Remote Control** tab.
   2. In the left-hand pane, select **Server Power Control** > **Power ON Server** > **Perform Action**.

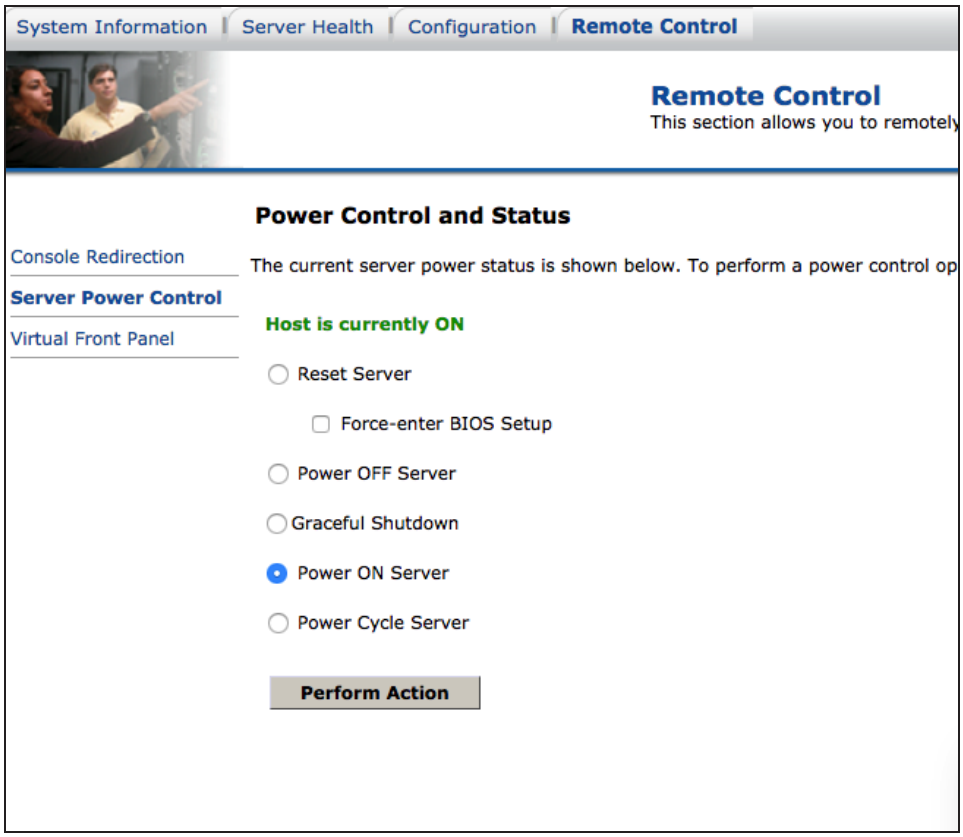

   C3500

   Log in to the IPMI and power on each node by performing the steps below:

   1. Click the **Remote Control** tab.
   2. Select **Power Control** > **Power On Server** > **Perform Action**.

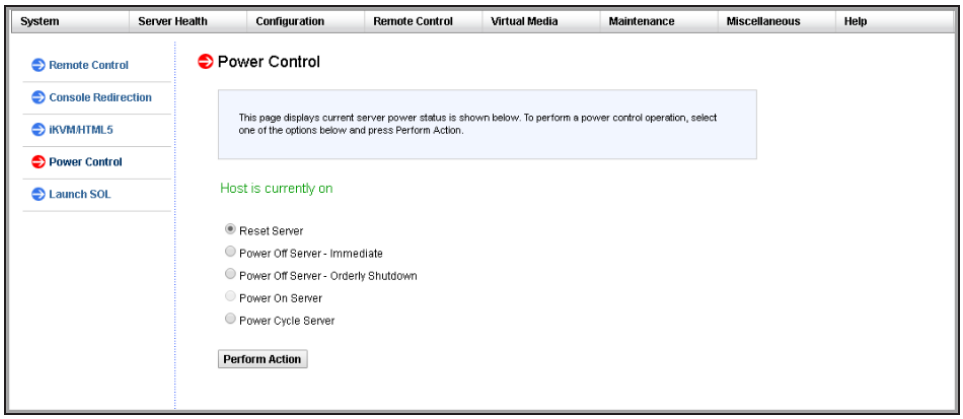

C4000 Series

Log in to the IPMI and power on each node by performing the steps below:

1. In the left-hand pane, select **System Power Control**.
2. In the **Remote Power Control** section, from the **Operation** drop-down, select **Power On**.
3. Click **Perform**.

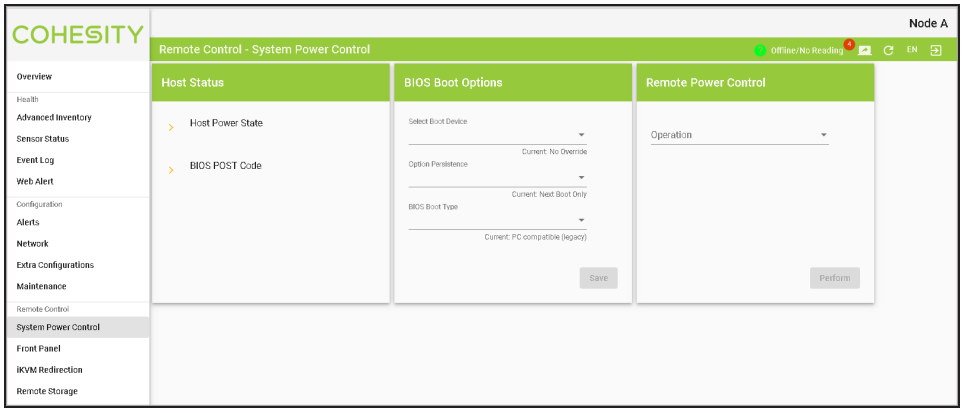

C5000

Log in to the IPMI and power on each node by performing the steps below:

1. In the left-hand pane, select **System Power Control**.
2. In the **Remote Power Control** section, from the **Operation** drop-down, select **Power On**.

3. Click **Perform**.

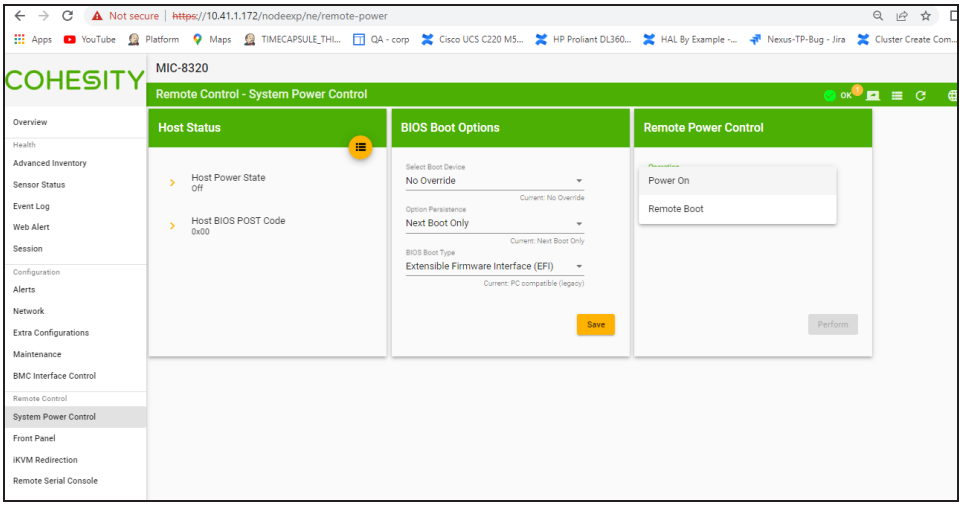

C6000 Series

Log in to the IPMI and power off each node by performing the steps below:

1. In the left-hand pane, select **Power Control**.
2. Under **Power Control** > **Power Actions**, select **Power On Server.**
3. Click **Perform Action**.

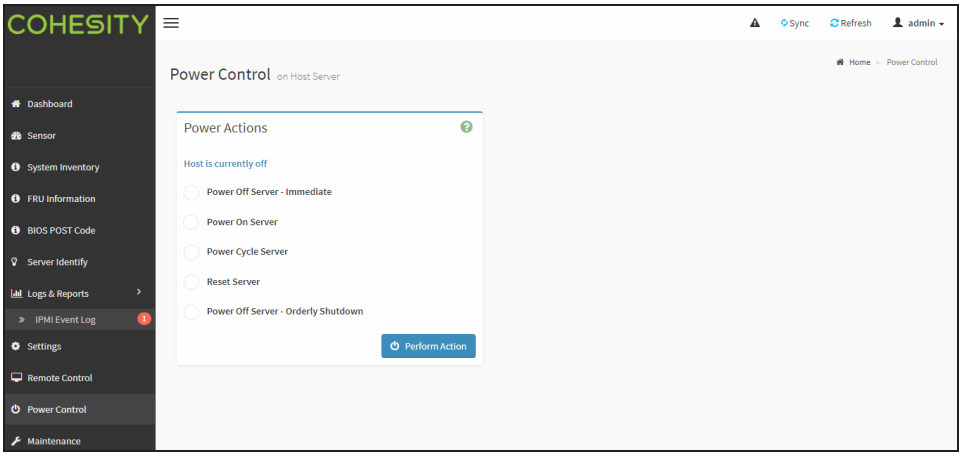

CX8000 Series

Log in to the IPMI and power off each node by performing the steps below:

1. In the left-hand pane, select **Power Control**.
2. Under **Power Control** > **Power Actions**, select **Power On**.

3. Click **Perform Action**.

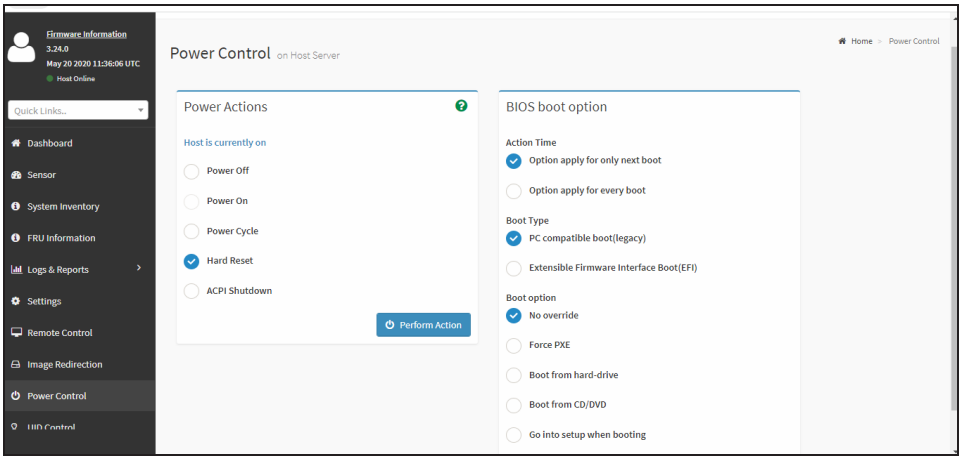

Virtual Edition Cluster

Log into your Hypervisor's Management Console to power on the node. For details on how to do this, see the documentation of your Hypervisor.

For example, if you are using a VMware VM:

1. Log into the vCenter Server using the vSphere Web Client.
2. Follow these steps to power on each VM in the Cohesity cluster:
   - From the **Hosts and Clusters** tab, right-click the VM and select **Power** > **Power On**.

     (or)
   - Click the **Summary** tab of the VM, and from the **ACTIONS** menu, select **Power** > **Power On**.

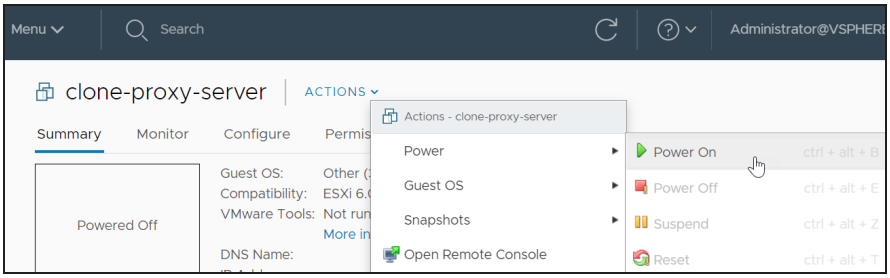

3. Repeat step 2 for all the nodes in the Cohesity cluster.

Cloud Edition Cluster

You can start your cloud instance using the Cloud Management Console of your Cloud Service Provider. For details on how to start a cloud instance, using the:

- **AWS**: Amazon EC2 console or the AWS CLI.
- **Azure**: Azure portal.

- **GCP**: Google Cloud console or the gcloud CLI.

2. Login to the Cohesity CLI.

   If your default user name and password are not admin, then login to Cohesity CLI using the `iris_cli -username=XXXX` command. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

3. Run the following command:

   To start a Cohesity cluster, use the command:

   ```
   cohesity>cluster start
   ```

   To start a single node on a Cohesity cluster, on the Linux CLI, run the:

   `node_start.sh script`

   The `node_start.sh script` is located at: `/home/cohesity/bin/node_start.sh`

4. Wait for 5 minutes for the services to start and run the status command to the view service status.

   In case of a Cohesity cluster, use command:

   ```
   cohesity>cluster status
   ```

   In case of a Cohesity node, use command:

   ```
   cohesity>node status
   ```

   For Cohesity cluster, output similar to the following is expected:

   ```
   NODE ID                        : 101954678100000
   NODE IPS                       : <Node IP>
   SOFTWARE VERSION               : <Software Version>
   ACTIVE OPERATION               :
   SERVICE NAME                   :
   alerts                      : 2927, 2936
   apollo                      : 2928, 3015
   athena                      : 17327, 17347
   atom                        : 17275, 17319
   bifrost                     : 3812, 11266
   bifrost_broker              : 11399, 11725
   bridge                      : 6382, 6686
   bridge_proxy                : 17450, 18200
   compass                     : 11421, 11480
   eagle_agent                 : 5269, 11234
   elrond                      : 11222, 11242
   etl_server                  : 10997, 24249
   gandalf                     : 25981, 25988
   groot                       : 1255, 1309
   ```

```
heimdall                    : 1214, 1242
icebox                      : 11470, 12072
iris                        : 4600, 13237
iris_proxy                  : 982, 14524
janus                       : 13425
keychain                    : 21531, 21563
librarian                   : 28018, 28021
logwatcher                  : 5403
magneto                     : 24228, 24254
newscribe                   : 11291, 11315
nexus                       : 2757
nexus_proxy                 : 27979, 27982
nfs_proxy                   : 21514, 21539
node_exporter               : 5398
patch                       : 1535, 21512
pushclient                  : 5399
rtclient                    : 24050, 24066
smb2_proxy                  : 21516, 21538
smb_proxy                   : 21533, 21577
stats                       : 2926, 2931
statscollector              : 5396
storage_proxy               : 21544, 21897
throttler                   : 11235, 11313
tricorder                   : 11237
vault_proxy                 : 21513, 21721
yoda                        : 21054, 21063
```

For the Cohesity node, output similar to the following is expected:

```
NODE ID                        : 101954678100000
NODE IPS                       : <Node IP>
NODE IN CLUSTER                : true
CLUSTER ID                     : 586212744040480
CLUSTER INCARNATION ID         : 1658841374325
SOFTWARE VERSION               : <Software Version>
LAST UPGRADED TIME             : <date, time>
NODE UPTIME                    : up 1 day, 20 hours, 22 minutes
ACTIVE OPERATION               :
REMOVAL STATE                  :
REMOVAL REASON                 :
APP NODE                       : false
SERVICE NAME                   :
alerts                      : 2927, 2936
apollo                      : 2928, 3015
athena                      : 17327, 17347
```

```
atom                    : 17275, 17319
bifrost                 : 3812, 11266
bifrost_broker          : 11399, 11725
bridge                  : 6382, 6686
bridge_proxy            : 17450, 18200
compass                 : 11421, 11480
eagle_agent             : 5269, 11234
elrond                  : 11222, 11242
etl_server              : 10997, 24249
gandalf                 : 25981, 25988
groot                   : 1255, 1309
heimdall                : 1214, 1242
icebox                  : 11470, 12072
iris                    : 4600, 13237
iris_proxy              : 982, 14524
janus                   : 13425
keychain                : 21531, 21563
librarian               : 28018, 28021
logwatcher              : 5403
magneto                 : 24228, 24254
newscribe               : 11291, 11315
nexus                   : 2757
nexus_proxy             : 27979, 27982
nfs_proxy               : 21514, 21539
node_exporter           : 5398
patch                   : 1535, 21512
pushclient              : 5399
rtclient                : 24050, 24066
smb2_proxy              : 21516, 21538
smb_proxy               : 21533, 21577
stats                   : 2926, 2931
statscollector          : 5396
storage_proxy           : 21544, 21897
throttler               : 11235, 11313
tricorder               : 11237
vault_proxy             : 21513, 21721
yoda                    : 21054, 21063
```

> **Note:** The numbers for each Cohesity cluster service may differ for each node or Cohesity cluster.

A cluster can take 15 to 20 minutes to become fully operational. After powering on, a healthy node should show **NODE IN CLUSTER** set to true, a positive **NODE UPTIME**

value, and empty values for both **REMOVAL STATE** and **ACTIVE OPERATION**, indicating that the node is neither being removed nor performing another operation. Additionally, most services should display two Process IDs (PIDs), which indicates normal operation. If a service shows one PID, it indicates that the service is still starting. If no PIDs are displayed, it means the service is not running and you should contact Cohesity Support.

Search the Cluster and Take Action

You can search the Cohesity cluster for objects, such as VMs, storage volumes and databases, refine the search results and then perform actions on the objects. You can find and protect unprotected objects, investigate failed Protection Groups, or recover or clone snapshots, all from the search results page. The actions launch the same familiar Cohesity workflow pages, but help is available for each page if you need it.

Example: Find and Protect with a Job

The following example demonstrates how to find unprotected VMs and protect them by creating a Protection Group.

1. Type `vm` in the search box at the top of the dashboard and press enter. You can search for any partial name or type an asterisk * wildcard.

   The Search page displays a list of objects that contain `vm` in the name.

2. Optionally, refine the list by selecting checkboxes in the left pane.
   a. Under **Type**, select the **Virtual Infrastructure** checkbox. Additional options are displayed based on the infrastructure you select.
   b. Under **Virtual Infrastructure Type**, select the **VMware** checkbox.
   c. Under **Status**, select the **Unprotected** checkbox.

   The list of objects automatically updates to display only unprotected VMware VMs.
3. Select a VM in the center pane.

   Details about the VM display in the right pane.
4. Click **Protect** and select **Protection Group**.
5. Click **New Job**.
6. The New Job page is displayed and you can create a job to protect the VM.

> **Tip:** For an explanation of the options on the New Job page, click the help icon (?) in the top right corner of the dashboard.

7. When you are finished configuring the Protection Group, click **Protect** to create the job and return to the Search page.

Example: Clone or Recover a Snapshot

To clone or recover a snapshot previously created by a Protection Group:

1. Type a partial job name or * in the search box at the top of the dashboard and press enter.

   The Search page is displayed with a list of objects.

2. Select the **Protected** check box in the left pane.
3. Select a job from the center pane.

   Details about the job are displayed in the right pane.

4. Click the actions icon ⋮ in the right pane and select **Clone** or **Recover**.

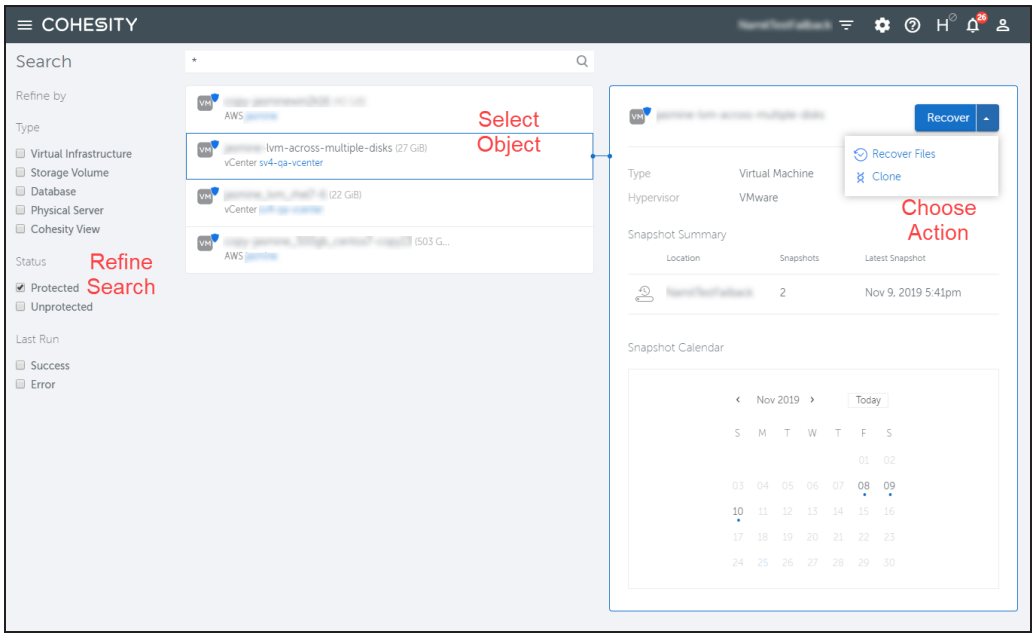

5. Depending on which action you click, the **Clone** or **Recover** page is displayed.

> **Tip:** For an explanation of the options on the **Clone** or **Recover** page, click the help icon (?) in the top right corner of the dashboard.

6. When you finish configuring the clone or recovery options, click **Finish** to run the task and return to the Search page.

Manage Storage Domains

From the Manage Storage Domains page, you can view, add and manage the Storage Domains that have been created on the Cohesity cluster. For detailed instructions, see Create or Edit Storage Domains.

The columns for the table displayed on the Manage Storage Domains page are described below.

Column Name	Description / Action
Storage Domain Name	Drill-down on this link to view the details and statics about the Storage Domain. You can also edit the Storage Domain from the linked page.
Organization	Drill-down on this link to get details of the organization that is associated with the Protection Group.
Physical Used	The size of the data physically stored on the Cohesity cluster, after the data has been reduced by change block tracking, compression and deduplication, including the disk space used for data resiliency (fault tolerance and erasure coding).
Physical Quota	The maximum physical capacity of the Storage Domain if its Physical Quota option is on. If the option is off, the field indicates `n/a`.
Logical Managed	The logical data managed, as computed by the Cohesity cluster without reduction by change-block tracking, compression and deduplication.
Redundancy	The Storage Domain's redundancy, either Redundancy Factor (RF2 or RF3) or Erasure Coding (EC) and the number of data stripes : number of failures to tolerate.
Deduplication	Indicates if Deduplication is on for the Storage Domain and when the deduplication occurs. If Deduplication is on, the Cohesity cluster eliminates duplicate blocks of repeating data stored on the Cohesity cluster, thus reducing the amount of storage space needed to store data. If Inline Deduplication is on, deduplication occurs when the Cohesity cluster writes the data to the Cohesity View. If Inline Deduplication is off, deduplication occurs after the Cohesity cluster writes the data to the Cohesity View.

Column Name	Description / Action
Compression	Indicates if Compression is on for the Storage Domain and when the compression occurs. If Compression is on, the Cohesity cluster compresses all the data stored in the Storage Domain, thus reducing the amount of storage space needed to store the data. If Inline Compression is on, compression occurs when the Cohesity cluster saves the blocks to the Cluster. If Inline Compression is off, compression occurs after the Cohesity cluster writes the data to the Cohesity cluster.
Encryption	Indicates if Encryption is on for the Storage Domain. **Note:** Encryption can be enabled at the cluster level only when creating the cluster and cannot be disabled or modified later. If Encryption is enabled at the Cluster level, all Storage Domains are encrypted. However, encryption can be enabled at the Storage Domain level after creating the cluster.
Cloud Tier	If Cloud Tier is on, the Cohesity cluster moves cold (rarely accessed) data to the designated External Target. Once Cloud Tier is on, it cannot be turned off.
Actions	To manage an existing Storage Domain, select one of the following options: For a Multitenancy Cohesity cluster, you can only modify the objects that belong to your organization. • **Edit** ( )–Configure the settings of the existing Storage Domain. For detailed instructions, see Create or Edit Storage Domains. • **Create View in this Storage Domain** ( )–Create a new View in this Storage Domain. For details instructions, see Create or Edit Storage Domains.

Create or Edit Storage Domains

A Storage Domain, formerly known as a View Box, is a named storage location on a Cluster. A Storage Domain defines the policy and frequency for deduplication among other configurations such as Compression, Erasure Coding, and Replication Factor. A Storage Domain contains Views. When you configure a Protection Group, you specify a Storage Domain, and during a Protection Group Run, the Cohesity cluster stores the Snapshots in that Storage Domain.

To create or edit a Storage Domain:

1. Navigate to **Settings > Summary** and select the **Storage Domains** tab.
2. On the Storage Domains page, complete one of the following actions:
   - Click **Add Storage Domain** located at the top right of the page.
   - Edit an existing Storage Domain by clicking the Actions Menu (⋮) and selecting **Edit** (✎).
3. **Storage Domain Name:** Enter a name for the Storage Domain.
4. **Associated External Target**: Select the external target to be associated with the Storage Domain. This associated external target stores all the data landing on this Storage Domain. Each Storage Domain can only be associated with one external target and vice versa. Once set, the associated external target cannot be changed.

   NOTE: This field is only available for GCP Next-Gen Cloud Edition.
5. **Authentication Provider:** You can map the Storage Domain to an AD domain if the Cohesity cluster, version 6.5.1d or higher, is AD domain-joined. Mapping to a specific AD domain means that users from only that AD domain have permission to view the Storage Domain's contents when accessing through SMB. By default, a Storage Domain is accessible to users from all joined AD domains.
6. **Deduplication:** Optionally toggle on. If Deduplication is on, the Cohesity cluster eliminates duplicate blocks of repeating data stored on the Cohesity cluster, thus reducing the amount of storage space needed to store data.

   > **Note:** Toggling **Deduplication** on for a non-Deduplication Storage Domain that uses Cloud Tier may result in a large number of reads from the Cloud Tier if Storage Domain data has already been sent to the Cloud Tier.

   Inline Deduplication: Optionally toggle on. If Inline Deduplication is on, deduplication occurs when the Cohesity cluster writes the data to the Cohesity View. If Inline Deduplication is off, deduplication occurs after the Cohesity cluster writes the data to the Cohesity View.
7. **Compression:** Optionally toggle on. If Compression is on, the Cohesity cluster compresses all the data stored in the Storage Domain, thus reducing the amount of storage space needed to store the data. If Compression is on and then turned off later, the current data on the Storage Domain is not decompressed, but any new data stored on the Storage Domain is not compressed.

   Inline Compression: Optionally toggle on. If Inline Compression is on, compression occurs when the Cohesity cluster saves the blocks to the Cluster. If Inline Compression is off, compression occurs after the Cohesity cluster writes the data to the Cohesity cluster.

> **Note:** If both Compression and Inline Deduplication are on, Cohesity recommends you turn **Inline Compression** on.

8. **Encryption:** Toggle to enable encryption for the Storage Domain. Encryption can be enabled only when you are creating the Storage Domain. Once the Storage Domain is created, the options to enable encryption or modify its encryption settings are disabled in the Cohesity cluster.

   By default, Cohesity uses **Internal KMS** for encryption. If encryption is enabled at the cluster level, then data within the Storage Domains is also encrypted, and you don't have to enable encryption for each Storage Domain. You can register an external KMS and select one of the following from the drop-down list:

   - AWS
   - KMIP Compliant

   For more information, see External Key Management Service.
9. Expand **Show Advanced Settings** to see the list of additional View settings.
10. **Physical Quota:** Optionally toggle on to limit the Storage Domain's physical capacity and then set the capacity. The capacity must be a whole number.

    Toggle **Alert** on to trigger an Alert when the Storage Domain reaches the specified threshold.
11. **Logical Quota:** Optionally toggle on to set a default logical capacity limit for Views created in this Storage Domain and set the capacity. Logical data is fully hydrated and expanded.

    Toggle **Alert** on to trigger an Alert when the View reaches the specified threshold.
12. **Cloud Tier:** Optionally toggle on. If Cloud Tier is on, the Cohesity cluster moves cold (rarely accessed) data to the designated External Target. Once Cloud Tier is on, it cannot be turned off.

    To turn Cloud Tier off, contact Cohesity Support to disable the feature.

    If Cloud Tier is on, select an External Target or register a new External Target.

    Set the coldness of data to tier.

    > **Note:** If the **Physical Quota** toggle is off, the cluster's default capacity and space utilization threshold are used for tiering.

    Optionally, toggle on **Override Cluster Threshold** and set the **Space Utilization Threshold** to activate cloud tiering for the Storage Domain.

13. **Fault Tolerance:** Select the number of hard drive, node, chassis, or rack failures that can occur before a Cluster Health Alert is triggered. The available choices are determined by the number of nodes, chassis, or racks in the Cluster. The Cluster's global **Custom Failure Domain** setting determines the failure domain (Node, Chassis or Rack) and the number of failures tolerated at the Storage Domain level. For more details about global cluster level failure domain settings, see Configure Cluster Settings.

    > **Note:** Fault Tolerance is not available for single node Cohesity cluster (ROBO). By default, the Storage Domain is configured with RF1.

    The following table provides the fault tolerance options you can select based on the failure domain set at the cluster level and the number of components available in the Cohesity cluster.

Minimum Failures Tolerated	Description
0D:0N	No hard disk failure or node failure can occur without triggering the alert. This option provides Redundancy Factor 1 (RF1). If you select this option, you are opting for no replication and indicating that you trust the underlying datastore medium's resiliency. This option is available only in the Virtual Edition clusters. **Note:** If you select this option, Erasure Coding will be disabled. If RF1 is enabled on a multi-node Virtual Edition Cluster, during the cluster upgrade, data will be temporarily unavailable until the associated node comes online.
1D:0N	Tolerate failure of 1 drive. No node failure is tolerated. This is only available on physical ROBO clusters, it provides RF2 like redundancy and is set by default.
1D:1N	Tolerate failure of 1 drive or 1 node. The available options to select are RF2 or EC x:1.

Minimum Failures Tolerated	Description
2D:1N	Tolerate failure of 2 drives or at least 1 node. The available options to select include RF3 and EC x:2. This option is available only if the Cluster has three or more nodes. Cohesity will always make the best effort to place each EC stripe unit to be on separate nodes when sufficient nodes are available. Only when sufficient nodes are not present the EC stripe units will be placed within the same node but on separate disks.
2D:2N	Tolerate failure of 2 drives or 2 nodes. This setting is only available if there are at least four nodes. With this setting, the available options are RF3 or EC x:2. Cohesity will always try to place each EC stripe unit on separate nodes when sufficient nodes are available. > **Note:** You cannot change the fault tolerance setting to 2D:2N if the metadata usage is >=50% of the available storage. For more information, see the Cohesity white paper: Cohesity Fault Tolerance – Data Integrity for Modern Web-scale Environments. > **Note:** The Erasure Coding (EC) setting is not recommended for Clusters with fewer than five nodes because the data cannot be healed after one Node failure. Your data will be at risk until the failed Node is replaced.
3D:1N	Tolerate failure of 3 drives or 1 node. This setting is only available if there are at least three nodes. With this setting, the available options are RF4, EC x:3, or EC x:4.
3D:2N	Tolerate failure of 3 drives or 2 nodes. This setting is only available if there are at least seven nodes. With this setting, the available options are RF4, EC x:3, or EC x:4.

Minimum Failures Tolerated	Description
0D:0C	No hard disk failure or chassis failure can occur without triggering the alert. This option provides Redundancy Factor 1 (RF1). If you select this option, you are opting for no replication and indicating that you trust the underlying datastore medium's resiliency. This option is available only in the Virtual Edition clusters. If you select this option, Erasure Coding will be disabled. If RF1 is enabled on a multi-node Virtual Edition Cluster, during the cluster upgrade, data will be temporarily unavailable until the associated chassis comes online.
1D:1C	Tolerate failure of 1 drive or 1 chassis. This setting is only available if there are at least three chassis. With this setting, the available options are RF2 or EC x:1.
2D:1C	Tolerate failure of 2 drives or 1 chassis. This setting is only available if there are at least three chassis. With this setting, the available options are RF3 or EC x:2.
2D:2C	Tolerate failure of 2 drives or 2 chassis. This setting is only available if there are at least four chassis. With this setting, the available options are RF3 or EC x:2.
3D:1C	Tolerate failure of 3 drives or 1 chassis. This setting is only available if there are at least three chassis. With this setting, the available options are RF4, EC x:3, or EC x:4.
3D:2C	Tolerate failure of 3 drives or 2 chassis. This setting is only available if there are at least seven chassis. With this setting, the available options are RF4, EC x:3, or EC x:4.
1D:1R	Tolerate failure of 1 drive or 1 rack. This setting is only available if there are at least three racks. With this setting, the available options are RF2 or EC x:1.
2D:1R	Tolerate failure of 2 drives or 1 rack. This setting is only available if there are at least three racks. With this setting, the available options are RF3 or EC x:2.
2D:2R	Tolerate failure of 2 drives or 2 racks. This setting is only available if there are at least four racks. With this setting, the available options are RF3 or EC x:2.
3D:1R	Tolerate failure of 3 drives or 1 rack. This setting is only available if there are at least three racks. With this setting, the available options are RF4, EC x:3, or EC x:4.
3D:2R	Tolerate failure of 3 drives or 2 racks. This setting is only available if there are at least seven racks. With this setting, the available options are RF4, EC x:3, or EC x:4.

The 2D in 2D:1N or 2D:2N has several effects at the same time:

- 2D allows selecting higher EC X:Y levels with fewer nodes. For example, an X+Y complete EC unit can be formed from n/2 nodes (assuming 2 disks per node taken for the single EC unit).
- Survive 2 simultaneous disk (HDD) failures OR 2 simultaneous node failures with 2D:2N with EC X:2. The disk failures are per cluster and not per node.
- Data stripe units from a single EC stripe are spread across as many nodes as possible. For example, if there are 10 nodes and EC 5:2 has been configured, the EC stripes will be spread across 7 separate nodes even if the protection level is 2D:1N. If there are 4 nodes, the EC 5:2 bits will be spread across 4 nodes with a 2D:1N protection level, with 3 of the 4 nodes having 2 stripe units each on separate disks.
- 2 node failures can be tolerated with EC X:2 if the number of nodes is at least X+2.
- Tolerate SSD (metadata) failures of up to 2.

For more information, see the Cohesity white paper: Cohesity Fault Tolerance — Data Integrity for Modern Web-scale Environments.

14. **Erasure Coding:** Toggle on for better fault tolerance and capacity utilization. You can also select an Erasure Coding setting from the drop-down list. The available choices are determined by the **Fault Tolerance** setting and the number of nodes, chassis, or racks in the Cohesity cluster. The settings indicate the ratio of the data stripes to the code stripes, expressed as x:y. For example, 2:1, 2:2, 3:1, 3.2, 4:1, 4:2, 4:3, 5:1, 5:2, or 5.3.

    Node Resilience Settings

    The following table shows possible EC settings that can be set for the required resiliency of hard drives and nodes. The table indicates the number of failures tolerated in disks *or* nodes at any EC or RF level. For example, if a 5-node Cohesity cluster is set to EC4:2, failures are tolerated in 2 disks or 1 node.

    For more details on Cohesity Erasure Coding, see the Cohesity-White-Paper-Fault-Tolerance-Data-Integrity.pdf.

    > **Note:** From Cohesity 6.6 onwards, you can use EC 8:2 for Azure, AWS, GCP, Imanis, MS SQL file-based, Physical Server, VMware, NAS, and SAN backup workloads.

Number of Nodes in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
3	1	1	2:1	RF2
	2	1	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
4	1	1	2:1, 3:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
5	1	1	2:1, 3:1, 4:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
6	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4

Number of Nodes in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
7	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
8	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
9	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Number of Nodes in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
10+	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Chassis Resilience Settings

The following table shows possible EC settings that can be set for the required resiliency of chassis.

> **Note:** From Cohesity 6.6 onwards, you can use EC 8:2 for Azure, AWS, GCP, Imanis, MS SQL file-based, Physical Server, VMware, NAS, and SAN backup workloads.

Number of Chassis in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
3	1	1	2:1	RF2
	2	1	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4

Number of Chassis in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
4	1	1	2:1, 3:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
5	1	1	2:1, 3:1, 4:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
6	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
7	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Number of Chassis in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
8	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
9	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
10+	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Rack Resilience Settings

The following table shows possible EC settings that can be set for the required resiliency of racks.

> **Note:** From Cohesity 6.6 onwards, you can use EC 8:2 for Azure, AWS, GCP, Imanis, MS SQL file-based, Physical Server, VMware, NAS, and SAN backup workloads.

Number of Racks in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Racks (R)	EC	RF (IF EC is Disabled)
3	1	1	2:1	RF2
	2	1	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
4	1	1	2:1, 3:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
5	1	1	2:1, 3:1, 4:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
6	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4

Number of Racks in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Racks (R)	EC	RF (IF EC is Disabled)
7	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
8	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
9	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
10+	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Post Processing Only: Toggle OFF to write erasure coding as data comes to the Cohesity cluster. When this option is ON, erasure coding is done post-process due to which data ingest performance rate is higher. .

ECX:4 is available only as a post processing option where the data is initially ingested as RF3 and then converted to X:4. Changing the viewbox configuration for EC X:Y (with Y < 4) to EC 12:4 or EC 16:4 is not permitted and not displayed in the viewbox editing page. However, you can switch from EC12:4 to EC16:4 and back, but the configuration only takes effect for any new data.

> **Note:** EC12:4 and EC16:4 are Early Access features. Contact Cohesity Support to enable the feature.

For more information, see the Cohesity white paper: Cohesity Fault Tolerance — Data Integrity for Modern Web-scale Environments.

15. Click **Create Storage Domain**.

Storage Domain Details

The **Storage Domain** details page has been redesigned to improve the user experience. This page provides granular details about the storage domain.

From the storage domain details page, you can do the following:

- Click the Actions menu (⋮) and select **Edit** to go to the **Edit Storage Domain** page.
- Click the Actions menu (⋮) and select **Create View in Storage Domain** to create a new view that uses this storage domain.

To view the storage domain details page:

1. Navigate to **Settings > Summary** and click the **Storage Domains** tab.
2. On the **Storage Domains** page, click a storage domain to navigate to the corresponding storage domain details page. This page displays the following details:

- **Storage**—Displays the current total quota, data consumed, space available, and the space saved owing to data deduplication and data compression:
- **Performance**—Displays the performance statistics of the selected storage domain:
    - **Throughput**—Amount of data transferred through external network interfaces in the cluster.
    - **IOPS**—Average number of read, write, or metadata operations serviced per second.
    - **Latency**—Average response time for read and write operations for views configured with TestAndDev High QoS policy.

- Partition details
- Authentication provider details
- Type of deduplication enabled on the storage domain—Inline or Post Process
- Type of compression enabled on the storage domain—Inline or Post Process
- Encryption status
- Quota and alert threshold details for the storage domain
- Quota and alert threshold details for views created in this storage domain
- Cloud Tier status

Delete Storage Domains

You can delete a Storage Domain to free up storage space if required. When deleting a Storage Domain, Cohesity clusters also delete the following entities associated with the Storage Domain:

- Protection groups and snapshots
- Recoveries in progress
- Views
- Dev/test clones

> **Note:** Cohesity cluster cannot delete views that have DataLock enabled or snapshots (and their owning protection groups) that have legal hold.

Before you Begin

Ensure the following prerequisites are met before you delete the Storage Domain:

- In a multitenant environment:
  - Only the Service Provider administrators can delete Storage Domains, not Tenants.
  - Service Provider administrators must remove all tenants associated with the Storage Domain before deleting the Storage Domains.
- If cloud archival is set up, delete the archive policy before deleting the Storage Domain.

To delete a Storage Domain:

1. Navigate to **Settings** > **Summary**.
2. Select the **Storage Domain** tab.

3. Click the Actions menu (⋮) of the Storage Domain you want to delete and click **Delete**.
4. The **Delete Storage Domain** dialog displays the list of Protection Groups, Recoveries, Clones, and Views that are dependent on the Storage Domain on separate pages one after another.

   Review the list of dependent objects on each page, select the check box to confirm the deletion, and click **Next**.

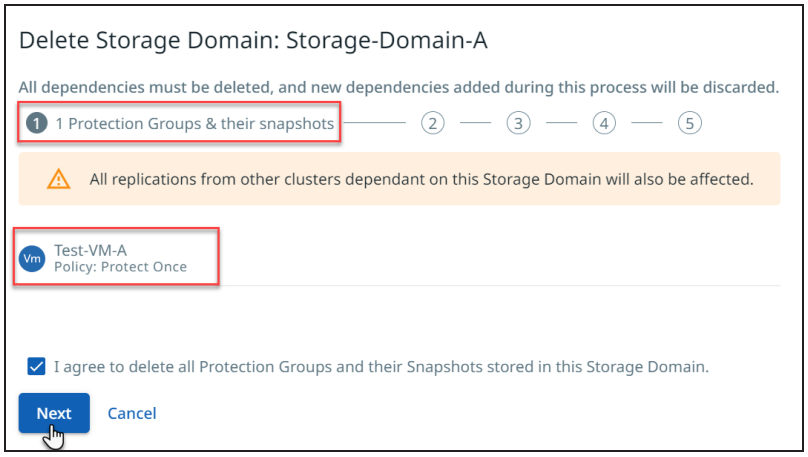

> **Note:** Once you click **Delete Storage Domain** in the next step, you cannot undo the action, associate entities to that Storage Domain, or restore the deleted data.

5. The last page in the **Delete Storage Domain** dialog displays the summary of the dependent objects. To delete the Storage Domain, enter **Yes** and click **Delete Storage Domain**.

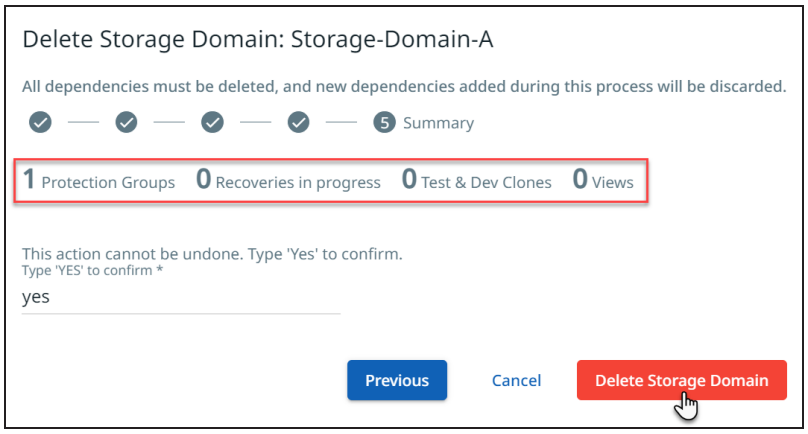

The Storage Domain deletion starts. The duration of the deletion process depends on the size of the Storage Domain.

When the delete succeeds, the Storage Domain and associated objects are deleted and no longer displayed under the Storage Domain tab. Cohesity clusters generate alerts (such as CE02806012) if the deletion fails or is unresponsive.

> **Note:** When you delete a storage domain with CloudTier enabled, the cold data moved to the external target is deleted along with the hot data in the storage domain.

Disk and Node Resilience Settings

Cohesity provides the following protection levels for disk or node resiliency that can be set separately for each storage domain. The allowed values depend on the fault tolerance setting chosen. Cohesity supports node and disk failures, and the number of failures tolerated depends on the configured resiliency. The user selects the minimum failures to tolerate and the system adjusts the available options for RF and EC accordingly.

The following table shows possible EC settings that can be set for the required resiliency of hard drives and nodes. The table indicates the number of failures tolerated in disks *or* nodes at any EC or RF level. For example, if a 5-node Cohesity cluster is set to EC4:2, failures are tolerated in 2 disks or 1 node.

For more details on Cohesity Erasure Coding, see the Cohesity-White-Paper-Fault-Tolerance-Data-Integrity.pdf.

> **Note:** From Cohesity 6.6 onwards, you can use EC 8:2 for Azure, AWS, GCP, Imanis, MS SQL file-based, Physical Server, VMware, NAS, and SAN backup workloads.

Number of Nodes in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
3	1	1	2:1	RF2
	2	1	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4

Number of Nodes in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
4	1	1	2:1, 3:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
5	1	1	2:1, 3:1, 4:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
6	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
7	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Number of Nodes in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
8	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
9	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
10+	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Chassis Resilience Settings

Cohesity provides the following protection levels for chassis resiliency that can be set separately for each storage domain. The allowed values depend on the fault tolerance setting chosen. The user selects the minimum failures to tolerate and the system adjusts the available options for RF and EC accordingly.

The following table shows possible EC settings that can be set for the required resiliency of chassis. The table indicates the number of failures tolerated in disks *or* nodes at any EC or RF level. For example, if 5 chassis is set to EC4:2, failures are tolerated in 2 disks or 1 node.

For more details on Cohesity Erasure Coding, see the Cohesity-White-Paper-Fault-Tolerance-Data-Integrity.pdf

Number of Chassis in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
3	1	1	2:1	RF2
	2	1	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
4	1	1	2:1, 3:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	2	2	2:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
5	1	1	2:1, 3:1, 4:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
6	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4

Number of Chassis in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Nodes (N) or SSD	EC	RF (IF EC is Disabled)
7	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
8	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
9	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
10+	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Rack Resilience Settings

Cohesity provides the following protection levels for rack resiliency that can be set separately for each storage domain. The allowed values depend on the fault tolerance setting chosen. The user selects the minimum failures to tolerate and the system adjusts the available options for RF and EC accordingly.

The following table shows possible EC settings that can be set for the required resiliency of racks.

Number of Racks in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Racks (R)	EC	RF (IF EC is Disabled)
3	1	1	2:1	RF2
	2	1	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
4	1	1	2:1, 3:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	2	2	2:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
5	1	1	2:1, 3:1, 4:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
6	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2	RF3
	3	1	3:3, 4:3, 5:3	RF4

Number of Racks in the Cluster	Number of Failures Tolerated		Storage Domain Resilience Settings	
	Disks (D)	Racks (R)	EC	RF (IF EC is Disabled)
7	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
8	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
9	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4
10+	1	1	2:1, 3:1, 4:1, 5:1	RF2
	2	1	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	2	2	2:2, 3:2, 4:2, 5:2, 6:2, 8:2	RF3
	3	1	3:3, 4:3, 5:3	RF4
	3	2	3:3, 4:3, 5:3	RF4

Manage Nodes

The nodes summary displays the details about all the hardware and cloud nodes that are part of the Cohesity cluster. You can assign nodes to one or more chassis on the nodes summary dashboard and set the node status.

The following topics provide guidance to help you manage hardware and cloud nodes:

Configure a Rack

Using the Cohesity Dashboard or the Cohesity CLI, you can:

- Add a rack and assign it to one or more chassis
- Edit a rack configuration
- View a list of racks, chassis, and nodes

> **Note:** After adding a rack, you cannot delete it. You can edit the rack configuration.

Add a Rack

You can add a rack and associate it with one or more chassis.

> **Note:** Every available chassis must be assigned to a rack to set the rack as a failure domain at the global cluster level. For more details about setting rack as the failure domain, see Configure Cluster Settings.

To add a rack and assign it to one or more chassis:

1. Navigate to **Settings > Summary** and click the **Nodes** tab.
2. On the **Nodes** page, click **Configure Rack**.

   Alternatively, select **Settings > Summary** and click **Configure** to display the Edit Cluster Setting page. Under Custom Failure Domain, click **Configure Rack**.
3. In the Assign Racks page, click **Add** and enter the following information:
   - **Name:** Name of the rack.
   - **Location:** Location of the rack.
   - **Chassis:** Select one or more chassis from the drop-down menu.
4. Click **Assign**.

   The rack is assigned successfully to the selected chassis.

Edit a Rack Configuration

You can edit a rack name and location. After a rack is associated with a chassis, you cannot remove a chassis or associate the chassis with another rack. You cannot delete a rack after you add and assign it to one or more chassis.

To edit or update a rack configuration:

1. Navigate to **Settings > Summary** and click the **Nodes** tab.
2. On the **Nodes** page, click **Configure Rack**.

   Alternatively, select **Settings > Summary** and click **Configure** to display the Edit Cluster Setting page. Under Custom Failure Domain, click **Configure Rack**.
3. In the Assign Racks page, edit or update the following information:
   - **Name:** Name of the rack.
   - **Location:** Location of the rack.
4. Click **Assign**.

   The rack configuration is updated.

View a List of Racks

To view a list of racks and the associated chassis and nodes:

1. Select **Settings > Summary**.
2. Click the **Nodes** tab.

   You can view the chassis and the rack names. You can also view a list of nodes that are part of the chassis.

You can also use Cohesity CLI to add, edit, and view rack details. For more information, see the Cohesity CLI Reference Guide.

Manage Hardware Nodes

A Cohesity node contains CPU with spinning disk or flash storage that is preinstalled with the Cohesity software. When you initially set up a Cohesity cluster that is hosted on hardware natively, the setup process assigns the nodes you selected to the Cohesity cluster. A Cohesity cluster hosted directly on supported hardware must have at least three nodes. A Cohesity cluster Virtual Edition only supports a single node. Additional nodes cannot be added. For more information about supported hardware, see Initial Cluster Setup.

When you connect an additional chassis, nodes will be detected automatically on the Cohesity cluster.

Using the nodes summary dashboard, you can do the following:

- Adding Nodes to the Cohesity cluster
- Removing Nodes from the Cohesity cluster

- Removing HDD Disk from a Hardware Node
- Adding HDD Disk to a Hardware Node
- Removing SSD Disk from a Hardware Node
- Adding SSD Disk to a Hardware Node

Adding Nodes to the Cohesity cluster

To add Nodes to the Cohesity cluster:

> **Note:**
> - If Cluster usage is above 85%, Cohesity recommends adding multiple Nodes to the Cohesity cluster.
> - VLAN VIPs need to be provisioned first before adding the nodes. Typically one per node in each VLAN. VLAN VIPs need not be added for Primary Interface VLAN.
> - You must ensure connectivity to all the existing sources on the newly added nodes.

1. Navigate to **Settings > Summary** and select the **Nodes** tab.
2. Click **Add Node**. The Cohesity cluster looks for Nodes to add to the Cohesity cluster and displays the Node list in the Cohesity Dashboard. You assign Nodes from the list to the Cohesity cluster.

   > **Note:** The cluster can accommodate a maximum of 8 nodes when adding them at the same time.

   To know the software version on the new node, SSH to the node and run the following command:

   iris_cli node status

   A Cohesity cluster Virtual Edition only supports a single node. Additional nodes cannot be added. The **Add Node** option is only available for Cohesity clusters that are hosted directly on supported hardware.

   a. Select the node(s) you want to add to the Cohesity cluster and click **Next**.
   b. Specify IP addresses and IPMI IP addresses for each of the new Nodes and click **Next**. You must open specific ports in the firewall to allow the Cohesity cluster for backup and recovery operations. For more information, see Manage Firewall Ports.

> **Note:** If a Node's software version is different from the Cohesity cluster's software version, the Node's software is automatically updated to the Cohesity cluster's version when you add the Node. Also, the IPMI configuration is optional. You can configure IPMI after adding nodes to the Cohesity cluster using the Cohesity CLI or the Cohesity Dashboard.

Removing Nodes from the Cohesity cluster

A Cohesity cluster node may have multiple Cohesity services such as Data Protection service, I/O operations service, and Indexing service running. Removing the node without gracefully handling these existing services may result in data loss. To prevent this, Cohesity recommends a pre-check of the nodes. During this pre-check process, the resources used by a service running on that node are optimally redistributed to other available nodes, and then the service is stopped. After all the services have stopped, the node is removed, thereby handling the node removal in an orderly manner.

Before you Begin:

- Do not remove drives from the Cohesity cluster unless it is required to fix a verified issue such as a failed drive as described later in this topic.
- Do not upgrade your Cohesity cluster while replacing the drive.
- Do not add new nodes while replacing the drive.

To remove nodes from the Cohesity cluster:

1. Navigate to **Settings > Summary** and click the **Hardware** tab.
2. Select the node you want to remove, click the **Actions** menu (⋮) next to the node, and click **More Information** to view the node information.
3. Click **Run Removal Precheck** to pre-check the nodes before removal.

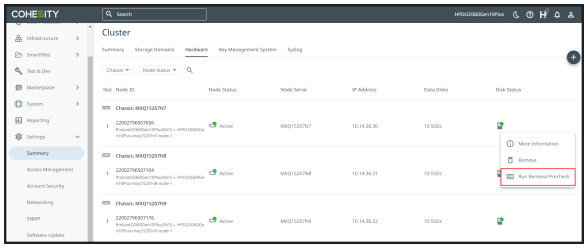

The status of the removal pre-check is displayed, as shown below.

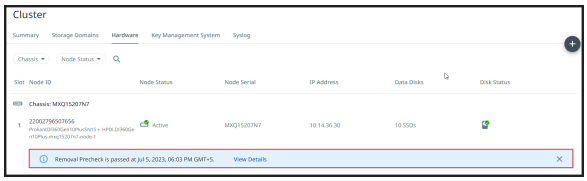

4. Click **View Details** and see the pre-check details.

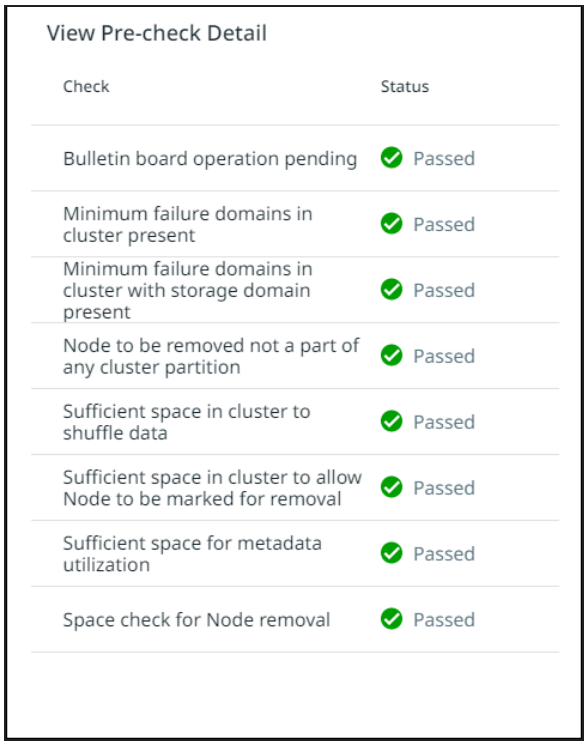

5. Verify the node information and click **Remove**.

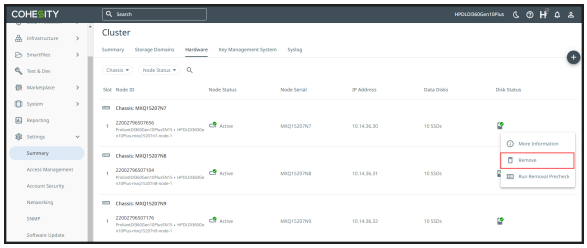

The nodes are pre-checked before the removal to ensure that they are not stopped abruptly.

6. Once the pre-check is passed and the nodes are ready for removal, click **Remove**.

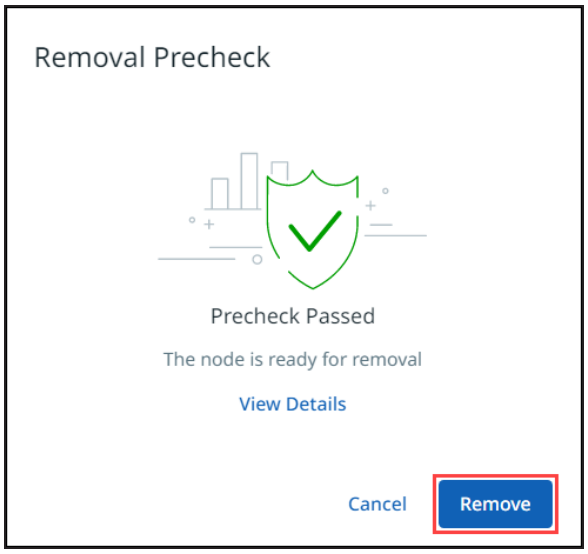

> **Note:** Node will be unusable once it is marked for removal.

After you confirm the node removal, the `NodeMarkedForRemoval` label is observed.

Once the node is removed, the `RemoveNodeSuccess` label is displayed with the details, as shown below.

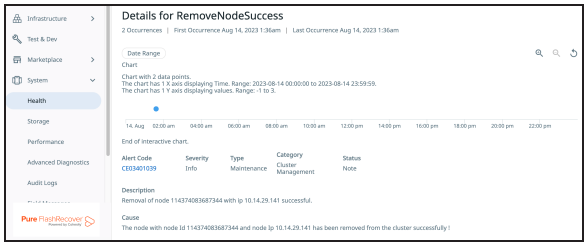

Removing HDD Disk from a Hardware Node

Before you Begin:

- Do not remove drives from the Cohesity cluster unless it is required to fix a verified issue such as a failed drive as described later in this topic.
- Do not upgrade your Cohesity cluster while replacing the drive.
- Do not add new nodes while replacing the drive.

To remove the HDD disk from a hardware node:

1. Navigate to **Settings > Summary > Hardware**.
2. Select a node to see the HDD disks associated with the node.

3. Identify the failed disk by lighting the disk LED. Select **Settings > Summary> Nodes**, and click **LED On** in the Actions menu (⋮) to identify the disk. Once turned on, "**LED turned on successfully**" notification will be observed.
4. Select the HDD disk you want to remove and click the Actions menu (⋮) next to the HDD disk.
5. Click **Run Removal Precheck** to pre-check the nodes before removal.
6. Verify the information and click **Remove**.

   > **Note:** The HDD disk will be unusable once it is marked for removal.

   The disk status is marked as **Removing**.

   `MarkDiskForRemoval` alert is generated and displayed on the Health page in the Cohesity cluster.

   After the disk is removed, the HDD disk will be marked as **Ready for Physical Removal**.

   Once the removal is finished, `DiskOkToRemove` alert is generated and displayed on the Health page in the Cohesity cluster.

   Once the disk is removed physically, `DiskRemoved` alert is generated and displayed on the Health page in the Cohesity cluster.

Adding HDD Disk to a Hardware Node

To add HDD disk to a hardware node:

1. Physically install the new hard drive and insert the new HDD disk to a node on the Cohesity cluster.
2. The `NewDiskFound` alert is generated and displayed on the Health page in the Cohesity cluster.

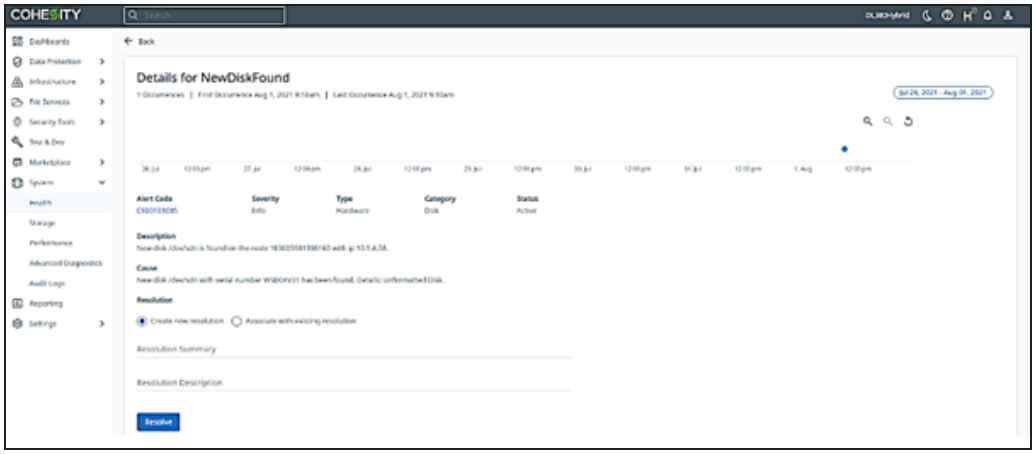

3. Navigate to **Settings > Summary** and click the **Nodes** tab.
4. Select a node to which you have added the new HDD disk. The status of the newly added disk is displayed as **Ready for Activation**.

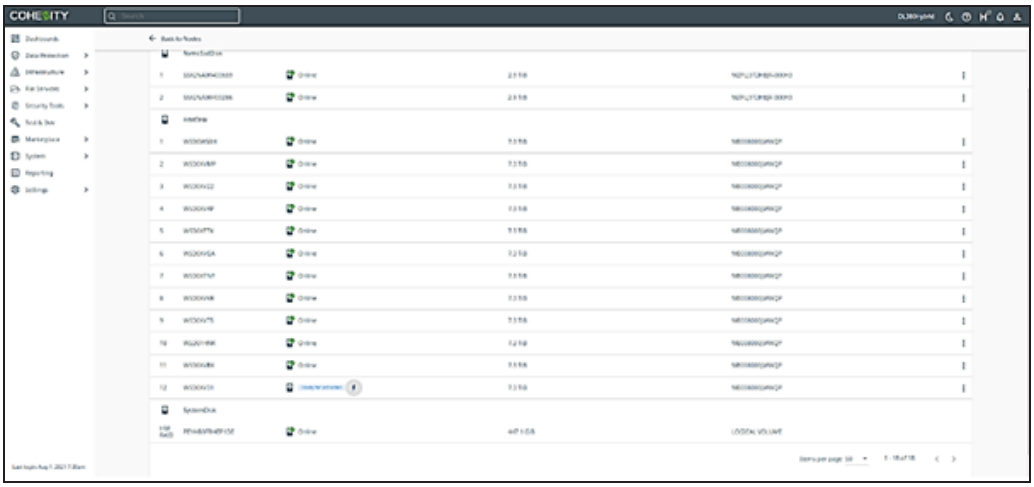

5. Select the newly added HDD disk and click the Actions menu (⋮) next to the HDD disk and click **Activate**.

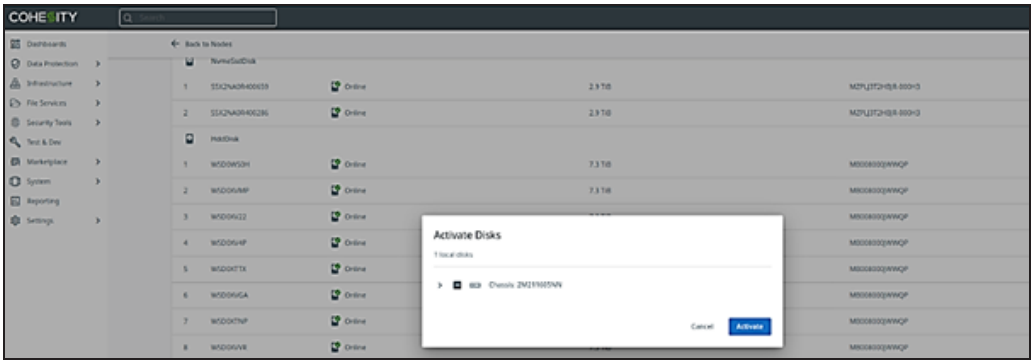

Once the disk activation is completed, the `DiskAddedToCluster` alert is displayed on the Health page in the Cohesity cluster.

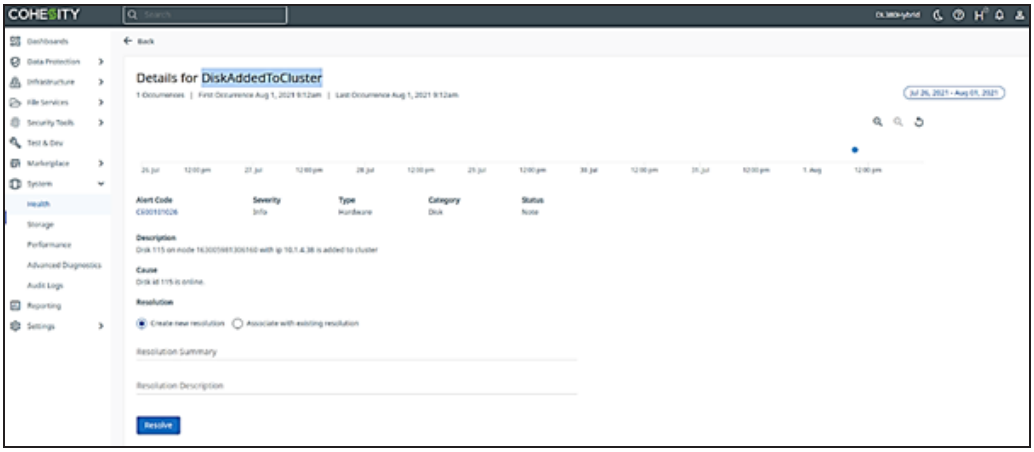

Removing SSD Disk from a Hardware Node

Before you Begin:

- Do not remove drives from the Cohesity cluster unless it is required to fix a verified issue such as a failed drive as described later in this topic.
- Do not upgrade your Cohesity cluster while replacing the drive.
- Do not add new nodes while replacing the drive.

To remove SSD disk from a hardware node:

1. Navigate to **Settings > Summary > Hardware**.
2. Select a node to see the SSD disks associated with the node.
3. Identify the failed disk by lighting the disk LED. Navigate to **Settings > Summary> Hardware**, and click **LED On** in the Actions menu (⋮) to identify the disk. Once turned on, "**LED turned on successfully**" notification will be observed.

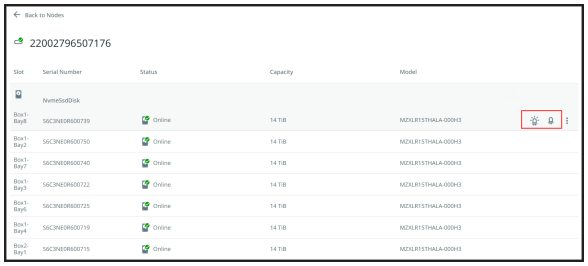

4. Select the SSD disk you want to remove and click the **Actions** menu (⋮) next to the SSD disk.
5. Click **Run Removal Precheck** to pre-check the nodes before removal.

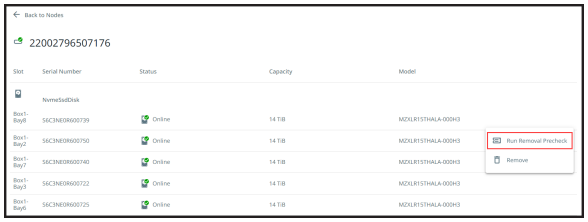

The status of the removal pre-check is displayed, as shown below.

6. Click **View Details** and see the pre-check details.

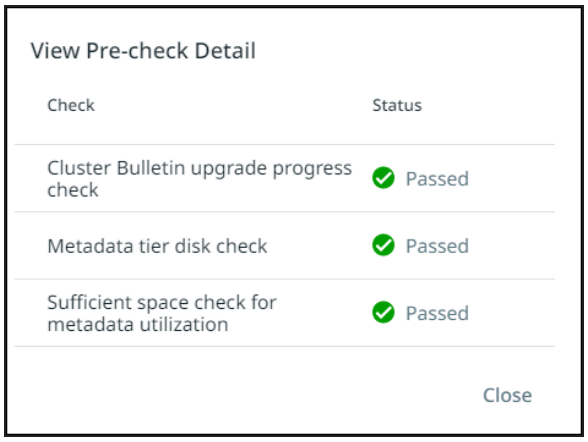

7. Verify the information and click **Remove**.

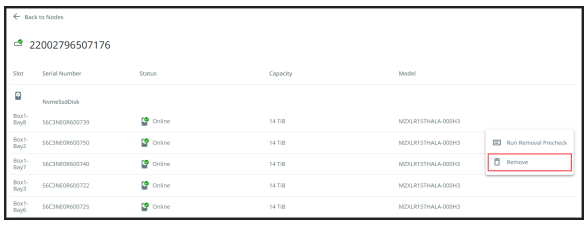

The disks are pre-checked before the removal to ensure that they are not stopped abruptly.

8. Once the pre-check is passed and the disks are ready for removal, click **Remove**.

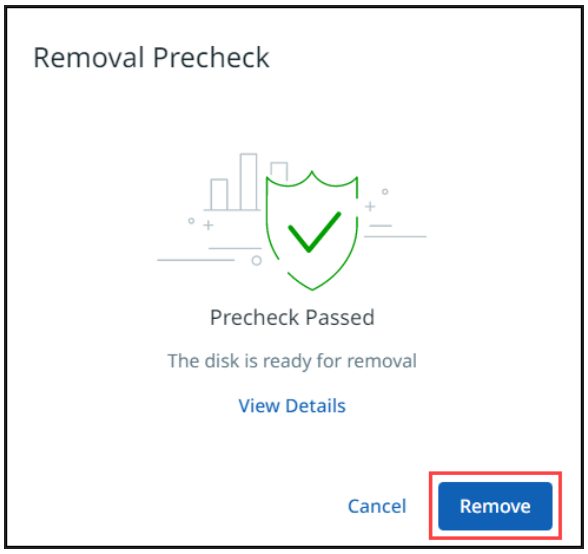

9. Enter '**Yes**' to confirm removal and click '**Start Removing**'.

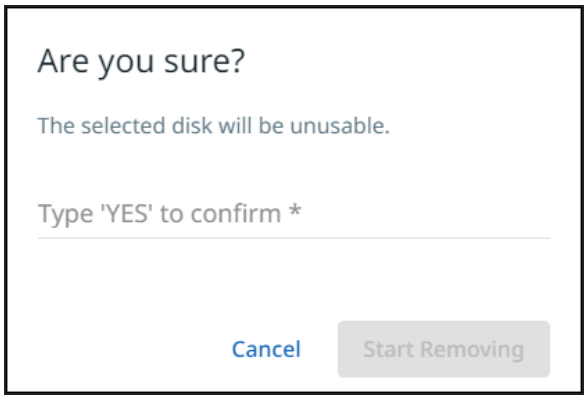

The status changes to 'Removing initiated' with the timestamp.

10. Click the status to view more information on the disk removal. The component information and the status of removal appear, as shown below.

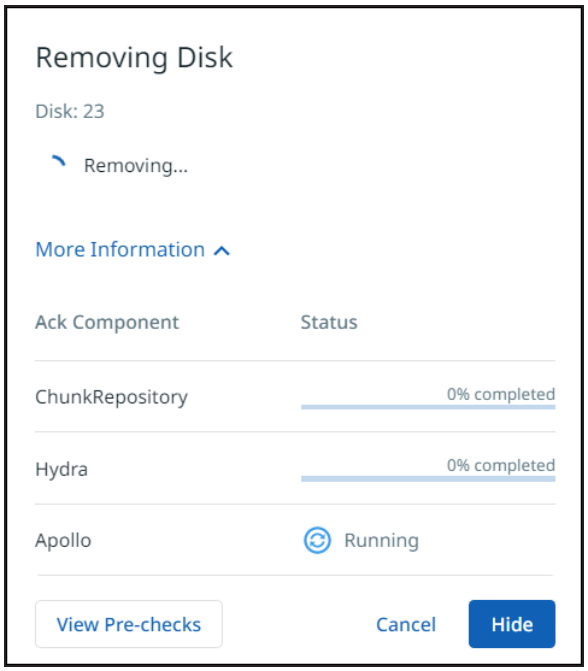

11. Click **Hide** to hide the information.
12. To cancel the disk removal, click **Cancel** and click **Yes**.

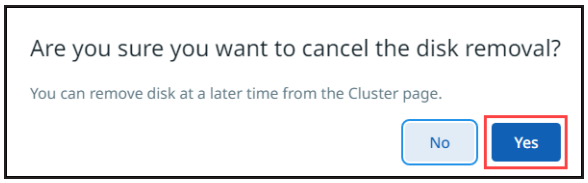

> **Note:** The SSD disk will be unusable once it is marked for removal.

The disk status is marked as **Removing**.

`MarkDiskForRemoval` alert is generated and displayed on the Cohesity cluster UI.

After the disk is removed, the SSD disk will be marked as **Ready for Physical Removal** on the Cohesity cluster.

Once the removal is finished, `DiskOkToRemove` alert is generated and displayed on the Cohesity cluster.

Once the disk is removed physically, `DiskRemoved` alert is generated and displayed on the Health page in the Cohesity cluster.

Adding SSD Disk to a Hardware Node

To add SSD disk to a hardware node:

1. Physically install the new hard drive and insert the new SDD disk to a node on the Cohesity cluster.
2. After the scanning is completed, the `NewDiskFound` alert is displayed on the Health page in the Cohesity cluster.

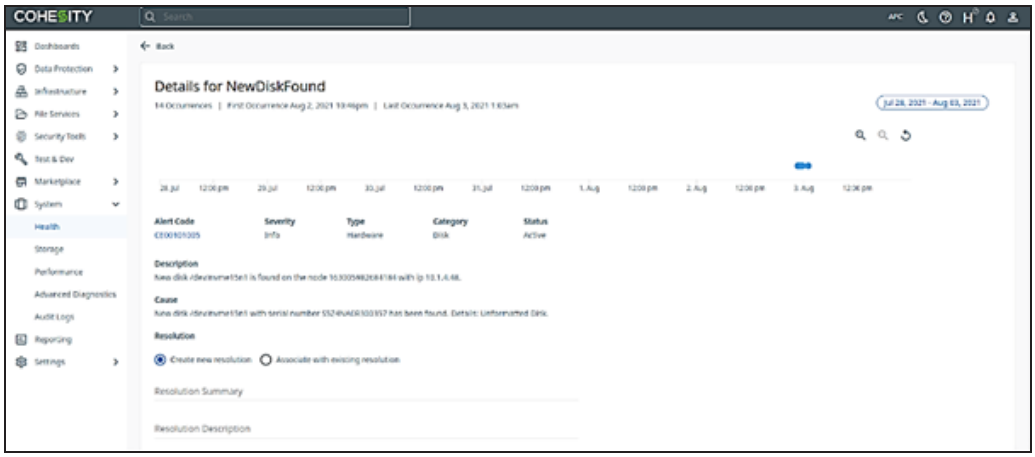

3. Navigate to **Settings > Summary** and click the **Nodes** tab.
4. Select a node to which you have added the new SSD disk. The status of the newly added disk is displayed as **Ready for Activation**.

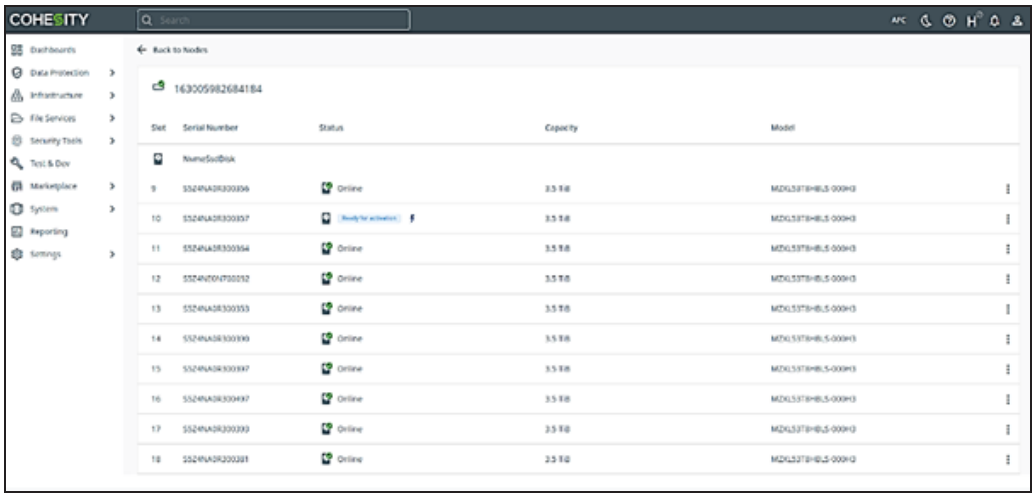

5. Select the newly added SSD disk and click the Actions menu (⋮) next to the SSD disk and click **Activate**.

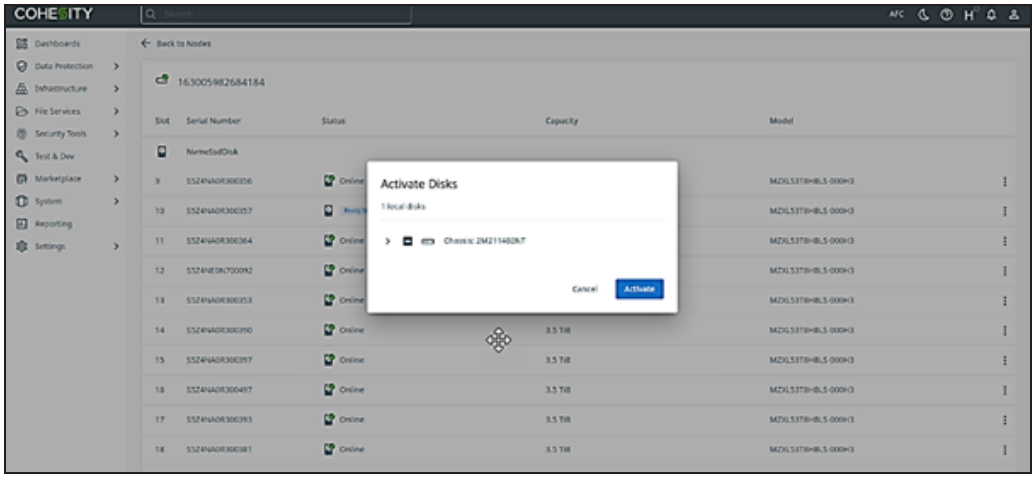

After the new disk is added to the node, the `DiskAddedtoCluster` alert is displayed on the Health page in the Cohesity cluster.

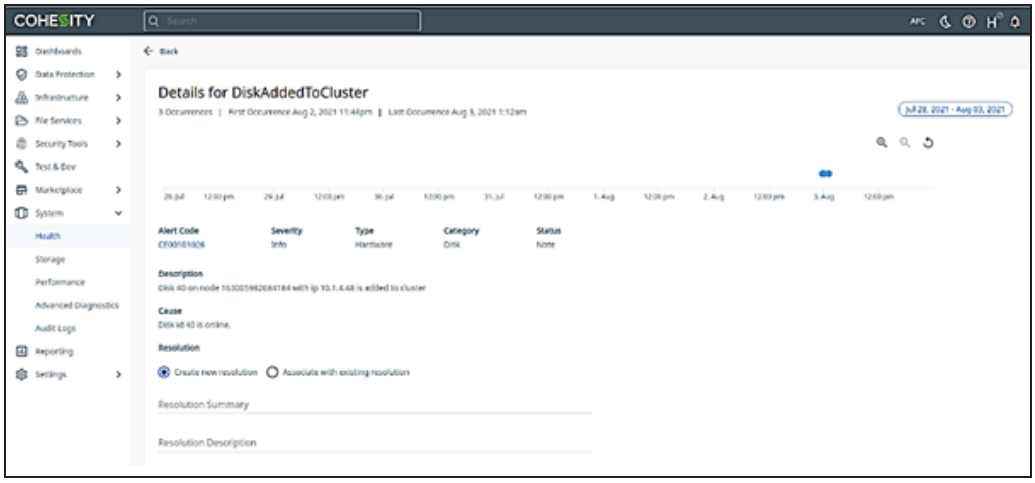

Manage Multi-Node Virtual Edition

A Cohesity Virtual Edition(VE) node is a virtual machine with disk storage that is preinstalled with the Cohesity software. When you initially set up a Cohesity cluster that is hosted on a hypervisor, the setup process assigns the nodes you selected to the Cohesity cluster. A Cohesity Virtual Edition multi-node (8TB and 16TB) cluster must have at least three nodes.

> **Note:** A Cohesity Virtual Edition single-node cluster only supports a single node. Additional nodes cannot be added.

Add Nodes to the Multi-Node Virtual Edition Cluster

To add a new node to the Virtual Edition multi-node cluster, the new node must be deployed with the same product model OVA used for the Virtual Edition cluster.

For example:

- For 8TB Virtual Edition cluster, deploy the new node with 8 TB OVA/vmdk image
- For 16TB Virtual Edition cluster, deploy the new node with 16 TB OVA/vmdk image.

The new node and the cluster nodes must be accessible to each other through a network.

> **Note:** If cluster usage is above 85%, Cohesity recommends adding multiple nodes to the Cohesity cluster.

You can add new nodes either using the Cohesity Dashboard or Cohesity CLI.

Perform the following steps to add nodes using Cohesity Dashboard:

1. Select **Settings > Summary** and select the **Nodes** tab.
2. Click **Add Node**. The Cohesity cluster looks for Nodes to add to the Cohesity cluster and displays the Node list in the Cohesity Dashboard. You assign Nodes from the list to the Cohesity cluster.

   To know the software version on the new node, SSH to the node and run the following command:

   ```
   iris_cli node status
   ```

   A Cohesity Virtual Edition single-node cluster only supports a single node. Additional nodes cannot be added. The **Add Node** option is only available for clusters that are hosted directly on supported hardware and Virtual Edition multi-node (8TB and 16TB) clusters. For more information about supported hardware, see Initial Cluster Setup.
3. Select the node(s) you want to add to the Cohesity cluster and click **Next**.
4. Specify IP addresses for each of the new nodes and click **Next.**

> **Note:** If a node's software version is different from the cluster's software version, the node's software is automatically updated to the cluster's version when you add the node. Also, the IPMI configuration is optional. You can configure IPMI after adding nodes to the Cohesity cluster using the Cohesity CLI or the Cohesity Dashboard.

5. Enter new VIPs if any and click **Finish**.

Perform the following steps to add nodes using Cohesity CLI:

Use the vSphere Web Client 6.5.0 to deploy and configure the Virtual Edition cluster.

> **Note:** The procedure you use may be different depending on the vSphere release and the vSphere Client type.

1. Obtain the Clustered Virtual Edition for VMware OVA file from the Cohesity Download Site. The OVA package file contains the Virtual Edition OVF template and VMDK file. It can be imported into a vCenter Server like an OVF template file. For more information, see Cohesity Support.
2. Using a VMware vSphere Web Client, log in to a vCenter Server that will host the Virtual Edition Virtual Machines:
    1. Enter the IP address or host name of the vCenter Sever into the address bar of a browser.
    2. If prompted, select the **vSphere Web Client (Flash)** option.
3. In the Inventory located in the left panel, navigate to your vCenter server, right-click on the vCenter root, and select **Deploy OVF Template...**.

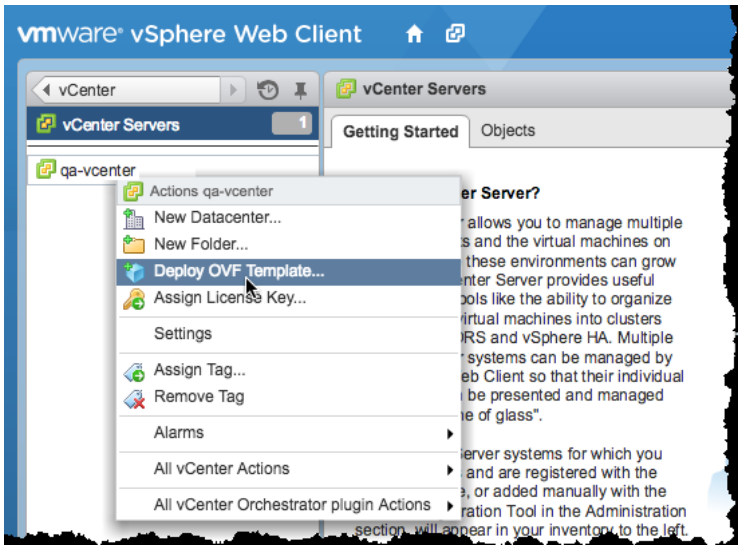

4. Under Local file, click **Browse** and find the CohesityVirtual Edition OVA file and click **Next** and **Open**.

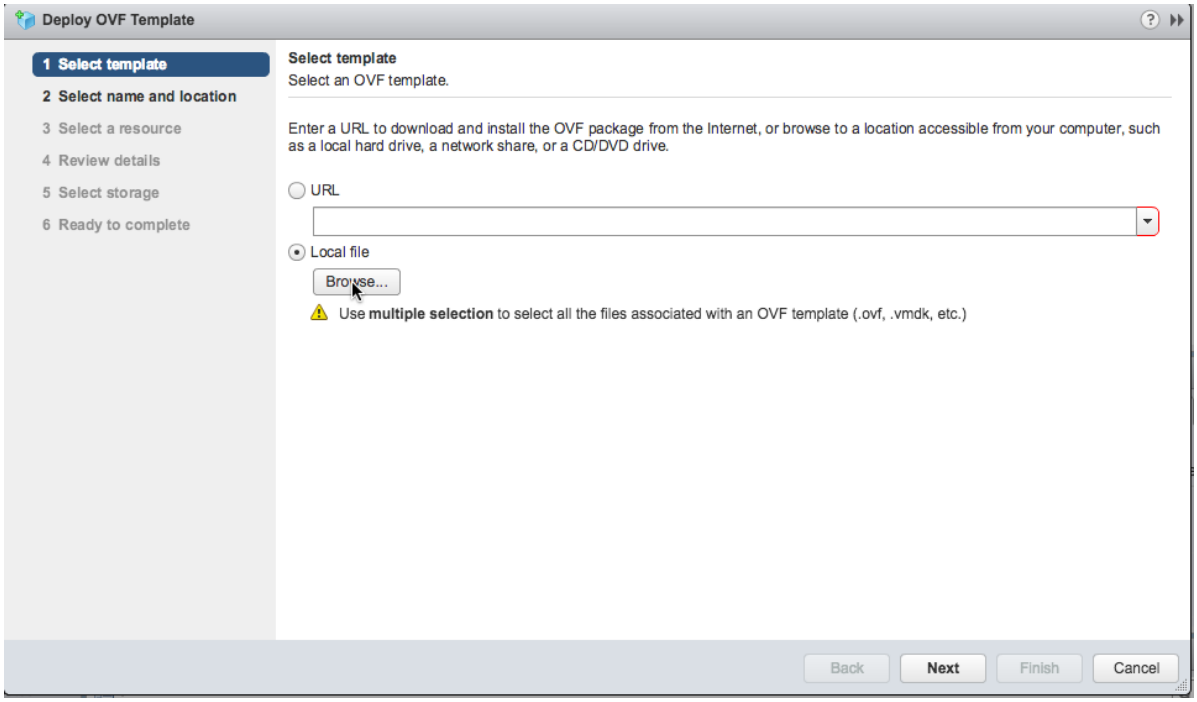

5. In the **Select name and location** panel, enter a unique name for the Virtual Machine such as `Virtual_Edition1` in the **Name** field. Get the unique Virtual Machine names that were saved in the Worksheet. When you create a new Virtual Machine, you must specify a unique Virtual Machine name for each node in the cluster. Each name must be unique to distinguish it from the existing VMs in the parent folder or Datacenter, and must not exceed 80 characters.
6. **Click Next**.
7. In the **Select a resource** panel, select where the VM should run and click **Next**.
8. In the **Review details** panel, click **Next**.
9. From the deployment **Configuration** drop-down, select an appropriate deployment configuration provided by Cohesity:
    - The small (8 TB) configuration supports Virtual Machines with 4 vCPUs, 32 GB of memory and 64 GB of virtual disk to store the operating system as the minimum configuration.
    - The large (16 TB) configuration supports Virtual Machines with 8 vCPUs, 64 GB of memory and 64 GB of virtual disk to store the operating system as the minimum configuration.

      The Metadata drive size must be smaller than the Data Tier drive size. A single Virtual Cohesity cluster must contain all nodes of small configurations or all nodes of large configuration only.

10. Click **Next**.
11. In the **Select storage** panel, select a VMFS datastore to store the new virtual machine files. Set the **Thick Provision Lazy Zeroed** virtual disk format. Click **Next**.
12. In the **Select networks** panel, for the **Destination** drop-down, select networks. Select the type of IP Address allocation, either dynamic **DHCP** or **Static - Manual**.
13. Click **Next**.
14. In the **Customize template** panel, if you are using static (manual) networking, provide values for the following Data Network Properties:
    - Network IP Address
    - Network Netmask
    - Default Gateway—If a Secondary Network interface is configured using the Cohesity CLI during the Virtual Edition installation procedure, the Secondary Network is used as the default gateway for the Cohesity cluster.

    If you are using DHCP networking, leave these three fields blank.

    > **Note:** If the VM is deployed over ESXi without vApp support, then the network information such as IP address, gateway, netmask for both primary and secondary network provided cannot be applied to the VM.

15. In the **Customize template** panel, you can optionally specify a Secondary Network interface. To configure a Secondary Network, you must specify a static manual settings for the following properties:
    - Network IP Address
    - Network Netmask
    - Default Gateway—If a Secondary Network interface is configured using the Cohesity CLI during the Virtual Edition installation procedure, the Secondary Network is used as the default gateway for the Cohesity cluster.

    If you are using DHCP networking, leave these three fields blank.
16. Click **Next**.
17. Click **Next** and then **Finish**.

    The process to deploy the VM starts.
18. Wait until the VM is deployed before continuing to the next procedure. Wait until all the following tasks are completed:
    - Upload disk
    - Reconfig VM
    - Import VAPP
19. Run the following command:

```
cluster add-nodes node-ips=<list of node_ips to be added>
```

Manage Cloud Edition (Legacy) Cluster Nodes on AWS

Cohesity Cloud Edition nodes in AWS is an AWS EC2 Instance with Cohesity software installed on it. Nodes are assigned to a Partition and the Partition uses the Nodes to store data.

When you initially set up a cluster, you specify the number of nodes and cluster size: small, medium or large. Small is 2TB HDD, medium is 6TB HDD and large is 12TB HDD. SSD and instance sizes remain the same for all cluster sizes.

> **Important:** To add a new node to a Cohesity cluster, both the node and Cohesity cluster must have Cohesity 6.0 or later installed. Upgrade the Cohesity cluster to at least 6.0x before adding a new Node.

Adding Nodes to Cohesity Cloud Edition (Legacy) in AWS

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the private (VPN-only) subnet of the AWS Virtual Private Cloud (VPC).

2. Log in to the Control VM.
   a. Open a terminal window on your machine.
   b. Ssh to the Control VM instance:

      ```
      $ ssh cohesity@<control_vm_ip_address>
      ```

      Enter the password for the `cohesity` account provided by Cohesity.
   c. If prompted, answer `yes` to continue connecting.
3. Change to the following directory.

   ```
   cd /home/cohesity/data/aws
   ```

4. Make a copy of the configuration file `cluster_config_params.json` that was used to create the cluster, for example:

   ```
   $ cp cluster_config_params.json copy_cluster_config_params.json
   ```

5. Edit the copied file with a Linux editor, such as `vi`.

   ```
   $ vi copy_cluster_config_params.json
   ```

6. Update the following parameters and save the file:

Property in JSON File	Description
`cohesity_aws_num_vms`	The number of additional nodes you are adding to the existing cluster.
`cohesity_aws_vm_ip_addresses`	The same number of IP addresses as the number of nodes specified. Provide the same number of IP addresses as the number of nodes specified (cohesity_aws_num_vms). All IP addresses must be available and must be in the private (VPN-only) subnet of the AWS Virtual Private Cloud (VPC). Specify these IP addresses in a comma separated list.
`cohesity_aws_s3_bucket_name`	An existing S3 bucket name is required for larger clusters. If you are creating a cluster (or adding nodes to an existing cluster), and the resulting cluster will have: • 20 or fewer nodes, the cluster will be deployed using the inline, default CloudFormation template, which is limited to 64KB and 200 resources. No bucket name is required. Specify two quotes "" with no space between them. • 21 or more nodes, provide the name of an existing S3 bucket. The Cohesity software will prepare a CloudFormation template with sufficient resource limits, upload the template to the bucket, deploy it from the bucket, and then delete the template after successful deployment.

7. Create the new Node VMs. Run the Setup Utility with the `create_vms` option.

```
/home/cohesity/software/bin/cohesity_aws_setup create_vms -cohesity_
aws_params_file=/home/cohesity/data/aws/copy_cluster_config_params.json
```

If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above command.

```
-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>
-http_proxy_username=<serverUsername> -http_proxy_
password=<serverPassword>
```

Wait until the process is completed. During the process, the IP addresses of the VMs are displayed in the output.

For more information about the commands to manage Cohesity cluster, see Commands to Manage Cohesity Cohesity Cloud Edition Cluster.

8. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.
9. Use the following command to add the new VMs to the cluster:

```
admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...>
```

Specify the IP addresses that you specified in the `cohesity_aws_vm_ip_addresses` JSON property.

Do not specify any spaces between the comma-separated IP addresses.

Commands to Manage Cloud Edition Cluster (legacy) on AWS

The following commands are available in both Control VM and Cohesity node.

Description	Command
To validate the `cohesity_aws_params_file`	`validate_cluster -cohesity_aws_params_file=<required>`
To start a cluster	`start_cluster -cohesity_aws_params_file=<required>`
To stop a cluster	`stop_cluster -cohesity_aws_params_file=<required>`
To create a cluster	`create_cluster -cohesity_aws_params_file=<required>`
To restart a cluster	`restart_cluster -cohesity_aws_params_file=<required>`
To delete a cluster	`destroy_cluster -cohesity_aws_params_file=<required>` Alternatively, you can delete a cluster from the Amazon AWS console; select the CloudFormation service, select the stack that has the name specified on the `cohesity_aws_cluster_name` property in the params file, and select **Delete**.
To create new nodes to be added to a cluster	`create_vms -cohesity_aws_params_file=<required>`
To delete VMs that have been removed from the cluster	`delete_vms -cohesity_aws_params_file=<required>`
Note: If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information with the following options on the above commands. `-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort> -http_proxy_username=<serverUsername> -http_proxy_password=<serverPassword>`	

Syntax to run the command in the Control VM:

```
/home/cohesity/software/bin/cohesity_aws_setup <command> -cohesity_aws_params_file=<customer_config_params.json>
```

An example to stop a cluster from the Control VM is as follows:

```
/home/cohesity/software/bin/cohesity_aws_setup stop_cluster -cohesity_aws_
params_file= /home/cohesity/data/aws/awsce.json
```

Syntax to run the command on a Cohesity node:

```
/home/cohesity/bin/tools/cohesity_aws_setup <command> -cohesity_aws_params_
file=<customer_config_params.json>
```

An example to stop a cluster on a Cohesity node is as follows:

```
/home/cohesity/bin/tools/cohesity_aws_setup stop_cluster -cohesity_aws_
params_file=awsce.json
```

Manage Cloud Edition (Legacy) Cluster Nodes on Azure

DataProtect for Cloud (Legacy) nodes in Azure is a Microsoft Azure VM that has the Cohesity software installed on it. A Partition and Storage Domain are created during the initial Cohesity cluster creation.

When you initially set up a cluster, you specify the number of nodes and cluster size: small, medium or large. Small is 2TB HDD, medium is 6TB HDD and large is 12TB HDD. SSD and instance sizes remain the same for all cluster sizes.

> **Important:** To add a new node to a Cohesity cluster, both the node and Cohesity cluster must have Cohesity 6.0 or later installed. Upgrade the Cohesity cluster to at least 6.0x before adding a new Node.

Adding Nodes to Cloud Edition (Legacy) Cluster Nodes on Azure

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the VPN subnet of the VPN network.

2. Log in to Control VM that was used during initial cluster setup. This VM must have Internet access so the Cohesity utility can call the Cloud Service APIs to create new Nodes.

   a. Open a terminal window on your machine.

   b. Ssh in to the Control VM:

      ```
      $ ssh cohesity@<control_vm_ip_address>
      ```

      Enter the password for the `cohesity` account provided by Cohesity.

   c. If prompted, answer `yes` to continue connecting.

3. Change to the following directory.

   ```
   cd /home/cohesity/software
   ```

4. Make a copy of the configuration file `params.json` that was used to create the cluster, for example:

```
$ cp params.json copy_params.json
```

5. Edit the copied file with a Linux editor, such as `vi`.

```
$ vi copy_params.json
```

6. Update the following parameters and save the file.

JSON Property	Description
`cohesity_azure_num_vms`	The number of additional nodes you are adding to the existing cluster.
`cohesity_azure_vm_ip_addresses`	The IP addresses to assign to the nodes (VMs) of the Cohesity cluster. Provide the same number of IP addresses as the number of nodes specified. All IP addresses must be available and must be in the VPN subnet of the VPN network. Separate the IP addresses with commas.

7. Run the Setup Utility with the `validate` option to validate the Azure credentials. This is a dry run of the command and the Node VMs are not created.

```
/home/cohesity/software/cohesity_azure_setup validate -cohesity_azure_
setup_params_file=/home/cohesity/software/copy_params.json
```

8. If no errors are reported and utility reports the Azure credentials as successfully validated, run the Setup Utility to create the VMs.

```
/home/cohesity/software/cohesity_azure_setup create_vms -cohesity_
azure_setup_params_file=/home/cohesity/software/copy_params.json
```

   If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above command. (Password protected proxy servers are not supported.)

```
-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>
```

   Wait until the process completes. During the process, the IP addresses of the VMs are displayed in the output. In addition, the name of the Cohesity cluster and the resource group that contains the Cohesity cluster are displayed in the output.

9. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.

10. Using the following command to add the new VMs to the cluster:

```
admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...>
```

Specify the IP addresses that you specified in the `cohesity_azure_vm_ip_addresses` JSON property.

Do not specify any spaces between the comma-separated IP addresses.

Commands to Manage Cloud Edition (Legacy) Cluster Nodes on Azure

The following setup utility commands are available on both the Control VM and nodes in the cluster.

Description	Command
To validate the `cohesity_azure_params_file`	`validate -cohesity_azure_setup_params_file=<required>`
To create a cluster	`create_cluster -cohesity_azure_setup_params_file=<required>`
To delete a cluster	`destroy_cluster -cohesity_azure_setup_params_flle=<required> -cohesity_azure_resource_group=<required>`
To start a cluster	`startup_cluster -cohesity_azure_setup_params_file=<required> -cohesity_azure_resource_group=<required>`
To stop a cluster	`shutdown_cluster -cohesity_azure_setup_params_file=<required> -cohesity_azure_resource_group=<required>`
To delete node (VMs) that have been removed from the cluster	`delete_vms -cohesity_azure_setup_params_file=<required> -cohesity_azure_resource_group=<required> -cohesity_azure_vm_list_csv=<required>`
To create new nodes (VMs) to be added to a cluster	`create_vms -cohesity_azure_setup_params file=<required>`
Note: If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above commands. (Password protected proxy servers are not supported.) `-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>`	

Manage Cloud Edition (Legacy) Cluster Nodes on GCP

Cloud Edition (Legacy) cluster nodes on GCP is a Google Compute Engine instance that has the Cohesity software installed on it.

When you initially set up a cluster, you specify the capacity of the Cohesity cluster as small, medium or large. Small is 2TB HDD, medium is 6TB HDD and large is 12TB HDD. SSD and instance sizes remain the same for all cluster sizes.

> **Important:** To add a new node to a Cohesity cluster, both the node and Cohesity cluster must have Cohesity 6.0 or later installed. Upgrade the Cohesity cluster to at least 6.0x before adding a new node.

Adding Nodes to Cloud Edition (Legacy) Cluster Nodes on GCP

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the VPC subnet of your VPC network.

2. Log in to Control VM. The VM must have internet access so the Cohesity utility can call the Cloud Service APIs to create new nodes.
   a. Open a terminal window on your machine.
   b. Ssh in to the Control VM:

      ```
      $ ssh cohesity@<control_vm_ip_address>
      ```

      Enter the password for the `cohesity` account provided by Cohesity.
   c. If prompted, answer `yes` to continue connecting.

3. Change to the following directory.

   ```
   cd /home/cohesity/data/gcp
   ```

4. Make a copy of the configuration file `cluster_config_params.json` that was used to create the cluster, for example:

   ```
   $ cp cluster_config_params.json copy_cluster_config_params.json
   ```

5. Edit the copied file with a Linux editor, such as `vi`.

   ```
   $ vi copy_cluster_config_params.json
   ```

6. Update the following parameters and save the file:

Property in JSON File	Description
`gcp_num_vms`	The number of additional nodes you are adding to the existing cluster.

Property in JSON File	Description
`gcp_vm_ip_addresses`	Specify the IP addresses to assign to the new nodes (VMs) to add to the Cohesity cluster. All IP addresses must be available and must be in the VPC subnet of your VPC network. Separate the IP addresses by commas.

7. To create the new node VMs to add to the Cohesity cluster, run the Setup Utility with the `create_vms` option.

   ```
   /home/cohesity/software/bin/cohesity_gcp_setup create_vms -cohesity_gcp_params_file /home/cohesity/data/gcp/copy_cluster_config_params.json
   ```

   Wait until the process completes. During the process, the IP addresses of the VMs are displayed in the output.
8. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.
9. Using the following command to add the new VMs to the Cohesity cluster:

   ```
   admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...>
   ```

   Do not specify any spaces between the comma-separated IP addresses.

   Specify the IP addresses that you provided in the `cohesity_gcp_vm_ip_addresses` JSON property.

Commands to Manage Cloud Edition (Legacy) Cluster Nodes on GCP

The following setup utility commands are available on both the Control VM and nodes in the Cohesity cluster.

Description	Command
To create a cluster	`create_cluster -cohesity_gcp_params_file=<required>`
To start a cluster	`start_cluster -cohesity_gcp_params_file=<required>`
To stop a cluster	`stop_cluster -cohesity_gcp_params_file=<required>`
To restart a cluster	`restart_cluster -cohesity_gcp_params_file=<required>`

Description	Command
To delete a cluster	`destroy_cluster -cohesity_gcp_params_flle=<required>` **Note:** Your service account requires the `compute.addresses.deleteInternal` permission to delete and free the static IP addresses associated with the node VMs.
To create new nodes (VMs) to be added to a cluster	`create_vms -cohesity_gcp_setup params file=<required>`
To delete node (VMs) that have been removed from the cluster	`delete_vms -cohesity_gcp_params_file=<required> -cohesity_gcp_vm_list_csv=<required>`

Manage Cohesity Cloud Edition Nodes in AWS

When you initially set up Cohesity Cloud Edition on AWS, you specify the number of nodes. Once the cluster is deployed, you can add or remove nodes as per your requirement.

Add Data Nodes to Cloud Edition Cluster on AWS

Perform the following steps to add data nodes to Cohesity Cloud Edition on AWS:

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the private (VPN-only) subnet of the AWS Virtual Private Cloud (VPC).

2. Log in to the Control VM.
   a. Open a terminal window on your machine.
   b. Ssh to the Control VM instance:

      ```
      $ ssh cohesity@<control_vm_ip_address>
      ```

      Enter the password for the `cohesity` account provided by Cohesity.
   c. If prompted, answer `yes` to continue connecting.
3. Change to the following directory.

   ```
   cd /home/cohesity/data/aws
   ```

4. Make a copy of the configuration file `dpcloud_cluster_config_params.json` that was used to create the cluster, for example:

```
$ cp dpcloud_cluster_config_params.json copy_dpcloud_cluster_config_
params.json
```

5. Edit the copied file with a Linux editor, such as `vi`.

```
$ vi copy_dpcloud_cluster_config_params.json
```

6. Update the following parameters and save the file:

Property in JSON File	Description
`cohesity_aws_num_vms`	The number of additional nodes you are adding to the existing cluster.
`cohesity_aws_vm_ip_addresses`	The same number of IP addresses as the number of nodes specified. Provide the same number of IP addresses as the number of nodes specified (cohesity_aws_num_vms). All IP addresses must be available and must be in the private (VPN-only) subnet of the AWS Virtual Private Cloud (VPC). Specify these IP addresses in a comma separated list.
`cohesity_aws_s3_bucket_name`	An existing S3 bucket name is required for larger clusters. If you are creating a cluster (or adding nodes to an existing cluster), and the resulting cluster will have: • 20 or fewer nodes, the cluster will be deployed using the inline, default CloudFormation template, which is limited to 64KB and 200 resources. No bucket name is required. Specify two quotes "" with no space between them. • 21 or more nodes, provide the name of an existing S3 bucket. The Cohesity software will prepare a CloudFormation template with sufficient resource limits, upload the template to the bucket, deploy it from the bucket, and then delete the template after successful deployment.

7. Create the new Node VMs. Run the Setup Utility with the `create_vms` option.

```
/home/cohesity/software/bin/cohesity_aws_setup create_vms -cohesity_
aws_params_file=/home/cohesity/data/aws/copy_dpcloud_cluster_config_
params.json
```

If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above command.

```
-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>
-http_proxy_username=<serverUsername> -http_proxy_
password=<serverPassword>
```

Wait until the process completes. During the process, the IP addresses of the VMs are displayed in the output.

For more information about the commands to manage Cohesity cluster, see Commands to Manage Cluster.

8. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.
9. Use the following command to add the new nodes to the cluster as data nodes:

```
admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...>
```

Specify the IP addresses that you specified in the `cohesity_aws_vm_ip_addresses` JSON property.

Do not specify any spaces between the comma-separated IP addresses.

Add Compute Nodes to Cloud Edition Cluster on AWS

Perform the following steps to add compute nodes to Cohesity Cloud Edition on AWS:

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the private (VPN-only) subnet of the AWS Virtual Private Cloud (VPC).
2. Log in to the Control VM.
   a. Open a terminal window on your machine.
   b. Ssh to the Control VM instance:

   ```
   $ ssh cohesity@<control_vm_ip_address>
   ```

   Enter the password for the `cohesity` account provided by Cohesity.
   c. If prompted, answer `yes` to continue connecting.
3. Change to the following directory.

```
cd /home/cohesity/data/aws
```

4. Make a copy of the configuration file `dpcloud_cluster_config_params.json` that was used to create the cluster, for example:

```
$ cp dpcloud_cluster_config_params.json copy_dpcloud_cluster_config_
params.json
```

5. Edit the copied file with a Linux editor, such as `vi`.

```
$ vi copy_dpcloud_cluster_config_params.json
```

6. Update the following parameters and save the file:

Property in JSON File	Description
`cohesity_aws_num_vms`	The number of additional nodes you are adding to the existing cluster.
`cohesity_aws_vm_ip_addresses`	The same number of IP addresses as the number of nodes specified. Provide the same number of IP addresses as the number of nodes specified (cohesity_aws_num_vms). All IP addresses must be available and must be in the private (VPN-only) subnet of the AWS Virtual Private Cloud (VPC). Specify these IP addresses in a comma separated list.
`cohesity_aws_s3_bucket_name`	An existing S3 bucket name is required for larger clusters. If you are creating a cluster (or adding nodes to an existing cluster), and the resulting cluster will have: • 20 or fewer nodes, the cluster will be deployed using the inline, default CloudFormation template, which is limited to 64KB and 200 resources. No bucket name is required. Specify two quotes "" with no space between them. • 21 or more nodes, provide the name of an existing S3 bucket. The Cohesity software will prepare a CloudFormation template with sufficient resource limits, upload the template to the bucket, deploy it from the bucket, and then delete the template after successful deployment.

7. Create the new Node VMs. Run the Setup Utility with the `create_vms` option.

```
/home/cohesity/software/bin/cohesity_aws_setup create_vms -cohesity_
aws_params_file=/home/cohesity/data/aws/copy_dpcloud_cluster_config_
params.json
```

If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above command.

```
-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>
-http_proxy_username=<serverUsername> -http_proxy_
password=<serverPassword>
```

Wait until the process completes. During the process, the IP addresses of the VMs are displayed in the output.

For more information about the commands to manage Cohesity cluster, see Commands to Manage Cohesity Cluster.

8. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.
9. Use the following command to add the new nodes to the cluster as compute nodes:

   ```
   admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...> use-as-compute-nodes=true
   ```

   Specify the IP addresses that you specified in the `cohesity_aws_vm_ip_addresses` JSON property.

   Do not specify any spaces between the comma-separated IP addresses.

Commands to Manage Cloud Edition Cluster on AWS

The following commands are available in both Control VM and Cohesity node.

Description	Command
To validate the `cohesity_aws_params_file`	`validate_cluster -cohesity_aws_params_file=<required>`
To start a cluster	`start_cluster -cohesity_aws_params_file=<required>`
To stop a cluster	`stop_cluster -cohesity_aws_params_file=<required>`
To create a cluster	`create_cluster -cohesity_aws_params_file=<required>`
To restart a cluster	`restart_cluster -cohesity_aws_params_file=<required>`
To delete a cluster	`destroy_cluster -cohesity_aws_params_file=<required>` Alternatively, you can delete a cluster from the Amazon AWS console; select the CloudFormation service, select the stack that has the name specified on the `cohesity_aws_cluster_name` property in the params file, and select **Delete**.
To create new nodes to be added to a cluster	`create_vms -cohesity_aws_params_file=<required>`
To delete VMs that have been removed from the cluster	`delete_vms -cohesity_aws_params_file=<required>`
Note: If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information with the following options on the above commands. `-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort> -http_proxy_username=<serverUsername> -http_proxy_password=<serverPassword>`	

Syntax to run the command in the Control VM:

```
/home/cohesity/software/bin/cohesity_aws_setup <command> -cohesity_aws_
params_file=<customer_config_params.json>
```

An example to stop a cluster from the Control VM is as follows:

```
/home/cohesity/software/bin/cohesity_aws_setup stop_cluster -cohesity_aws_
params_file= /home/cohesity/data/aws/awsce.json
```

Syntax to run the command on a Cohesity node:

```
/home/cohesity/bin/tools/cohesity_aws_setup <command> -cohesity_aws_params_
file=<customer_config_params.json>
```

An example to stop a cluster on a Cohesity node is as follows:

```
/home/cohesity/bin/tools/cohesity_aws_setup stop_cluster -cohesity_aws_
params_file=awsce.json
```

Manage Cloud Edition Cluster Nodes on Azure

When you initially set up Cohesity Cloud Edition on Azure, you specify the number of nodes. Once the cluster is deployed, you can add or remove nodes as per your requirement.

Add Data Nodes to Cloud Edition Cluster on Azure

Perform the following steps to add data nodes to Cohesity Cloud Edition on Azure:

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the VPN subnet of the VPN network.
2. Log in to Control VM that was used during initial cluster setup. This VM must have internet access so the Cohesity utility can call the Cloud Service APIs to create new Nodes.
   a. Open a terminal window on your machine.
   b. Ssh in to the Control VM:

      ```
      $ ssh cohesity@<control_vm_ip_address>
      ```

      Enter the password for the `cohesity` account provided by Cohesity.
   c. If prompted, answer `yes` to continue connecting.
3. Change to the following directory.

   ```
   cd /home/cohesity/software
   ```
4. Make a copy of the configuration file `params_dpcloud.json` that was used to create the cluster, for example:

   ```
   $ cp params_dpcloud.json copy_params_dpcloud.json
   ```
5. Edit the copied file with a Linux editor, such as `vi`.

```
$ vi copy_params_dpcloud.json
```

6. Update the following parameters and save the file.

JSON Property	Description
`cohesity_azure_num_vms`	The number of additional nodes you are adding to the existing cluster.
`cohesity_azure_vm_ip_addresses`	The IP addresses to assign to the nodes (VMs) of the Cohesity cluster. Provide the same number of IP addresses as the number of nodes specified. All IP addresses must be available and must be in the VPN subnet of the VPN network. Separate the IP addresses with commas.

7. Run the Setup Utility with the `validate` option to validate the Azure credentials. This is a dry run of the command and the Node VMs are not created.

```
/home/cohesity/software/cohesity_azure_setup validate -cohesity_azure_
setup_params_file=/home/cohesity/software/copy_params_dpcloud.json
```

8. If no errors are reported and utility reports the Azure credentials as successfully validated, run the Setup Utility to create the VMs.

```
/home/cohesity/software/cohesity_azure_setup create_vms -cohesity_
azure_setup_params_file=/home/cohesity/software/copy_params_
dpcloud.json
```

   If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above command. (Password protected proxy servers are not supported.)

```
-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>
```

   Wait until the process completes. During the process, the IP addresses of the VMs are displayed in the output. In addition, the name of the Cohesity cluster and the resource group that contains the Cohesity cluster are displayed in the output.

9. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.
10. Use the following command to add the new nodes to the cluster as data nodes:

```
admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...>
```

    Specify the IP addresses that you specified in the `cohesity_azure_vm_ip_addresses` JSON property.

    Do not specify any spaces between the comma-separated IP addresses.

Add Compute Nodes to Cloud Edition Cluster on Azure

Perform the following steps to add compute nodes to Cohesity Cloud Edition on Azure:

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the VPN subnet of the VPN network.

2. Log in to Control VM that was used during initial cluster setup. This VM must have internet access so the Cohesity utility can call the Cloud Service APIs to create new Nodes.

   a. Open a terminal window on your machine.

   b. Ssh in to the Control VM:

   ```
   $ ssh cohesity@<control_vm_ip_address>
   ```

   Enter the password for the `cohesity` account provided by Cohesity.

   c. If prompted, answer `yes` to continue connecting.

3. Change to the following directory.

   ```
   cd /home/cohesity/software
   ```

4. Make a copy of the configuration file `params_dpcloud.json` that was used to create the cluster, for example:

   ```
   $ cp params_dpcloud.json copy_params_dpcloud.json
   ```

5. Edit the copied file with a Linux editor, such as `vi`.

   ```
   $ vi copy_params_dpcloud.json
   ```

6. Update the following parameters and save the file.

JSON Property	Description
`cohesity_azure_num_vms`	The number of additional nodes you are adding to the existing cluster.
`cohesity_azure_vm_ip_addresses`	The IP addresses to assign to the nodes (VMs) of the Cohesity cluster. Provide the same number of IP addresses as the number of nodes specified. All IP addresses must be available and must be in the VPN subnet of the VPN network. Separate the IP addresses with commas.

7. Run the Setup Utility with the `validate` option to validate the Azure credentials. This is a dry run of the command and the Node VMs are not created.

   ```
   /home/cohesity/software/cohesity_azure_setup validate -cohesity_azure_
   setup_params_file=/home/cohesity/software/copy_params_dpcloud.json
   ```

8. If no errors are reported and utility reports the Azure credentials as successfully validated, run the Setup Utility to create the VMs.

   ```
   /home/cohesity/software/cohesity_azure_setup create_vms -cohesity_
   azure_setup_params_file=/home/cohesity/software/copy_params_
   dpcloud.json
   ```

   If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above command. (Password protected proxy servers are not supported.)

   ```
   -http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>
   ```

   Wait until the process completes. During the process, the IP addresses of the VMs are displayed in the output. In addition, the name of the Cohesity cluster and the resource group that contains the Cohesity cluster are displayed in the output.

9. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.

10. Use the following command to add the new nodes to the cluster as compute nodes:

    ```
    admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...> use-as-
    compute-nodes=true
    ```

    Specify the IP addresses that you specified in the `cohesity_azure_vm_ip_addresses` JSON property.

    Do not specify any spaces between the comma-separated IP addresses.

Commands to Manage Cloud Edition Cluster on Azure

The following setup utility commands are available on both the Control VM and nodes in the cluster.

Description	Command
To validate the `cohesity_azure_params_file`	`validate -cohesity_azure_setup_params_file=<required>`
To create a cluster	`create_cluster -cohesity_azure_setup_params_file=<required>`

Description	Command
To delete a cluster	`destroy_cluster -cohesity_azure_setup_params_ flle=<required> -cohesity_azure_resource_ group=<required>`
To start a cluster	`startup_cluster -cohesity_azure_setup_params_ file=<required> -cohesity_azure_resource_ group=<required>`
To stop a cluster	`shutdown_cluster -cohesity_azure_setup_params_ file=<required> -cohesity_azure_resource_ group=<required>`
To delete node (VMs) that have been removed from the cluster	`delete_vms -cohesity_azure_setup_params_file=<required> -cohesity_azure_resource_group=<required> -cohesity_ azure_vm_list_csv=<required>`
To create new nodes (VMs) to be added to a cluster	`create_vms -cohesity_azure_setup_params file=<required>`
Note: If the Control VM uses an HTTP proxy server to connect to the internet, provide the server information by adding the following options to the above commands. (Password protected proxy servers are not supported.) `-http_proxy_ip=<proxyServerIP> -http_proxy_port=<proxyServerPort>`	

Manage Cloud Edition Cluster Nodes on GCP

When you initially set up Cohesity Cloud Edition on GCP, you specify the number of nodes. Once the cluster is deployed, you can add or remove nodes as per your requirement.

Adding Nodes to Cloud Edition Cluster on GCP

1. Determine unique IP addresses for *each* new Node that will be the added to the existing cluster.

   All IP addresses must be available and must be in the VPC subnet of your VPC network.

2. Log in to Control VM. The VM must have internet access so the Cohesity utility can call the Cloud Service APIs to create new nodes.

   a. Open a terminal window on your machine.

   b. Ssh in to the Control VM:

   ```
   $ ssh cohesity@<control_vm_ip_address>
   ```

Enter the password for the `cohesity` account provided by Cohesity.

c. If prompted, answer `yes` to continue connecting.

3. Change to the following directory.

```
cd /home/cohesity/data/gcp
```

4. Make a copy of the configuration file `dpcloud_cluster_config_params.json` that was used to create the cluster, for example:

```
$ cp dpcloud_cluster_config_params.json copy_dpcloud_cluster_config_
params.json
```

5. Edit the copied file with a Linux editor, such as `vi`.

```
$ vi copy_dpcloud_cluster_config_params.json
```

6. Update the following parameters and save the file:

Property in JSON File	Description
`gcp_num_vms`	The number of additional nodes you are adding to the existing cluster.
`gcp_vm_ip_addresses`	Specify the IP addresses to assign to the new nodes (VMs) to add to the Cohesity cluster. All IP addresses must be available and must be in the VPC subnet of your VPC network. Separate the IP addresses by commas.

7. To create the new node VMs to add to the Cohesity cluster, run the Setup Utility with the `create_vms` option.

```
/home/cohesity/software/bin/cohesity_gcp_setup create_vms -cohesity_
gcp_params_file /home/cohesity/data/gcp/copy_dpcloud_cluster_config_
params.json
```

Wait until the process completes. During the process, the IP addresses of the VMs are displayed in the output.

8. Start the Cohesity CLI remotely or locally as described in Using the Cohesity CLI.

9. Using the following command to add the new VMs to the Cohesity cluster:

```
admin> iris_cli cluster cloud-join node-ips=<ip1,ip2,...>
```

Do not specify any spaces between the comma-separated IP addresses.

Specify the IP addresses that you provided in the `cohesity_gcp_vm_ip_addresses` JSON property.

Commands to Manage Cloud Edition Cluster on GCP

The following setup utility commands are available on both the Control VM and nodes in the Cohesity cluster.

Description	Command
To create a cluster	`create_cluster -cohesity_gcp_params_file=<required>`
To start a cluster	`start_cluster -cohesity_gcp_params_file=<required>`
To stop a cluster	`stop_cluster -cohesity_gcp_params_file=<required>`
To restart a cluster	`restart_cluster -cohesity_gcp_params_file=<required>`
To delete a cluster	`destroy_cluster -cohesity_gcp_params_flle=<required>` **Note:** Your service account requires the `compute.addresses.deleteInternal` permission to delete and free the static IP addresses associated with the node VMs.
To create new nodes (VMs) to be added to a cluster	`create_vms -cohesity_gcp_setup params file=<required>`
To delete node (VMs) that have been removed from the cluster	`delete_vms -cohesity_gcp_params_file=<required> -cohesity_gcp_vm_list_csv=<required>`

Manage Compute Nodes

Compute Nodes are special nodes in the Cohesity cluster that provide dedicated CPU and memory to run resource-intensive apps and services. These nodes do not store any data and do not require disks other than the boot disk. Nodes are automatically detected as compute nodes or can be designated as compute nodes at the time of adding nodes to a Cohesity cluster. Compute nodes are only required if you are running resource-intensive applications.

You can run multiple applications or multiple instances of a single application on a compute node.

System Requirements

The following are the system requirements for a compute node:

- System Disk : 1 x 240GB
  - Must be SSD (HDD not supported)
  - Recommended to have 2 system disks with the same capacity for high availability.

    > **Note:** Do not configure any hardware RAID on the two system disks, if the system disks are connected to a RAID controller. You must configure the RAID controller and both the system disks as JBOD mode in the hardware system setup.

  - System disk must be less than 800 GB
  - Any additional drives must be removed from the system
- Memory : 128 GB (minimum)
- Intel CPU 32 Cores (minimum) with each core at 2.1GHz or above

Supported Hardware Platforms

Cohesity also extends support for hardware platforms from the following partners as long as they meet the minimum system requirements.

Cohesity certified platforms	Reference Links
C5K CNG2 Compute Node	C5000 Series Setup Guide
Bring your own Hardware (BYOH)	HPE, Cisco, DELL, Fujitsu, or Lenovo - any x86 that meets compute node requirements.

Bring Your Own Hardware (BYOH)	Minimum Specifications*
CPU	2 x 16 Cores (Intel or AMD), min 2.0GHz
Memory	128GB
Boot drives	2 x 240GB SATA/NVMe M.2 SSDs Note: Boot SSD drive size should not be more than 800GB (max limit)
NIC	4 x 10/25GbE (BaseT, SFP, SFP+, etc)

*The requirements can be more unless explicitly specified.

> **Note:** If you use a hardware model other than the Cohesity-certified platforms, it is recognized as 'Product Name : Unknown Platform' and supported as long as it meets the minimum system requirements.

Manage Compute Nodes

> **Note:** Nodes configured with VIPs cannot be designated as or automatically added as compute nodes.

Automatic Addition of Compute Nodes

Adding a physical compute node is seamless for the user with no additional configuration. If a physical node is added, and it meets the minimum specifications for a compute node, it will be automatically added as a Compute Node, unless it is a pre-certified data node platform.

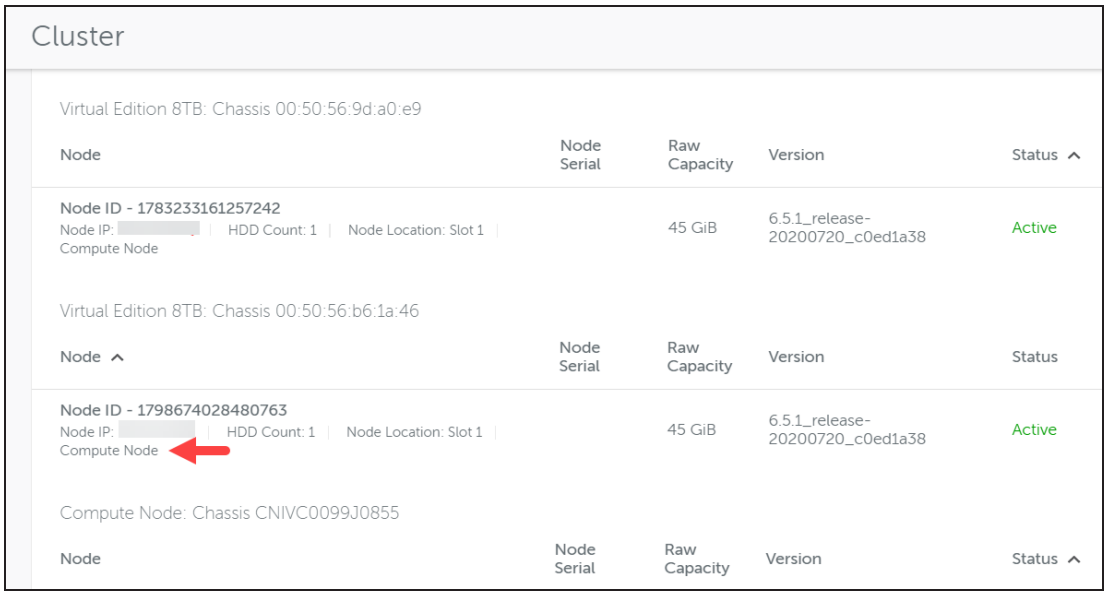

Delete Compute Nodes

Deleting a compute node is like deleting any other node. For more information, see node delete.

Gracefully Remove Nodes from the Cohesity Cluster

A Cohesity cluster node may have multiple Cohesity services such as Data Protection service, I/O operations service, and Indexing service running. Removing the node without gracefully handling these existing services may result in data loss. To prevent this, Cohesity recommends that you remove a node gracefully. During graceful removal, the resources

used by a service running on that node are optimally redistributed to other available nodes, and then the service is stopped. After all the services have stopped, the node is removed, thereby handling the node removal in an orderly manner.

A node is gracefully removed by using the web UI or by running a Cohesity CLI command. The CLI command output displays the progress of the redistribution of resources used by each service. The output helps you understand how long you have to wait for the graceful removal to complete.

> **Note:** Physical Hardware nodes, Virtual Edition nodes, and Cloud Edition nodes can be gracefully removed.

Prerequisites

- You need CLUSTER_MANAGEMENT privilege to perform the graceful removal of nodes.
- Any ongoing upgrade must be canceled before initiating graceful removal. For details on how to cancel an upgrade, see Cancel Upgrade.

Gracefully Remove a Node

To gracefully remove a node from the Cohesity cluster:

1. To view the list of nodes in a Cohesity cluster and identify the unique ID of the node, do the following:
    1. Log into the Cohesity cluster.
    2. Navigate to **Settings** > **Summary**.
    3. Select the **Hardware** tab. The details of all the nodes in the Cohesity cluster are displayed. Make a note of the ID of the node you want to delete.
2. To remove the node, do one of the following:
    1. Using the web UI
        1. Hover over the node you want to remove and click the kebab menu and select **Remove**.
        2. In the **Removal Precheck** dialog, click **Remove** to confirm.
    2. Using the Cohesity CLI
        1. Start the Cohesity CLI. For details on how to do this, see Using the Cohesity CLI.
        2. Run the following command to mark a node in the Cohesity Cluster for removal using the ID of the node you made a note of in Step 1:

```
admin@127.0.0.1>node rm id=<ID of the node>
```

Example

```
admin@127.0.0.1>node rm id=1111222255
```

Sample Output

```
Success: Node ID:1111222255 marked for removal successfully.
```

Check Removal Status

Run the following command to check the removal status of the node.

```
node ls ids=<node id> include-marked-for-removal=true show-stats=true
```

Sample Command Output:

```
node ls ids=1111222255 include-marked-for-removal=true show-stats=true

NODE ID                          : 1111222255

NODE IP                          : 1.11.11.111

REMOVAL STATE                    : MARKED FOR REMOVAL

REMOVAL REASON                   : GRACEFUL REMOVAL

SERVICES ACKED                   : Athena, Librarian, Gandalf, Yoda, Magneto,
Groot, All

SERVICES NOT ACKED               : Hydra,ChunkRepository,Scribe,Apollo

DATA HEALING STATUS              :

Hydra                        : Number of epochs 4

ChunkRepository              : 3.84 MiB(2 files) remaining

Scribe                       : Number of buckets per entity 814
```

The output of the CLI command displays the progress of the node removal as follows:

- REMOVAL STATE— Displays MARKED FOR REMOVAL if removal is in progress. Changes to "OK TO REMOVE" after all the services have stopped and then to "REMOVED" when graceful removal is complete.
- REMOVAL REASON— Is always GRACEFUL REMOVAL.
- SERVICES ACKED—Indicates which services have successfully redistributed the resources and acknowledged the removal.
- SERVICES NOT ACKED—Indicates which services are yet to successfully redistribute the resources and acknowledge the removal.
- DATA HEALING STATUS— Indicates the progress of the ongoing redistribution of the resources for each service. For example, ChunkRepository indicates the progress of the reallocation of files used by the I/O service.

Cancel Removal of a Node

You can still cancel the removal of a node if one or more cluster services are in progress with the redistribution of the resources and are yet to send the acknowledgment.

Run the following command to cancel removal of a node:

```
node rm id=<node-id> cancel=true
```

Example

```
node rm id=1111222255 cancel=true
```

Sample Output

```
Success: Node ID:1111222255 marked-for-removal is cancelled
```

Gracefully Remove Disks from the Cohesity Cluster

A Cohesity cluster node from which you are removing an HDD or SSD disk may have multiple Cohesity services such as Data Protection service, I/O operations service, and Indexing service running. Removing the disks from such nodes without gracefully handling these existing services may result in data loss. To prevent this, Cohesity recommends that you remove a disk gracefully. During graceful removal, the resources used by a service running on the node from which the disk is removed are optimally redistributed to other available nodes. Then the disk is available for removal, thereby facilitating the disk removal in an orderly manner.

A disk is gracefully removed by running a Cohesity CLI command. The CLI command output displays the progress of the redistribution of resources used by each service. The output helps you understand how long you have to wait for the graceful removal to complete.

Prerequisites

- You need CLUSTER_MANAGEMENT privilege to perform the graceful removal of a disk.
- Any ongoing upgrade must be canceled before initiating graceful removal. For details on how to cancel an upgrade, see Cancel Upgrade.

Gracefully Remove a Disk

To gracefully remove a disk from the HDD or SSD disk from the Cohesity cluster:

1. Start the Cohesity CLI. For details on how to do this, see Using the Cohesity CLI.
2. To view the list of disks in a Cohesity cluster and identify the unique ID of the disk, run the following:

   ```
   disk ls
   ```

   From the command output, make a note of the ID of the disk you want to delete.
3. Run the following command to remove the HDD or SSD disk:

```
admin@127.0.0.1>disk remove disk-id=<disk ID>
```

Example

```
admin@127.0.0.1>disk remove disk-id=16
```

Sample Output

```
Success: Disk ID:16 marked for removal.
```

The reallocation of resources starts immediately. The disk is removed once all the services have successfully redistributed the resources and acknowledged the removal.

Check Removal Status

Run the following command to check the removal status of the disk.

```
disk ls ids=<disk id> include-marked-for-removal=true show-stats=true
```

Sample Command Output:

```
disk ls ids=16 include-marked-for-removal=true show-stats=true

DISK ID                          : 16

CURRENT NODE ID                  : 128635482687442

CURRENT NODE IP                  : 10.1.1.11

REMOVAL STATE                    : MARKED FOR REMOVAL

REMOVAL REASON                   : GRACEFUL REMOVAL

SERVICES ACKED                   : Groot, Gandalf, Magneto, All, Athena

SERVICES NOT ACKED               : ChunkRe-
pository,Scribe,Apollo,Yoda,Hydra,Librarian

DATA HEALING STATUS              :

ChunkRepository             : 490.95 MiB(38 files) remaining

Scribe                      : Number of buckets per entity 768
```

The output of the CLI command displayed the progress of the disk removal as follows:

- REMOVAL STATE— Displays MARKED FOR REMOVAL if removal is in progress. Changes to “Removed” when graceful removal is complete.
- REMOVAL REASON— Is GRACEFUL REMOVAL or HEALTHFAIL. The removal reason is HEALTHFAIL when the disks are removed if a disk is found inoperable during a self-initiated health check.
- SERVICES ACKED—Indicates which services have successfully redistributed the resources and acknowledged the removal.

- SERVICES NOT ACKED—Indicates which services are yet to successfully redistribute the resources and acknowledge the removal.
- DATA HEALING STATUS— Indicates the progress of the ongoing redistribution of the resources for each service. For example, ChunkRepository indicates the progress of the reallocation of files used by the I/O service.

> **Note:** Upon completion of the scribe shuffle, the newscribe process may experience a crash. This is a normal occurrence on nodes when the scribe shuffle enables them to utilize the new buckets. However, this does not result in a restart of all services on the node.

Cancel Removal of a Disk

You can still cancel the removal of a disk if one or more cluster services are in progress with the redistribution of the resources and are yet to send the acknowledgment.

Run the following command to cancel the removal of a disk:

```
disk remove disk-id=<disk-id> cancel=true
```

Example

```
disk remove disk-id=16 cancel=true
```

Sample Output

```
Success: Disk ID:16 marked-for-removal is cancelled
```

Hardware Node Detection Issues

If you encounter issues while creating Nodes, read the following sections for possible resolutions:

Verify the following:

- Multicast is enabled on the network switch that the Cohesity cluster Nodes are connected to.
- All of the following are reachable from the 10GbE switch:
  - Default gateway
  - DNS server
  - NTP server

No Node is Detected

Possible cause

The network switch configuration prevented mDNS from being received/transmitted.

Resolution

Attempt to connect the laptop to another Node's management port to narrow down which network cable, switch port, or Node network port has an issue so you can resolve it.

Some Nodes are Not Displayed

Possible cause

The Node’s network connection has an issue.

Resolution

Check the Node’s cable connection, network switch port or the Node’s network interface. Checking indicator lights may help you identify the issue.

> **Note:** If your network switching is configured for VPC (Virtual Port Channel) or MLAG (Multi-Chassis Link Aggregation), verify that the bond mode on each node is set to mode 4. If the bond mode has not been changed to mode 4 to match the VPC/MLAG configuration, the node may not appear during detection.

Use Cohesity Data Cloud (Self-managed) in DoD Mode

Cohesity Data Cloud (Self-managed) is DoDIN APL certified.

To use Cohesity Data Cloud (Self-managed) as per DoDIN APL requirements and recommendations, you need to enable DoD mode in the Cohesity cluster.

As part of the DoDIN APL requirements, Cohesity provides the following security features:

- Login banner support
- Transfer audit logs to a Syslog server over a secure TLS connection.
- Certificate mapping-based authentication support.
- Support to configure the banner on the Cohesity Dashboard.
- Server-side session management.
- Node-to-node traffic encryption.
- Certificate path validation.

Enable DoD Mode

You can enable DoD mode using the Cohesity CLI.

> **Warning:** Once you have enabled DoD mode, you cannot disable it. For node-to-node traffic encryption, you must configure the required certificates before enabling DoD mode. For more information, see Node-to-Node Encryption.

When you enable DoD mode, by default:

- The **cohesity** user is disabled, and you cannot use the **cohesity** user account to access the Cohesity bash shell using SSH. Instead, you need to use the support user account to log into the Cohesity bash shell using SSH.
- You can change the **support** user password only once within a day.
- The **support** user can run the "`sudo`" prefixed commands only after providing the **support** user password. The sudo session is valid for 15 minutes, and you will be prompted to enter the password again.
- The **cohesity_console** account is disabled, and instead, you can use the root account to log into the console.
- The data transfer between the nodes in the Cohesity cluster is encrypted.

Prerequisites

Dedicated partitions are required to support DODIN RHEL STIG requirements for Operating System partitions. Before you enable the DOD mode, you must run a script provided by the Cohesity Data Cloud (Self-managed) to check if the system disk on the Cohesity node contains the dedicated partitions. To run the script:

1. Log on to the Cohesity node bash shell using SSH.
2. Run the following script:

```
/home/cohesity/bin/validate_dodin_compliant_partitions.sh
```

The script validates if the following dedicated partitions are present in the the system disk on the Cohesity node:

- `/var`
- `/spare/var`
- `/var/log/audit`
- `/cohesity_users_home`

If the validation fails, it implies that the Cohesity node does not contain the dedicated partitions. Therefore, you need to perform ISO installation on the Cohesity node to get all the security benefits listed by DODIN RHEL STIG for Operating System partitions.

Enable DOD Mode

To enable DoD mode:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.
2. Run the following command to enable DoD mode:

```
$iris_cli security-config update-security-mode security-mode-dod=true
Configuring security mode...Success: Security mode DOD will be turned
on after rolling reboot.
```

The Cohesity cluster goes into a rolling reboot. Wait until the rolling reboot is complete.

3. Verify if the DoD mode is enabled. Run the following command:

```
$iris_cli cluster info
```

Node-to-Node Encryption

By default, when you enable DoD mode on the Cohesity cluster, the user data exchanged between the nodes are encrypted.

> **Note:** Internal data such as statistics data or other non-critical data is not encrypted.

Before you enable DoD mode you need to configure the nodes to use certificates to encrypt the data exchanged between the Cohesity cluster nodes.

Configure Certificate for Node-to-Node Encryption

To use a certificate for node-to-node encryption, you must manually create a certificate signing request, generate the certificate, and import the certificate on each node in the Cohesity cluster.

Prerequisites

Before you configure a certificate for node-to-node encryption, ensure the following prerequisites:

- Verify if the domain name is set for your Cohesity cluster using the following command:

```
$ cluster ls-domains
```

For example:

```
$ cluster ls-domains
CLUSTER ID                : 4759089959284457
CLUSTER NAME              : cohesity
DOMAIN NAMES              : eng.cohesity.com, corp.cohesity.com
```

- Update the hostnames of the nodes in the Cohesity cluster using the following command:

```
$ node-csr update-hostnames node-ids=<node_id1,node_id2> host-
names=<hostname_1, hostname_2>
```

Where:

- `node-ids` - The node id of the nodes in the Cohesity cluster. Provide a comma separated list of the node ids.
- `host-names` - Define the hostnames for the nodes in the Cohesity cluster. Provide a comma separated list of the hostnames for the node ids.

- Ensure that the certificate used for node-to-node encryption includes SSL Client and SSL Server key usage. If these key usages are missing, node-to-node communication will not work correctly.

Step 1: Create a Certificate Signing Request

Perform the following to create a certificate signing request (CSR) on each node in the Cohesity cluster.

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.
2. Run the following command to create a certificate signing request:

```
$node-csr create city=<city_name> state=<state_name> country=<country_
name> organization=<org_name> organization-unit=<org_unit_name> dns-
names=<dns_name>
```

Sample Output:

```
CSR :-----BEGIN CERTIFICATE REQUEST-----
MIIC5TCCAc0CAQAwVTELMAkGA1UEBhMCVVMxETAPBgNVBAoMCENvaGVzaXR5MRIw
EAYDVQQLDAlDb3JlSW5mcmExCzAJBgNVBAgMAkNBMRIwEAYDVQQHDAlDdXBlcnRp
bm8wggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQDWDkZokJkovKdSj6Ni
5Shqkjx+BUsgs0DNS6Kfpo1ZlFEeDbwQ9ThgycNU8PWtNyb8bNwzPlMvbPtB1o4C
u+ACukfASH9blJyOBP4JrUAouHTHK73gDU0hlnN+PrVjMNtgRvAh68vJtLudtK8j
3pcGE3eWEMolNC8M3h5bLUF+b9d9UyCvrJIjcwrUkI+9EI5OJNCQIxLdZ1XjTQdf
khSnarh2wNTrF1AuBrmJdTY0S5FG9s0TD8UdDE5c90vfNVVlLN02NK06s3wVtskL
j+eF7DNUJrL1gftx7i0s5ksdLPYxzVo8ap5VW4OqSZDtPv/zllBK9u/6FAqqDbC1
q+nXAgMBAAGgSzBJBgkqhkiG9w0BCQ4xPDA6MDgGA1UdEQQxMC+CLWNvaGVzaXR5
LTAwNTA1Njg5MzFlOC1ub2RlLTQuZW5nLmNvaGVzaXR5LmNvbTANBgkqhkiG9w0B
AQsFAAOCAQEAt+vgpETeZqPKJaQA6I+jEKi6hcHaNUFTHrPrb7QMdzzzB9AIyoYc
onCctHFkpt77NNqmmq+3K3+J2YBlpIwDVjKu5vf1rY3BMeeo6RlOsWJH62cdQbS9
V3xr+dkF2tnVXrsoMizh2IaVBSVfbzllIVLwfayuRZzKeAt13XVQsqAnBFNdSyFO
d9R42rwPMfJ1MzfFz/jRBFIrbknkDKNdOJ+80AyRKFYQSvNvygD+0t+rww9EcBgq
2Jw06WFDmFMNxPnSLe0/Mt7UUt/WCTU5bs9Coq8cat4nBNeebJ7GgubE+7G3QhcZ
KzyUXKkcY/Md/n/k2EHDrF5thGX8X8qugw==
-----END CERTIFICATE REQUEST-----
```

The following table lists the attributes required to create the certificate signing request:

Attribute	Required/Optional	Description
city	Required	Identifies the city where the company is located.
common-name	Optional	Specifies a context for the certificate. For example, the name of the node to which the certificate must be assigned. Ensure that the common name matches the node name, including uppercase and lowercase letters. Any difference in letter case will cause node-to-node communication to fail.
country	Required	Identifies the country where the state is located. You can specify it as two letter code defined by the ISO standard.
dns-names	Optional	Concatenate the hostname with the first domain name. For more information, see Prerequisites. For example: `cohesity-0050568931e8-node-2.eng.cohesity.com` Though `dns-names` attribute is optional. Cohesity recommends that you use the `dns-names` attribute.
email-address	Optional	An alternative subject name component that must be included in the certificate.
host-ip	Optional	An alternative subject name component to be included in the certificate. It is used to uniquely identify the node.
organization	Required	Specifies the name of the company.
organization-unit	Required	Specifies the department or business unit in the company that owns the node.
state	Required	Identifies the state where the city is located.

Step 2: Sign the CSR with a CA

You can sign the CSR using a third-party Certificate Authority (CA) or a self-signed CA using OpenSSL commands.

Step 3: Verify the CA Certificate

Run the following command to check if the CA certificate is valid. The CA certificate is valid if the output is OK.

```
openssl verify -verbose -CAfile <path to ca cert>  <path to signed cert>
```

Example

```
openssl verify -verbose ca.crt server.crt
```

Sample Output for a valid CA certificate

```
ca.crt: C = US, ST = CA, L = SJ, CN = Cohesity Inc Server, O = Cohesity
OK
```

Sample Output for an invalid CA certificate

```
server.crt: C = US, ST = CA, L = SanJose, CN = localhost, O = Cohesity
error 20 at 0 depth lookup:unable to get local issuer certificate
```

Step 4: Import the Signed Certificates to the Nodes

After the certificate signing request (CSR) corresponding to the node is signed, import the signed certificate and the CA certificate onto each node in the Cohesity cluster.

Perform the following on all the nodes in the Cohesity cluster:

1. SCP or copy the signed certificate to node in the Cohesity cluster. Enter that you copy both the signed and CA certificates to the node.
2. Change the certificate file permission to 644.
3. Import the signed and CA certificates to the node using the following command:

   $node-csr import signedcert-path=<signed_cert_path> cacert-path=<ca_cert_path>

   Where:

   - `<signed_cert_path>` - The absolute path of the signed certificate file that must be imported to the node.
   - `<ca_cert_path>` - The absolute path of the CA certificate that was used to sign the CST for the node.

Certificate Path Validation

Certificate path validation is necessary to confirm if the certificate path is valid so that the Cohesity node can accept the end entity certificate. Cohesity uses Online Certificate Status Protocol (OCSP) for obtaining the revocation status of an X.509 digital certificate. Issuer certificate is required to recreate OCSP request for certificate path validation.

> **Note:** Cohesity supports only HTTP protocol URL for issuer certificates.

Configure a Login Banner

You can customize the login banner of Cohesity Dashboard to present a statement to users during the login process, requiring them to acknowledge and agree to the message before logging in.

When you configure the login banner, the message will appear before the user logs on to the Cohesity Dashboard. To configure a login banner, see ***Optional Settings*** in Configure Settings.

Configure Banner Message

You can present a banner message to a user who is using the Cohesity Dashboard. You can configure the banner and the message using the Cohesity CLI.

To configure the banner message:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.
2. Run the following command to configure the banner and the message:

```
$iris_cli security-config update is-data-classified=true classified-
data-msg=”This is a sample message”
DATA CLASSIFICATION       :
IS DATA CLASSIFIED        : true
CLASSIFIED DATA MSG       : This is a sample message
```

Certificate Mapping-Based Authentication

If the Cohesity cluster is joined to an Active Directory (AD) domain, Active Directory (AD) users can authenticate to the Cohesity cluster using certificate mapping-based authentication.

In certificate mapping-based authentication, the AD user is authenticated based on the attribute mappings you have configured for the certificate and AD. In the Cohesity cluster, you can define the following attributes of the certificate and AD that must map for the user authentication:

- Certificate Attributes:
  - CommonName
  - EmailAddress

    - UserPrincipalName
- Active Directory Attributes:
    - SamAccountName
    - UserPrincipalName

For more information, see Certificate Mapping-Based Authentication.

Server-Side Session Management

Cohesity provides various session management related settings for server-side session security.

The following table lists the settings:

Settings	Description	Example
Session absolute timeout	A timeout after which a session is closed no matter whether there is user activity or not. You can specify the absolute session expiration time in seconds. The default value is 86400 seconds and the DoD mode value is 3600.	`$iris_cli security-config update session-absolute-timeout=3600` `SESSION CONFIGURATION : SESSION ABSOLUTE TIMEOUT : 3600`
Session inactivity timeout	A timeout after which a session is closed if the user is inactive. You specify the inactivity time in seconds after which the session expires. The default value is 3600 seconds and the DoD mode value is 600.	`$iris_cli security-config update session-inactive-timeout=600` `SESSION CONFIGURATION : SESSION INACTIVITY TIMEOUT : 600`
Limit sessions	Restrict the number of active sessions allowed per user and system-wide. The default value is set to false.	`$iris_cli security-config update limit-sessions=true` `SESSION CONFIGURATION : LIMIT SESSIONS : true`
Session limit per user	The maximum number of active sessions allowed per user. The default value is set to 10. **Note:** Ensure that you have enabled limit sessions.	`$iris_cli security-config update session-limit-per-user=40` `SESSION CONFIGURATION : SESSION LIMIT PER USER: 40`

Settings	Description	Example
Session limit system-wide	The maximum number of active sessions allowed system-wide. The default value is set to 100000 and the DoD mode value is 60000. **Note:** Ensure that you have enabled limit sessions.	`$iris_cli security-config update session-limit-system-wide=60000` `SESSION CONFIGURATION : SESSION LIMIT SYSTEM WIDE: 60,000`

Revoke active sessions	Revoke all the active sessions of a user or all active system-wide sessions.	If you provide the user-sid of the user, all active sessions of the user are revoked.
		
		For example:
		
		`$iris_cli user delete-sessions user-sid="S-1-100-21-86254287-1950289-11"`
		`Success: Deleted all active sessions of user with sid S-1-100-21-86254287-1950289-11`
		
		If you do not provide the user-sid of the user, then system-wide active sessions are revoked.
		
		For example:
		
		`$iris_cli user delete-sessions`
		`Success: Deleted system wide all active user sessions`

Note:
These are the two APIs used to get and update the session related settings.

1. GET /v2/security-config
2. PUT /v2/security-config

For more information, see Create Session.

Securely Transfer Audit Logs to an External Syslog Server

You can add a Syslog server to the Cohesity cluster and securely transfer the following logs to an external Syslog server:

- Cluster audit logs
- Operating System audit logs
- Filer audit logs
- Web Server audit logs

The logs are transferred over a TLS connection from the Cohesity cluster to the external Syslog server.

You can add a Syslog server to the Cohesity cluster using the Cohesity CLI.

To securely transfer logs to an external Syslog server:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.
2. Run the following command to add a Syslog server:

   ```
   $iris_cli syslog add address=<ip_address/hostname> port=<port>
   protocol=<tcp/udp> ca-certificate-path=<certificate-absolute-path>
   ```

   Where:

   - `<ip_address/hostname>` - The IP address or hostname of the Syslog server.
   - `<port>` - The port for data transfer.
   - `<tcp/udp>` - The protocol used for the data transfer. The supported values are TCP or UDP.
   - `<certificate-absolute-path>` - The absolute path of the CA certificate on the Syslog server.

   > **Note:** The Syslog server CA certificate file path can be local to the current node.

About Node-to-Node Encryption

The Zero Trust Architecture of the Cohesity Data Cloud (Self-managed) is strengthened with certificate-based authentication for communications between services on a Cohesity cluster. Cohesity provides X509 certificates that are used to:

- Authenticate services for communication with another service
- Encrypt the user data exchanged between the nodes on a Cohesity cluster.

The communication between services of the same or different nodes happen through TLS and Mutual TLS (mTLS).

> **Important:** In new Cohesity cluster from 7.1.1 deployments, node-to-node encryption is enabled by default. However, if you are upgrading your existing Cohesity clusters from version 6.8 or later, you need to manually enable node-to-node encryption. For details, refer to Upgrade Scenarios.

This applies when a service communicates with another service on the same or a different node, it first verifies the other service through certificates and then encrypts the data before communication.

> **Note:** If node-to-node-encryption is enabled on your Cohesity cluster before upgrade, then it will remain enabled after upgrade. For details, refer to Upgrade Scenarios.

Each service on a Cohesity node is assigned a unique X509 certificate that holds the private key for authentication during a TLS handshake. This certificate is rotated every 12 hours automatically .

> **Note:** The mTLS-based node-to-node encryption is not enabled by default when you create a new cluster. Once you enable node-to-node encryption on your Cohesity cluster, it cannot be disabled.

Upgrade Scenarios

Cohesity Cluster Version	DoD Mode	Node-to-node Encryption	Post-Cluster Upgrade Actions
Latest patch update for 6.8.1 or any earlier version	Disabled	Disabled	Follow the steps to Enable Node-to-node Encryption
Latest patch update for 6.8.1 or any earlier version	Disabled	Enabled	Run the command on Cohesity CLI to enable mTLS-based node-to-node encryption: `iris_cli cluster tpenc enable-encryption=true`

Cohesity Cluster Version	DoD Mode	Node-to-node Encryption	Post-Cluster Upgrade Actions
Latest patch update for 6.8.1 or any earlier version	Enabled	Enabled	Run the command on Cohesity CLI to enable mTLS-based node-to-node encryption: `iris_cli cluster tpenc enable-encryption=true` For more information, see *Node-to-Node Encryption* in Use Cohesity Data Cloud (Self-managed) in DoD Mode.

Enable Node-to-node Encryption

> **Important:** By default, new Cohesity clusters deployed with version 7.1.1 or later have node-to-node encryption enabled. However, when upgrading a Cohesity cluster from version 6.8 or later to version 7.1.1 or above, only Gandalf RPC encryption is enabled by default. You need to manually enable internode traffic encryption. This ensures that all internode data traffic is encrypted, making the Cohesity Platform compliant with Zero Trust principles.

To enable node-to-node encryption:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command to enable node-to-node encryption:

   ```
   $iris_cli cluster stop
   $iris_cli cluster tpenc enable-encryption=true
   $iris_cli cluster start
   ```

Verify Service and Node-to-node encryption is Enabled

You can run the following commands on Cohesity CLI to check:

1. Services on your Cohesity cluster are associated with X509 certificates.
2. Node-to-node encryption is enabled on your Cohesity cluster.

Command 1

```
iris_cli -username=<Value> -password=<Value> service-identity status
```

Sample Response

```
[cohesity@cluster-bqkp45044016-node-3 ~]$ iris_cli service-identity status
Username: admin
'admin' Password:
GLOBAL-ENABLED : true
[cohesity@cluster-bqkp45044016-node-3 ~]$
```

Quorum Protected Operations

Quorum approval on Helios is an authorization model that ensures that a predefined quorum of approvers must approve sensitive or privileged operations requested by Helios before those operations are executed.

Quorum approvals help you eliminate the risk associated with a unitary, highly-privileged administrator - specifically to prevent a rogue, poorly trained, or compromised administrator from performing sensitive or privileged operations on the Cohesity platform without authorization or oversight.

Any Cohesity cluster-related operations that are quorum-protected in Helios cannot be initiated from the Cohesity Platform UI or CLI. Instead, you need to initiate these operations from Helios. For more information on Quorum Approvals on Helios, see Quorum Approvals.

> **Important:** If you disconnect the Cohesity cluster from Helios, the quorum-protected cluster operations continue to be restricted from initiation from the Cohesity Platform UI or CLI.

Antivirus Software to Protect Cluster Data

You can enable antivirus software to protect your data on the Cohesity cluster from viruses. Once an antivirus software scans a file and detects a threat, the Cohesity cluster reports the threat on the Cohesity Dashboard. The administrators can either quarantine, unquarantine, reset, or delete the infected file.

Perform the following steps to enable antivirus scanning and manage threats:

1. Add an Antivirus Provider
2. Enable Antivirus Scan on a View
3. View and Manage Threats

Add an Antivirus Provider

To perform an antivirus scanning on your data in the Cohesity cluster, you must first add the antivirus software to the Cohesity cluster. You can add the software in one of the following ways:

- Register External Antivirus Software
- Install the Antivirus App from the Cohesity Marketplace

Register External Antivirus Software

1. Navigate to **SmartFiles > Antivirus** and click the **Providers** tab.
2. Click **Add External Provider** in the upper-right corner of the page.
3. In the **New External Provider** field, enter the name of the antivirus software provider of your choice.
4. (Optional) Describe the provider.
5. Enter the ICAP Server URL in the following format: `icap://<server>:<port>/<uri>`

Option	Description
icap	Indicates the required protocol identifier that enables communication with external servers hosting third-party antivirus software.
server	Specifies the domain name or IP address of the external server (ICAP server) hosting the third-party antivirus software.
port	Specifies the port that the external server monitors for incoming requests. If not specified, the default icap port, 1344 will be used.
uri	Specifies the path of the ICAP service.

6. (*Optional*) Click **Add** to register more than one external ICAP server.
7. Click **Check ICAP Server Connection Status** to check the status of the external antivirus software.

   If the connection is not successful, ensure that you have provided the correct ICAP URL and the external server is up and running.
8. Click **Add**.

   The antivirus software is registered on the Cohesity cluster.

Install the Antivirus App from the Cohesity Marketplace

1. Navigate to **SmartFiles > Antivirus** and click the **Providers** tab.
2. Click **Install Antivirus App** at the top right of the page.

You will be directed to the Cohesity App Marketplace.

3. Click **GET APP** for the Antivirus App you want to install on the Cohesity cluster.

   You will be directed to the Helios login page.

4. Log in to Cohesity Helios, select the Cohesity cluster on which you want to install the Antivirus App, and click **Install**. You also have the option to install the App later.
5. Go back to the Cohesity cluster and go to **Marketplace > My Apps**.

   The page displays the list of Apps that are available on the Cohesity cluster.

6. Click **Run App** to create an instance of the Antivirus App.

   The instance created will be displayed as an Antivirus provider on the **Providers** tab.

Enable Antivirus Provider

1. Navigate to **SmartFiles > Antivirus** and click the **Providers** tab.

   The page displays the list of antivirus providers that can be enabled for antivirus scanning.

2. From the Action menu, toggle on the **Enable** option.

   You can enable only one antivirus provider group at a time.

Enable Antivirus Scan on a View

Once you install the antivirus software on the Cohesity cluster, you can enable the antivirus scan on a View.

To enable antivirus scan on a view:

1. Navigate to **SmartFiles > Antivirus** and click the **Secured Views** tab.
2. Click **Add a View** at the top right of the page.
3. From the View drop-down list, select a view for which you want to perform the scanning of files.

   Antivirus scanning is supported only on Views with SMB protocol. Therefore, the drop-down list displays the Views which support SMB protocol.

4. Provide the following details either by clicking **Edit All** or by clicking the **Edit** option

for each scan settings option:

Option	Description
Scan on	Antivirus scans can be configured to send the files for scanning when they are closed, opened, or both. In each case, all the files to be scanned are added in a queue and sent to the antivirus provider(s) for scanning, as per the queue order. • **Close**: Files are scanned after they are closed (Default) With this configuration, you can access the files without any delay as file scanning is done after the user's file operation. This type of scanning is faster but less secure. • **Open**: Files are scanned before they are opened. With this configuration, you can access the files once they are scanned. You must wait for the file scan to complete. This type of scanning is slower but more secure. • **Open and Close**: Scans the files before they are opened and after they are closed. This type of scanning is the slowest and most secure, as the scanning queue includes all the live files (opened and closed). Cohesity recommends scanning the file after they are closed as this configuration ensures that the files are accessible even if the antivirus provider(s) or scanning network is down, unavailable or unreachable.
Files to Scan	• **All Files**: Select this option to scan all the files in the View. • **All Files Except**: Select this option to exclude files in the view from scanning. You can specify the file extensions (comma-separated) to be excluded in the **File extensions to skip** field. • **Only These Files**: Select this option to scan only the selected files. You can specify the file extensions (Comma-separated) to be included in the **File extensions to scan** field.
	Exclude files based on file size: Enable this option to exclude files based on the file size. In the **Max File Size** field, you can specify the file size limit for scanning. The default file size limit is 50 MB. The maximum file size limit is 125 MB.
Scan Timeout	Set the maximum amount of time (in seconds) that a scan can take before timing out.
If Scan Cannot Be Completed	Blocks file access if the antivirus servers are unavailable or are unable to scan the file due to any issue. This option is disabled by default to ensure that the files stay accessible even if the antivirus provider(s) or scanning network is down, unavailable, or unreachable.

5. Click **Add**.

You can also enable antivirus scan while creating a View or by editing an existing View.

View and Manage Threats

After you enable the antivirus scan on the View, you can view the scanned report from the **Threats** page and take action on the infected files.

To view the scanned report:

Navigate to **SmartFiles > Antivirus** to view the **Threats** page.

Alternatively, on the **Secured Views** page, click the view for which you want to view the reports.

The Threats page displays the following details:

- The total number of quarantined files
- The total number of unquarantined files
- The total number of views that are affected by the virus
- Option to filter the report based on the status of the file and View
- Information on the files that are affected:

Field	Description
Filename	The files that are infected, along with the file path. A lock icon will be displayed along with the file name if you have enabled the **File DataLock** option while creating the View.
View	The View name of the files. A lock icon will be displayed along with the View name if you have cloned the View as DataLock.
Threat	The threat type or virus signature along with the ICAP server URL that detected the threat.
Provider Name	The antivirus software provider that detected the threat.
Scan Time	The time at which the file is scanned for viruses.
Status	The status of the threat: • **Quarantine**: The file is determined to be unsafe. All access to the file is blocked until the administrator manually changes the file state. • **Unquarantine**: Allow file access. The file is not scanned again until the file state is reset.

Manage Threats

You can manage threats by performing the following actions from the **Threats** page:

Action	Description	How to
Unquarantine Files	Move the files from the quarantined state to allow file access. Unquarantined files are not scanned again.	Select the checkbox next to file names and then click **Unquarantine Files**.
Delete Files	Deletes the quarantined files. You cannot delete unquarantine files.	Select the checkbox next to file names and then click **Delete Files**.
Reset Files	Reset the unquarantine files. You cannot reset quarantine files	Select the checkbox next to file names and then click **Reset Files**.

Cohesity Licensing

The Cohesity Licensing system helps customers and partners ensure that Cohesity software is being used in accordance with the terms and conditions of the license and the usage is in compliance with the entitlement. The Cohesity licensing model gives the flexibility to use certain features or services of the Cohesity software, depending on which license(s) the customer is entitled to.

All Cohesity clusters must be licensed. License usage is metered on each Cohesity cluster and the usage is aggregated in Cohesity Helios. Alerts are triggered or GUI notifications are displayed in one of the following scenarios:

- License usage violation—If the aggregate metered usage is nearing the entitlement or has exceeded the entitlement
- A cluster is set up without a license
- Clusters are not reporting their usage information to Helios
- The license file on a particular cluster has expired or is about to expire

You must set up licensing on the Cohesity cluster. For more information on setting up licensing on each cluster, follow the instructions below.

You can monitor the license usage and view entitlement details in Helios. Starting from the Cohesity 6.6.0c release, details about license usage and entitlement are displayed only on the **License Usage** page in Helios. On the Cohesity cluster, the **License** page does not display details about the license usage and entitlement. For more information, see View License Usage and Entitlements.

License Modes

Before you start using the Cohesity licensing feature, use the descriptions in the following table to decide which Licensing Mode is suitable for your business requirement:

Cohesity cluster State	How Licensing Works	How Audits are Conducted
Manage in Cohesity Data Cloud (SaaS)– Cluster remains connected with Helios at all times. All usage data is automatically shared with Helios.	License is automatically generated in Helios. The license is then reflected in the connected cluster automatically. There is no need for any manual intervention of operations. For details on how to configure this mode, see Choose the Licensing Mode .	Audit reports are automatically uploaded to Helios on a daily basis. **Note:** The audit report contains details such as cluster ID, timestamp, license usage data, and entitlements.

Cohesity cluster State	How Licensing Works	How Audits are Conducted
Manage on Cluster–If you are unable to share cluster management information with Cohesity Data Cloud (SaaS) or if you cannot connect your cluster to the internet for security reasons, you can consider keeping your cluster disconnected from Helios. Cohesity provides the following options: • Manually license your cluster–Every 180 days, you must manually download the audit report from the Cohesity cluster, upload the audit report on Helios, generate the new license file on Helios, and upload the new license file on the Cohesity cluster. • Helios Smart Licensing - Enabling Helios Smart Licensing may enable transmitting some telemetry data to Cohesity. For details on turning off this feature, see page 6 in Cohesity Cluster Secure Phone-Home Channel Security Brief.	**Manual Licensing**–Cluster remains disconnected from Helios and you manage the licenses manually. Perform the following steps: 1. Download the audit report which captures the license usage information from the cluster and upload it to Helios on a monthly basis. 2. Download the updated license files from Helios and upload them to the cluster. For details on how to configure this mode, see Configure Manual Licensing Mode.	Generate and download the license file from Helios. The license file allows usage on the cluster for up to 180 days. Every month, Cohesity recommends that you: 1. Retrieve the audit report from the cluster. 2. Upload it manually to Helios. 3. Generate a new license file. 4. Apply the new license file on the designated cluster.
	Helios Smart Licensing–You connect your cluster with Helios for licensing purposes. If you wish to keep your cluster connected with Helios only for licensing purposes, you can avoid the above steps, and the licensing data between your cluster and Helios is automatically synchronized. For details on how to configure this mode, see Configure Helios Smart Licensing Mode .	Audit reports are automatically uploaded.

Configure License During Cluster Creation

During the initial setup of the Cohesity cluster, upon logging in for the first time, you are presented with the End User License Agreement (EULA). Following the acceptance of the EULA, a prompt to configure the license appears. You can opt to configure the license immediately or defer it for later configuration. Regardless of the chosen timing, it is advisable to explore the available licensing modes and select a mode aligned with your business needs before initiating the license configuration.

> **Note:** If the Cohesity cluster is disconnected from Helios, the LicenseAboutToExpire alert is triggered when the license is about to expire. For more information, see the KB article.

Prerequisites

Before you configure the license, confirm the following access:

- On the cluster, you must be logged in as a user with the Admin role, or a custom role that includes the Manage Cluster privilege.
- In Helios, you must be logged in as a Helios administrator with access to **Home** > **License Usage**.

> **Note:** A Helios Super Admin role is not required to configure licensing. Super Admin controls Helios account access management and is separate from the cluster-side Admin role described above.

Choose the Licensing Mode

Follow these steps to choose between configuring the license immediately by selecting a license mode and deferring the license configuration for a later time:

1. Log in to the Cohesity Cluster.
2. Read the **EULA License Agreement** and click **Agree**.
3. Do one of the following:
    - If you want the cluster to remain connected to Helios and automatically synchronize operations such as downloading audit reports and uploading license files, perform the following:

1. Click **Manage in Helios SaaS** and then **Connect to Helios SaaS**.

2. On the **Connect with Helios** page, select the appropriate access permission—**Full Access** or **View Only** and click **Continue**.
3. You will be redirected to the Helios login page.
4. Enter valid Helios credentials and click **Connect** to establish connection with Helios.

   Once the clusters are connected to Helios successfully, all operations such as generating the license file in Helios, downloading the license file on the cluster, and uploading this license file in Helios are all performed automatically. No manual intervention is required.

   You can access the **License** page on Helios to check the license usage against the purchased entitlement.

- If you are unable to share cluster management information with Cohesity Data Cloud (SaaS) or if you cannot connect your cluster to the internet for security reasons and are considering keeping your cluster disconnected from Helios, perform the following:

1. Click **Manage on Cluster**. Only the **Single-Cluster Management** icon remains enabled.
2. Click **Connect to Cluster**.
3. On the **License** page, select one of the following licensing modes: Helios Smart Licensing, Manual Licensing, and Dark Site.

   To manage the licenses later, select **Skip licensing at this time**. For details on how to configure the license at a later point in time, see Configure License After Cluster Creation.

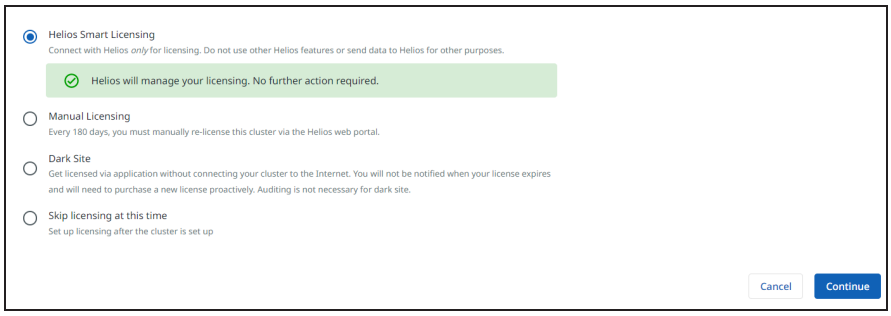

Configure Helios Smart Licensing Mode

Perform the following steps to configure Helios Smart Licensing:

1. Follow the steps in the Choose the Licensing Mode section and select Helios Smart Licensing at Step 3.
2. On the **Management Options** page that appears, select the **Connect to Helios only for licensing** option.
3. Click **Continue**.

   You are redirected to the Helios page to enter your credentials.
4. Click **Connect** to establish the connection with Helios.

   Only the licensing information is shared with Helios.

Configure Manual Licensing Mode

Configuring the Manual Licensing Mode requires you to download the License Usage Report from the Cohesity cluster and upload it to Helios. You must also generate the License File in Helios and upload it to the Cohesity cluster. As these steps require switching between the Cohesity cluster and Helios, Cohesity recommends that you remain logged in to both the Cohesity cluster and Helios until you complete the Manual License configuration.

Perform the following steps to configure Manual Licensing:

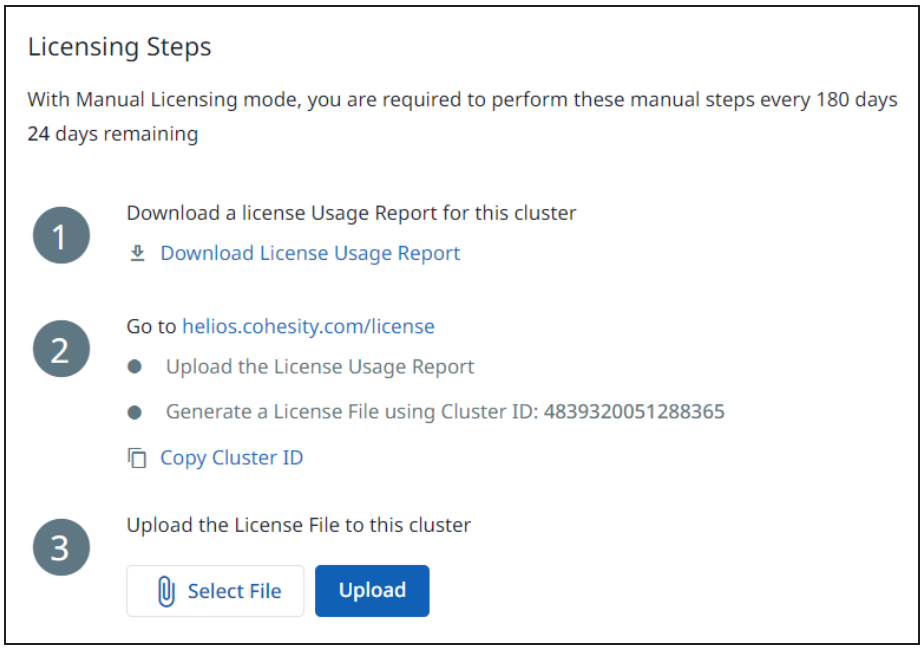

1. Follow the steps in the Choose the Licensing Mode section and select Manual Licensing at Step 3.
2. On the **Licensing Steps** dialog, click **Download License Usage Report**.

   The License Usage Report file or the Audit Report is downloaded to your local machine. The Audit Report contains details such as cluster ID, timestamp, license usage data, and entitlements.
3. To upload the Audit Report:
   1. Log in to Helios and navigate to **Home** > **License Usage**. The license summary of all clusters connected to Helios is displayed.
   2. Click **Upload Audit Report**.
   3. In the **Upload Audit Report** dialog, upload the Audit Report downloaded in step 2.
   4. Click **Submit**.

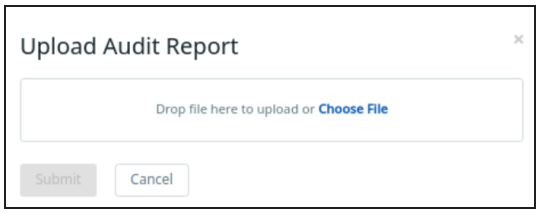

4. To generate the License File:
   1. In your Cohesity cluster, navigate to **Settings** > **License**. On the **Licensing Steps** page, select **Copy Cluster ID**.
   2. In Helios, click **Generate License File**.

The **Generate License File** dialog appears and prompts for the Cluster ID.

3. Paste the Cluster ID copied in Step 4 (a) and click **Generate**.

   The License File is downloaded to your local machine.

5. To upload the License File:
   1. In your Cohesity cluster, navigate to **Settings** > **License**.
   2. On the **Licensing Steps** page, click **Select File**.
   3. Browse to select the License File downloaded in Step 4 (c) and click **Upload**.

Once Manual Licensing is configured, you must download the audit report from the Cohesity cluster and upload it to Helios every 180 days.

Configure License After Cluster Creation

Before you configure the license for your Cohesity cluster, Cohesity recommends that you learn about the licensing modes (in the License Modes section) and choose a Licensing mode suitable for your business requirements *before* you configure the license.

Prerequisites

Before you configure the license, confirm the following access:

- On the cluster, you must be logged in as a user with the Admin role, or a custom role that includes the Manage Cluster privilege.
- In Helios, you must be logged in as a Helios administrator with access to **Home** > **License Usage**.

> **Note:** A Helios Super Admin role is not required to configure licensing. Super Admin controls Helios account access management and is separate from the cluster-side Admin role described above.

Choose a Licensing Mode after Cluster Creation

If you chose **Skip License** when you created the cluster, perform the following to configure licensing for your Cohesity cluster:

1. Log in to the Cohesity cluster.
2. Navigate to **Settings** > **License**.

   If you have not configured the license, the **License** page displays the Configure License link and the number of days within which you must configure the license. The license must be configured within 180 days of creating the cluster.
3. To configure the license, click **Configure License**.
4. In the floating menu that appears, select one of the following modes and click **Apply**:
   1. Helios Smart Licensing
   2. Manual Licensing
   3. Dark Site

      Follow the procedures below to configure a license mode.

Configure Helios Smart Licensing Mode

Perform the following steps to configure **Helios Smart Licensing**:

1. Follow the steps in the Choose a Licensing Mode after Cluster Creation section and select **Helios Smart Licensing** as the Licensing Mode.
2. On the **Management Options** page that appears, select the **Connect to Helios only for licensing** option.
3. Click **Continue**.

   You are redirected to the Helios page to enter your credentials.
4. Click **Connect** to establish the connection with Helios.

   Only the licensing information is shared with Helios.

Configure Manual Licensing Mode

Configuring the Manual Licensing Mode requires you to download the License Usage Report from the Cohesity cluster and upload it to Helios. You must also generate the License File in Helios and upload it to the Cohesity cluster. As these steps require switching between the Cohesity cluster and Helios, Cohesity recommends that you remain logged in to both the Cohesity cluster and Helios until you complete the Manual License configuration.

Perform the following steps to configure Manual Licensing:

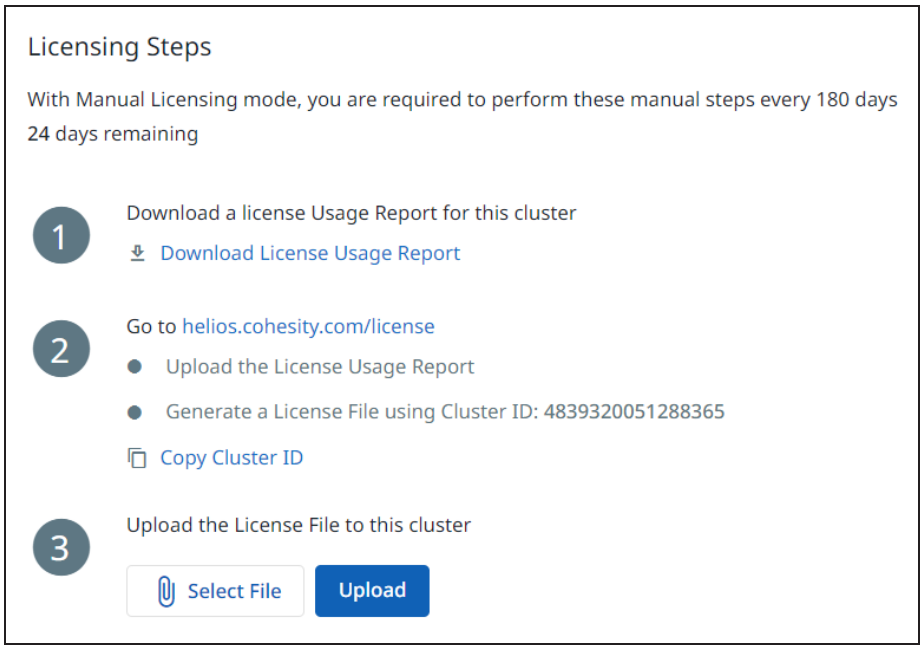

1. Follow the steps in the Choose a Licensing Mode after Cluster Creation section and select Manual Licensing as the Licensing Mode.
2. On the **Licensing Steps** dialog, click **Download License Usage Report**.

   The License Usage Report file or the Audit Report is downloaded to your local machine. The Audit Report contains details such as cluster ID, timestamp, license usage data, and entitlements.
3. To upload the Audit Report:
   1. Log in to Helios and navigate to **Home** > **License Usage**. The license summary of all clusters connected to Helios is displayed.
   2. Click **Upload Audit Report**.
   3. In the **Upload Audit Report** dialog, upload the Audit Report downloaded in step 2.
   4. Click **Submit**.
4. To generate the License File:
   1. In your Cohesity cluster, navigate to **Settings** > **License**. On the **Licensing Steps** page, select **Copy Cluster ID**.
   2. In Helios, click **Generate License File**.

      The **Generate License File** dialog appears and prompts for the Cluster ID.
   3. Paste the Cluster ID copied in Step 4 (a) and click **Generate**.

      The License File is downloaded to your local machine.
5. To upload the License File:

1. In your Cohesity cluster, navigate to **Settings** > **License**.
2. On the **Licensing Steps** page, click **Select File**.
3. Browse to select the License File downloaded in Step 4 (c) and click **Upload**.

Once Manual Licensing is configured, you must download the audit report from the Cohesity cluster and upload it to Helios every 180 days.

Configure Dark Site Mode

To set up a cluster as a dark site, you must contact your Cohesity Representative.

To set up the dark site:

1. Follow the steps in the Choose a Licensing Mode after Cluster Creation section and select Dark Site as the Licensing Mode.
2. To copy the license key from Helios:
    1. Log in to Helios and navigate to **Home** > **License Usage**. The alphanumeric license key is displayed in the License page.
    2. Copy the License key.
3. Log in to your Cohesity cluster.
4. Navigate to **Settings** > **License**.
5. On the **Licensing Steps** page, paste the license key copied from Helios, and click **Upload**.
   The message **License Activated** is displayed. This message indicates that the license has been activated on the cluster.

Troubleshooting Licenses Managed on the Clusters

Cohesity Platform Smart Licensing and Manual Licensing Issues

The following are some of the notifications for sites that are not connected to Cohesity Platform:

Notification	Scenario	Steps to Resolve
`Unregistered cluster`	The cluster has not been registered with Cohesity Platform within 30 days of creation.	1. Go to Cohesity Platform and enter your cluster ID. 2. Once the license file is downloaded, upload the license file to your cluster, and click **Submit**. The following message is displayed: `License file uploaded successfully.`

Notification	Scenario	Steps to Resolve
`Cluster pending audit report upload`	The audit report has not been uploaded within 90 days of the last upload.	Generate a new audit report: 1. Download the audit report from your cluster. 2. Log in to Helios and upload the audit report.

Dark Site Issues

If the license activation was unsuccessful for Dark Site, see the points listed below for possible corrective actions:

- Invalid cluster ID—Confirm the cluster ID for which you have generated the license key. Reconfirm if you have applied this cluster ID while generating the license key. Otherwise, generate the license again using the correct cluster ID.
- License key expired—If the license key has expired, you must generate a new license key. The validity of the license key is 30 days from the time of license generation. You can also confirm if the cluster node time is not more than 30 days from the Helios time.
- Invalid license key—If the license key is invalid, reconfirm the correct license key and enter it correctly.

Manage Apps

Apps are purpose-built applications developed by Cohesity or other third party vendors. You can run the apps directly on the Cohesity cluster to leverage data stored on clusters. For example, the Cohesity Insight app lets you search data stored in views on a cluster. Running apps directly on the Cohesity cluster lets you:

- Get business insights from your backup and unstructured data.
- Analyze data in place; no need to copy data from Cohesity cluster to another infrastructure for analysis.
- Achieve business-critical goals around compliance, eDiscovery and security.

After an app is installed on the Cohesity cluster, users who are assigned the Admin role with the Apps privilege can run the app. The Apps privilege can also be assigned to other roles.

Running an app automatically creates an app instance on the cluster. You can have several instances of the app running at the same time. For example, you might run multiple instances that access data in different views.

Prerequisites

- Marketplace Apps use ports 61001 through 63786 on the Cohesity cluster. Be sure to open this port range if you are using a firewall.
- Before you can install and run apps, the cluster must be enabled to support apps management. Navigate to **Marketplace** > **My Apps**. Click on the **My Apps** tab, click **App Settings** located at the top right of the page and toggle on **Enable Apps Management**.
  - **Allow external connectivity from apps**: Toggle on to allow the apps network access beyond the Cohesity cluster Nodes. Not all apps require external connectivity. For example, a backup app that needs to access a database will require external connectivity, but the Cohesity Insight app, which searches Views on the Cluster, does not require external connectivity.
- Enabling apps management requires each node to have a minimum of 12 CPUs and 55GB RAM. If the cluster does not meet the requirements, you may encounter an error about insufficient memory when you attempt to enable apps management.

Note:

- For clusters running on 6.8.1 and above, the supported subnet size is between /12 (maximum size) and /24 (minimum size).
- The private IP ranges supported are 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. The default IP range is 192.168.0.0/16.
  The RFC 6598 range of 100.64.0.0/10 is also supported. For more information, contact Cohesity Support.
- The alerts 'AppsModeDisabledDueToSubnetClash' or 'NodeNotPresnetInKubernetesCluster' are generated if the apps subnet is the same as the cluster subnet. If they are the same, then update the default app subnet to another private subnet to enable/disable apps management from cluster UI.

Consideration

- Only admin users can launch apps. System services such as Runbook, NoSQL and Hadoop Services, CyberScan cannot be launched by the non-admin user.
- Currently, Marketplace apps are not supported on Cloud Edition and Virtual Edition clusters.

> **Note:**
> - Only the CyberScan and Insights apps are supported on Virtual Edition clusters.
> - NoSQL app is supported from 6.8.1 on GCP.

Browse Apps

You can browse the apps in:

- App Marketplace— Scroll down the Marketplace page to view all the available apps. From the **Filter by** drop-down, select a category to view the apps specific to the selected category.
- Helios—Select **All Clusters** in the top left corner, select **More > Apps** and then click the **Browse Apps** tab.

Install and Run Apps

For information about installing and running apps, see the Marketplace Apps documentation.

Networking Workflows

Cohesity provides two methods to configure Cohesity settings: the Cohesity configure_network.sh script and the iris_cli. In general, using the iris_cli is the preferred method if available. The iris_cli, however, does not currently support all networking workflows. For more information, see Cohesity CLI Reference Guide.

For details about various networking deployment configurations, see Cohesity Networking Quick Start Guide. For details about network design choices, see Optimal Network Designs with Cohesity.

> **Note:** Cohesity recommends using the `configure_network.sh` script only for the initial cluster setup. To change the IP address of the node, contact Cohesity Support.

configure_network.sh

The configure _network.sh script is located at:

```
~/bin/network/configure_network.sh
```

Example from running the script:

```
[cohesity@mycluster-node-1 ~]$ ~/bin/network/configure_network.sh
1) Configure IPMI interface                7) Set Bond Mode
2) Configure IP Address on interface       8) Update Interface Gateway
3) Show Interface config                   9) Configure Default Gateway
4) Show IPMI Config                       10) Configure VLAN interface
5) Check Connectivity                     11) Configure Static Route
6) Remove IP Config                       12) Quit
Please enter your choice:
```

configure_network.sh Option	iris_cli Steps
1) Configure IPMI Interface	N/A
2) Configure IP Address on Interface	Configure Interfaces
3) Show Interface Config	Configure Interfaces
4) Show IPMI Config	Manage IPMI Configuration
5) Check Connectivity	N/A
6) Remove IP Config	Configure Interfaces
7) Set Bond Mode	Bonding Mode
8) Update Interface Gateway	Configure Interfaces
9) Configure Default Gateway	Configure Interfaces
10) Configure VLAN Interface	Configure a Non-native VLAN for Cohesity cluster Creation and Expansion
11) Configure Static Route	Add a Static Route

Related Topic

Cohesity Networking Quick Start Guide.

Bonding Mode

You can edit the bonding mode of the Cohesity cluster bond interfaces. The Cohesity cluster supports these bonding modes:

- **Bonding mode 1 (active-backup):** Places one of the interfaces into a backup state. If the active interface fails, the interface in the backup state becomes active.

This mode provides fault tolerance. By default, bonding mode 1 is configured on the Cohesity cluster bond interfaces.

- **Bonding mode 4 (LACP 802.3ad):** This mode is also known as Dynamic Link Aggregation mode. It creates aggregation groups with the same speed and duplex settings.

The physical and logical networking components in Cohesity cluster are illustrated in this image.

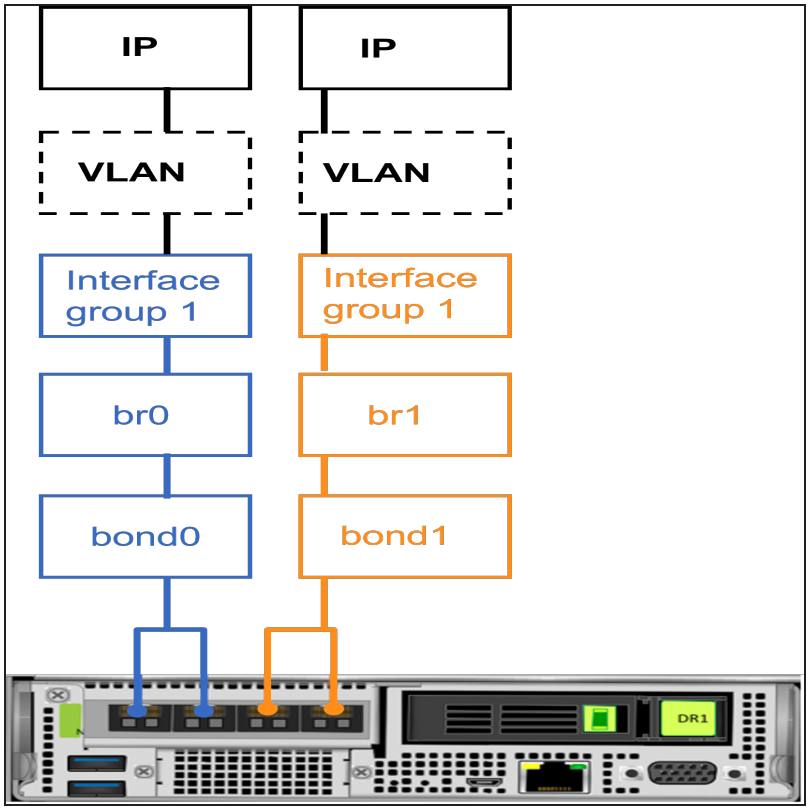

Edit the Bonding Mode

This section provides you steps to change the Cohesity cluster configuration and local ifcfg files to set the bonding mode to LACP.

On the Free Node

Perform these steps on the free node:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run this command:
   - To update the bonding mode from 1 to 4:

     ```
     iris_cli interface update interface-name=bond0 bonding-mode=4
     xmit-hash-policy=1
     ```

   - To update the bonding mode from 4 to 1:

     ```
     iris_cli interface update interface-name=bond0 bonding-mode=1
     ```

3. To verify the node configuration and connectivity, perform these steps:
   - Run this command and verify the output if bonding mode is configured on the bond interface:

     ```
     cat /proc/net/bonding/bond<x>
     ```

     Where x is the bond interface.

     Example:

     ```
     cat /proc/net/bonding/bond0
     ```

   - Ping the gateway IP as follows:

     ```
     ping <gateway IP address>
     ```

On the Cohesity Cluster

Perform these steps on one of the nodes in the Cohesity cluster:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Stop the cluster services running on the nodes before you edit the bonding mode:

   ```
   admin> cluster stop
   ```

3. Check the cluster status:

   ```
   admin> cluster status
   ```

   Verify the cluster status in the command output:

   ```
   SERVICE STATE SYNC: DONE
   CLUSTER ACTIVE OPERATION:
   CLUSTER HEAL STATUS: NORMAL
   ```

4. Set the existing bonding mode:

```
iris_cli interface update interface-name=<interface_group_name>
bonding-mode=<x> xmit-hash-policy=1
```

Example:

```
iris_cli interface update interface-name=intf_group1 bonding-mode=4
xmit-hash-policy=1
```

or

```
iris_cli interface update interface-name=intf_group1 bonding-mode=1
```

5. Change the upstream switch ports to LACP or move the network cable connections to LACP configured ports and wait for LACP on the switch and bond0 interface to come up.
6. To verify the node configuration and connectivity, perform these steps:
   - Run this command and verify the output if bonding mode is configured on the bond interface:

     ```
     cat /proc/net/bonding/bond<x>
     ```

     Where x is the bond interface.

     Example:

     ```
     cat /proc/net/bonding/bond0
     ```

   - Ping the gateway IP as follows:

     ```
     ping <gateway IP address>
     ```

7. Start the cluster services:

   ```
   admin> cluster start
   ```

8. Check the cluster status:

   ```
   admin> cluster status
   ```

   Verify the cluster status in the command output:

   ```
   SERVICE STATE SYNC: DONE
   CLUSTER ACTIVE OPERATION:
   CLUSTER HEAL STATUS: NORMAL
   ```

9. Verify the updated bonding mode:

   ```
   admin> interface list
   ```

Create Cross-Bonding Interface

Create Cross-Bonding for a Primary Interface

Perform these steps on one of the nodes in the cluster:

1. Connect the node in the Cohesity cluster to the IPMI interface via the IPMI Console. For details on how to configure IPMI, see Manage IPMI Configuration.
2. Run the following command in the IPMI console to get a list of all the connected NIC ports and check the 10GbE link.

```
list_all_nic_ports.sh
NODE ID : 181140266786854
NODE IP : 10.1.4.16
```

Sample Output

```
PORT NAME STATE MAC
I/O Module Slot Port 1 ens802f0 UP a4:bf:01:41:a0:26
I/O Module Slot Port 2 ens802f1 UP a4:bf:01:41:a0:27

Port Type: 1GbE
=====================================================
PORT NAME STATE MAC
Embedded 1G NIC Port 1 eno1 DOWN a4:bf:01:21:7e:c8
Embedded 1G NIC Port 2 eno2 DOWN a4:bf:01:21:7e:c9
=====================================================
```

3. Run the following command to stop a single node on a Cohesity cluster:

```
node_stop.sh script
```

The node_stop.sh script is located at */home/cohesity/bin/*

4. In the **ifcfg-ens1f1** file, replace **MASTER=bond0** with **MASTER=bond1**.

To change the bonding master for ens1f1 from bond0 to bond1, you can use one of the following methods:

a. Using nmcli command:

To change the bond master using nmcli, run the following command:

```
sudo nmcli connection modify ens1f1 master bond1
```

> **Note:** nmcli uses a space between the property name and value. Do not use master=bond1 on the command line.

b. Using ens1f1.nmconnection file:

To change the bond master using ens1f1.nmconnection file:

Replace master=bond0 with master=bond1.

> **Note:** If the master field is not available, check for the controller field and update controller=bond0 to controller=bond1.

The ens1f1.nmconnection file is located at /etc/NetworkManager/system-connections/.

5. In the ifcfg-ens7f1 file, replace **MASTER=bond0** with **MASTER=bond1**.

   To change the bonding master for ens7f1 from bond0 to bond1, you can use one of the following methods:

   a. Using nmcli command:

      To change the bond master using nmcli, run the following command:

      ```
      sudo nmcli connection modify ens7f1 master bond1
      ```

   b. Using ens7f1.nmconnection file:

      To change the bond master using ens7f1.nmconnection file:

      Replace master=bond0 with master=bond1.

      > **Note:** If the master field is not available, check for the controller field and update controller=bond0 to controller=bond1.

      The ens7f1.nmconnection file is located at /etc/NetworkManager/system-connections/.

6. Restart the node.
7. Use the following command to verify that the cross-bonding for the primary interface has been successfully created.

   ```
   List_all_nic_ports.sh
   ```

   Sample Output

   ```
   Port Type: 10GbE
   ================================================================================
   PORT    NAME    STATE   MAC
   PCIe Slot 1 Port 2      ens1f1          UP      48:df:37:de:02:48
   PCIe Slot 7 Port 1      ens7f0          UP      48:df:37:dd:f0:b0
   PCIe Slot 7 Port 2      ens7f1          UP      48:df:37:dd:f0:b8
   PCIe Slot 1 Port 1      ens1f0          UP      48:df:37:de:02:40

   Port Type: 1GbE
   ================================================================================
   PORT    NAME    STATE   MAC
   ```

```
Embedded 1G LOM Port 1          eno1    UP      ec:eb:b8:8d:4b:82
Embedded 1G LOM Port 2          eno2    UP      ec:eb:b8:8d:4b:83
```

```
[support@cross-bonding-mxq330bfx-node-1 ~]$ cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)

Bonding Mode: IEEE 802.3ad Dynamic link aggregation
Transmit Hash Policy: layer3+4 (1)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0

802.3ad info
LACP rate: slow
Min links: 0
Aggregator selection policy (ad_select): stable

Slave Interface: ens1f0
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 48:df:37:de:02:40
Slave queue ID: 0
Aggregator ID: 5
Actor Churn State: none
Partner Churn State: none
Actor Churned Count: 0
Partner Churned Count: 0

Slave Interface: ens7f1
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 48:df:37:dd:f0:b8
Slave queue ID: 0
Aggregator ID: 6
Actor Churn State: churned
Partner Churn State: none
Actor Churned Count: 1
Partner Churned Count: 0
[support@cross-bonding-mxq330bfx-node-1 ~]$ cat /proc/net/bonding/bond1
```

```
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)

Bonding Mode: IEEE 802.3ad Dynamic link aggregation
Transmit Hash Policy: layer3+4 (1)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0

802.3ad info
LACP rate: slow
Min links: 0
Aggregator selection policy (ad_select): stable

Slave Interface: ens1f1
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 48:df:37:de:02:48
Slave queue ID: 0
Aggregator ID: 5
Actor Churn State: none
Partner Churn State: churned
Actor Churned Count: 0
Partner Churned Count: 1

Slave Interface: ens7f0
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 48:df:37:dd:f0:b0
Slave queue ID: 0
Aggregator ID: 6
Actor Churn State: churned
Partner Churn State: churned
Actor Churned Count: 1
Partner Churned Count: 1
[support@cross-bonding-mxq330bfx-node-1 ~]$ iris_cli interface list
NODE ID                    : 80123552218480
CHASSIS NAME               : MXQ0330BFT
NODE SLOT NUMBER           : 1
```

```
INTERFACE NAME            : bond1
BONDING MODE              : 1
BOND SLAVE                : ens7f0, ens1f1
BOND SLAVE SLOT           : PCIe Slot 7, PCIe Slot 1
INTERFACE MAC ADDRESS     : 48:df:37:de:3d:78
INTERFACE MTU             : 1500
DEFAULT ROUTE             : false
INTERFACE ROLE            : secondary
INTERFACE SPEED           : 10Gbit/s
INTERFACE LINK UP         : true
INTERFACE SERVICES        : AvahiDiscoverService

INTERFACE NAME            : bond2
BONDING MODE              : 1
BOND SLAVE                : eno1, eno2
BOND SLAVE SLOT           : Embedded 1G LOM
INTERFACE MAC ADDRESS     : ec:eb:b8:8d:1f:3a
INTERFACE MTU             : 1500
DEFAULT ROUTE             : false
INTERFACE ROLE            : secondary
INTERFACE SPEED           : 1Gbit/s
INTERFACE LINK UP         : true
INTERFACE SERVICES        : AvahiDiscoverService

INTERFACE NAME            : bond0
INTERFACE GROUP           : intf_group1
BONDING MODE              : 4
BOND SLAVE                : ens7f1, ens1f0
BOND SLAVE SLOT           : PCIe Slot 7, PCIe Slot 1
INTERFACE MAC ADDRESS     : 48:df:37:de:3d:70
INTERFACE VIRTUAL IP      : 10.1.4.85
INTERFACE V4 STATIC IP    : 10.1.4.79
INTERFACE V4 SUBNET       : 10.1.0.0/20
INTERFACE V4 GATEWAY      : 10.1.0.1
INTERFACE MTU             : 1500
DEFAULT ROUTE             : true
INTERFACE ROLE            : primary
INTERFACE SPEED           : 10Gbit/s
INTERFACE LINK UP         : true
INTERFACE SERVICES        : ClusterDataService, AvahiDiscoverService

INTERFACE NAME            : lo
INTERFACE V4 STATIC IP    : 127.0.0.1
```

```
INTERFACE V4 SUBNET       : 127.0.0.0/8
INTERFACE MTU             : 1500
DEFAULT ROUTE             : false
INTERFACE ROLE            : secondary
INTERFACE SERVICES        : AvahiDiscoverService
NODE ID                   : 80123552203328
CHASSIS NAME              : MXQ330BFX
NODE SLOT NUMBER          : 1

INTERFACE NAME            : bond1
BONDING MODE              : 1
BOND SLAVE                : ens1f1, ens7f0
BOND SLAVE SLOT           : PCIe Slot 1, PCIe Slot 7
INTERFACE MAC ADDRESS     : 48:df:37:de:02:48
INTERFACE MTU             : 1500
DEFAULT ROUTE             : false
INTERFACE ROLE            : secondary
INTERFACE SPEED           : 20Gbit/s
INTERFACE LINK UP         : true
INTERFACE SERVICES        : AvahiDiscoverService

INTERFACE NAME            : bond2
BONDING MODE              : 1
BOND SLAVE                : eno1, eno2
BOND SLAVE SLOT           : Embedded 1G LOM
INTERFACE MAC ADDRESS     : ec:eb:b8:8d:4b:82
INTERFACE MTU             : 1500
DEFAULT ROUTE             : false
INTERFACE ROLE            : secondary
INTERFACE SPEED           : 1Gbit/s
INTERFACE LINK UP         : true
INTERFACE SERVICES        : AvahiDiscoverService
INTERFACE NAME            : bond0
INTERFACE GROUP           : intf_group1
BONDING MODE              : 4
BOND SLAVE                : ens1f0, ens7f1
BOND SLAVE SLOT           : PCIe Slot 1, PCIe Slot 7
INTERFACE MAC ADDRESS     : 48:df:37:de:02:40
INTERFACE VIRTUAL IP      : 10.1.4.82
INTERFACE V4 STATIC IP    : 10.1.4.81
INTERFACE V4 SUBNET       : 10.1.0.0/20
INTERFACE V4 GATEWAY      : 10.1.0.1
INTERFACE MTU             : 1500
```

```
DEFAULT ROUTE              : true
INTERFACE ROLE             : primary
INTERFACE SPEED            : 10Gbit/s
INTERFACE LINK UP          : true
INTERFACE SERVICES         : ClusterDataService, AvahiDiscoverService

INTERFACE NAME             : lo
INTERFACE V4 STATIC IP     : 127.0.0.1
INTERFACE V4 SUBNET        : 127.0.0.0/8
INTERFACE MTU              : 1500
DEFAULT ROUTE              : false
INTERFACE ROLE             : secondary
INTERFACE SERVICES         : AvahiDiscoverService
```

8. Repeat steps 1 to 7 for all the nodes in the Cohesity cluster.

Create Cross-Bonding for a Secondary Interface

Follow these steps to create Cross-Bonding for a secondary interface. Perform these steps on one of the nodes in the cluster:

1. Run the following command in the IPMI console to get a list of all the connected NIC ports and check the 10GbE link.

   ```
   list_all_nic_ports.sh
   NODE ID : 181140266786854
   NODE IP : 10.1.4.16
   ```

   Sample Output

   ```
   Port Type: 10GbE
   ====================================================
   PORT    NAME    STATE   MAC
   Network Card Slot Port 1        eno1    UP
   Network Card Slot Port 2        eno2    UP
   Network Card Slot Port 3        eno3    DOWN
   Network Card Slot Port 4        eno4    UP

   Port Type: 1GbE
   ====================================================
   PORT    NAME    STATE   MAC
   PCIe Slot 1 Port 1     enp59s0f0       DOWN
   PCIe Slot 1 Port 2     enp59s0f1       DOWN
   ```

2. Run the following command to verify the bond interface and its associated physical interface:

   ```
   ip addr
   ```

In this sample output, eno1 and eno2 are bond slaves of bond0, and eno3 and eno4 are bond slaves of bond1.

Sample Output

```
4: eno1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq
master bond0 state UP group default qlen 1000
link/ether e4:43:4b:2f:f7:10 brd ff:ff:ff:ff:ff:ff
5: eno2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq
master bond0 state UP group default qlen 1000
link/ether e4:43:4b:2f:f7:10 brd ff:ff:ff:ff:ff:ff
6: eno3: <NO-CARRIER,BROADCAST,MULTICAST,SLAVE,UP> mtu 1500 qdisc mq
master bond1 state DOWN group default qlen 1000
link/ether e4:43:4b:2f:f7:12 brd ff:ff:ff:ff:ff:ff
7: eno4: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq
master bond1 state UP group default qlen 1000
link/ether e4:43:4b:2f:f7:12 brd ff:ff:ff:ff:ff:ff
```

3. Remove the bond interface_group.

```
iris_cli interface-group delete group-name=<name of the interface
group>
```

4. Remove the bond to free the interfaces.

```
iris_cli bond delete bond-interface=<name of the bond>
```

5. Place the node in maintenance mode.

```
iris_cli node update id=<ID of the node> maintenance-mode=true
```

6. Use the following CLI command to update the cross-bonding interface for bond0 and bond1.

```
iris_cli bond update bond-interface=<name of the bond interface> bond-
slave-interface=<name of the bond slave interface>
```

Example:

Commands that set eno1 and eno3 as salve-interface bond0 and eno2 and eno4 as slave-interface for bond1:

```
iris_cli bond update bond-interface=bond0 bond-slave-interface=eno1
```

```
iris_cli bond update bond-interface=bond1 bond-slave-interface=eno4
```

```
iris_cli bond update bond-interface=bond0 bond-slave-
interface=eno1,eno3
```

```
iris_cli bond update bond-interface=bond1 bond-slave-
interface=eno2,eno4
```

7. Identify the ports that need to be part of the respective bond interface and create bond interfaces using the following command:

   ```
   iris_cli bond create bond-interface=<name of the bondinterface> bond-slaveinterface=<name of the bond slave interface>
   ```

   Example

   ```
   iris_cli bond create bond-interface=bond0 bond-slaveinterface=eno3,eno4
   ```

   ```
   iris_cli bond create bond-interface=bond1 bond-slaveinterface=eno1,eno2
   ```

Delete Cross-Bonding Interface

To delete the cross-bonding interfaces:

1. Run the following command to delete the interface groups that are associated with the cross-bonding interfaces:

   ```
   iris_cli interface-group delete group-name=<name of the interface group>
   ```

   > **Note:** Deleting an interface group fails if the underlying network bonding interface has static IPs, VIPs or VLANs configured.

2. Run the following commands to delete the cross-bonding interfaces:

   ```
   iris_cli bond delete bond-interface=<name of the bond interface>
   ```

   Example

   ```
   iris_cli bond delete bond-interface=bond1
   ```

3. To verify the delete is successful, check if the ifcfg-bond<x> file is deleted/not present in the */etc/sysconfig/network-scripts/* folder.

   Example: Ensure the ifcfg-bond1 file (where 1 is the bond number) is not in the network-scripts folder.

4. Run the following command to remove the node from the maintenance node:

   ```
   iris_cli node update id=<ID of the node> maintenance-mode=false
   ```

Configure Interfaces

Use the iris_cli to configure an interface. To edit the bonding mode, see Bonding Mode.

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run any of the following commands to perform the desired task.

> **Note:** You can change the IP address of the node on the Cohesity cluster. To change the IP address, contact Cohesity Support.

Configure the IP Address for an Interface in an Existing Cohesity cluster

Run the following command to set the IP address (IPv4 or IPv6) for the respective Secondary Interface.

For IPv4

```
iris_cli ip config interface-name=<interface_name> interface-ips=<ip_
addresses> node-ids=<node_ids> subnet-gateway=<gateway> subnet-mask-bit-
s=<subnet_mask>
```

Example:

```
iris_cli ip config interface-name=intf_group2 interface-ips=10.1.1.1 node-
ids=1234 subnet-gateway=10.1.1.100 subnet-mask-bits=24
```

For IPv6

```
iris_cli ip config interface-name=<interface_name> interface-ips=<ip_
addresses> node-ids=<node_ids> subnet-gateway=<gateway> subnet-mask-bit-
s=<subnet_mask> family=<ip address_family>
```

Example:

```
iris_cli ip config interface-name=intf_group2 interface-ips=2001::100 node-
ids=1234 subnet-gateway=2001::1 subnet-mask-bits=64 family=
```

> **Note:** The preceding commands cannot be used to update the IPv4 or IPv6 addresses of a Primary Interface.

Configure the IP Address for an Interface in a Free Node

Run the following command to set the IP address (IPv4 or IPv6) for the respective Secondary Interface (and not in a node in a Cohesity cluster).

```
iris_cli ip config interface-name=<interface_name> interface-ips=<ip_
addresses> subnet-gateway=<gateway> subnet-mask-bits=<subnet_mask>
```

Example:

```
iris_cli ip config interface-name=bond0 interface-ips=10.1.1.1 subnet-gate-
way=10.1.1.100 subnet-mask-bits=24
```

Update the IP address of a Primary Interface

If you are using Cohesity cluster 6.6 or later, run the configure_network.sh script to change the Primary Interface on the VLAN:

```
~/bin/network/configure_network.sh
```

Sample Output:

```
(1) Configure IPMI interface                 7) Set Bond Mode
2) Configure IP Address on interface        8) Update Interface Gateway
3) Show Interface config                     9) Configure Default Gateway
4) Show IPMI Config                         10) Configure VLAN interface
5) Check Connectivity                       11) Configure Static Route
6) Remove IP Config                         12) Quit
Please enter your choice:
IPv4 selected.
Enter the interface to configure [Default=bond0]:bond0
Choose interface:bond0
Enter the vlan id to configure:100
Creating the vlan interface file ...
Created the vlan interface file successfully
Do you want to set bond0.100 as primary interface? y
Setting the primary interface on bond-.100.
Do you want to configure IP info on this VLAN?: Y
Enter the IP address to configure: 20.1.5.18
Enter the interface Prefix: [Default=32]: 20
Do you want to configure the gateway? y
Enter the gateway to configure: 20.1.0.1
Enter the MTU to configure: 1500
```

```
--------MESSAGE-------------
Following is your configuration
DEVICE=xy0.200
TYPE=ABBridge
DEVICETYPE=ovs
ONBOOT=yes
ZONE-public
IPADDR=20.1.5.18
PREFIX= 20
GATEWAY=20.1.0.1
MTU=1500
Do you want to make config active? y
-----------MESSAGE-----------
Reset interface
-----------MESSAGE----------
Restarting Cohesity service....
```

> **Note:** The configure_network.sh prompts for information from you as it runs. The preceding example depicts what follows when you enter "y"(yes) when the script prompts you to configure an IP address on the VLAN. If you do not want to configure the IP address on the VLAN, enter "n"(no) when prompted by the script.

Update the IP address of a Secondary Interface

Configure the Secondary Interface VLAN IP

Follow these steps to configure the Secondary VLAN IP:

1. Run the following command to create the Secondary VLAN IP using Virtual IP:

   ```
   iris_cli vlan create id=<vlan number> interface-group=<name of the interface group> subnet-ip=<IP address of the subnet> subnet-mask-bits=<subnet mask bits> vlan-ips=<comma separated values of the VLAN IPs> subnet-gateway=<subnet gateway>
   ```

   Example

```
iris_cli vlan create id=200 interface-group=intf_group2 subnet-
ip=200.1.0.0  subnet-mask-bits=20 vlan-
ips=200.1.4.15,200.1.4.16,200.1.4.17 subnet-gateway=200.1.0.1
```

2. Run the following command to verify if the Secondary VLAN IP has been created successfully:

```
cluster linux cmd="ifconfig | grep -A 1 <name of the interface>"
```

Example

```
cluster linux cmd="ifconfig | grep -A 1 br1.200"
```

Sample Output

```
Username: admin
'admin' Password:
NODE ID                         : 181140263920230
NODE IP                         : 10.1.4.15
COMMAND OUTPUT                  :
================================
br1.200: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
inet6 fe80::a6bf:1ff:fe38:3c21  prefixlen 64  scopeid 0x20<link>
--
br1.200:3: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
inet 200.1.4.16  netmask 255.255.240.0  broadcast 200.1.15.255
--
br1.200:4: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
inet 200.1.4.15  netmask 255.255.240.0  broadcast 200.1.15.255
================================
NODE ID                         : 181140264822986
NODE IP                         : 10.1.4.18
COMMAND OUTPUT                  :
================================
br1.200: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
inet6 fe80::a6bf:1ff:fe21:8468  prefixlen 64  scopeid 0x20<link>
```

```
--
br1.200:5: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
inet 200.1.4.17  netmask 255.255.240.0  broadcast 200.1.15.255
================================
NODE ID                        : 181140266786854
NODE IP                        : 10.1.4.16
COMMAND OUTPUT                 :
================================
br1.200: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
inet6 fe80::a6bf:1ff:fe21:7ec8  prefixlen 64  scopeid 0x20<link>
================================
```

Configure the Secondary Interface IP (without VLAN)

Follow these steps to configure the Secondary Interface IP without using the VLAN IP:

1. Run the following command to create the Interface Group for the new bond interface:

   ```
   interface-group create group-name=<name of the interface group>
   interface-name=<name of the interface> node-ids=<comma separated values
   of the node ids>
   ```

   Example:

   ```
   interface-group create group-name=intf_group2 interface-name=bond1
   node-ids=181140264822986,181140266786854,181140263920230
   ```

   ```
   Success: Interface group created successfully.
   ```

2. Run the following command to check if the Interface Group is created in all the nodes of the Cohesity cluster:

   ```
   interface-group ls
   ```

   Sample Output

   ```
   INTERFACE GROUP NAME          : intf_group1
   INTERFACE GROUP ID            : 1
   INTERFACE GROUP TYPE          : Vlan
   NODE ID                    : 181140264822986
   INTERFACE NAME             : bond0
   NODE ID                    : 181140266786854
   ```

```
INTERFACE NAME             : bond0
NODE ID                    : 181140263920230
INTERFACE NAME             : bond0
INTERFACE GROUP MTU           : 1500
INTERFACE GROUP BONDING MODE  : 4
BONDING XMIT HASH POLICY      : 1
BONDING LACP RATE             : 0
INTERFACE GROUP NAME          : intf_group2
INTERFACE GROUP ID            : 2
INTERFACE GROUP TYPE          : Vlan
NODE ID                    : 181140264822986
INTERFACE NAME             : bond1
NODE ID                    : 181140266786854
INTERFACE NAME             : bond1
NODE ID                    : 181140263920230
INTERFACE NAME             : bond1
INTERFACE GROUP MTU           : 1500
INTERFACE GROUP BONDING MODE  : 1
```

3. Run the following command to assign an IP address to the interface:

```
ip config interface-ips=<comma separated IP address of the interface>
node-ids=<comma separated values of the node IDs> subnet-
gateway=<subnet gateway> subnet-mask-bits=<subnet mask bits> interface-
name=<name of the interface>
```

Example

```
ip config interface-ips=20.1.4.15,20.1.4.16,20.1.4.17 node-
ids=181140264822986,181140266786854,181140263920230 subnet-
gateway=20.1.0.1 subnet-mask-bits=20 interface-name=intf_group2
```

Remove an IP Configuration

Run the following command to unconfigure an interface.

For IPv4

```
iris_cli ip delete interface-name=<interface_name> node-ids=<node_ids>
```

Example:

```
iris_cli ip delete interface-name=intf_group2 node-ids=1234
```

For IPv6

```
iris_cli ip delete interface-name=<interface_name> node-ids=<node_ids> fam-
ily=<ip address_family>
```

Example:

```
iris_cli ip delete interface-name=intf_group2 node-ids=1234 family=2
```

Update the Interface Gateway

Run the following command to update an interface gateway.

```
iris_cli ip config interface-name=<interface_name> interface-ips=<ip_
addresse> node-ids=<node_ids> subnet-mask-bits=<subnet_mask> subnet-gate-
way=<gateway>
```

Example:

```
iris_cli ip config interface-name=intf_group2 node-ids=1234 subnet-mask-bits-
s=24 subnet-gateway=10.1.1.100
```

Use the preceding command to change the gateway if the primary interface is not involved. If traffic routes through the primary interface to reach the default gateway, use the following command instead:

```
iris_cli route add dest-network=<subnet>/<mask> interface-group=<intf_group>
next-hop=<gateway_ip_address>
```

Example:

If the subnet is 10.0.0.0/24, the interface group associated with this subnet is intf_group1, and the gateway address is 10.1.2.1:

```
iris_cli route add dest-network=10.0.0.0/24 interface-group=intf_group1
next-hop=10.1.2.1
```

Show the Configuration for Interfaces in a Cohesity cluster

Run the following command to show interface configurations in a Cohesity cluster.

```
iris_cli interface list
```

Configure Interface Group

The Interface Group can be configured only through Cohesity CLI. Run the following commands to configure and manage an interface group that will be mapped to different bonds. Once an interface group is created through CLI, it can also be managed (edit, remove, and so on) using the iris_cli.

Operation	Command

To create an Interface Group in the Cohesity cluster	`iris_cli interface-group add group-name=<interface_group_name>`
	or
	`iris_cli interface-group create group-name=<name of the interface group> interface-name=<name of the interface> node-ids=<comma separated values of the node ids>`

To list all the Interface Groups in the Cohesity cluster	`iris_cli interface-group list`
To delete an existing Interface Group **Note:** Deletion of an interface group fails if the underlying network bonding interface have static IPs, VIPs or VLANs configured	`iris_cli interface-group delete group-name=<interface_group_name>`
To edit an Interface Group	`iris_cli interface-group edit group-name=<interface_group_name>`

Add Secondary Interface Groups to New Nodes in Existing Cohesity cluster

If you add a new node to an existing 7.1.2 or later Cohesity cluster, the secondary interface group is not automatically added for the new node. Run these commands to add the secondary interface group to the new node:

Add new nodes and associate them with the existing secondary interface group

Run this command to add a new node to an existing Cohesity cluster and to also associate the new node with an existing secondary interface group:

Command

```
interface-group update group-name=<name of the existing interface group> interface-name=<name of the existing interface>  add-node-ids=<comma separated values of the new node ids>
```

Example

```
interface-group update group-name=intf_group2 interface-name=bond1 add-node-ids=181140264822986,181140266786854,181140263920230
```

Sample Output

```
Success: Interface group updated successfully.
```

Add a new secondary interface group to a new node

Run this command to add a new secondary interface group to a new node already added to a Cohesity cluster:

Command

```
interface-group create group-name=<name of the new interface group> inter-
face-name=<name of the interface> node-ids=<comma separated values of the
node ids>
```

Example

```
interface-group create group-name=intf_group3 interface-name=bond1 node-ids-
s=181140264822986,181140266786854,181140263920230
```

Sample Output

```
Success: Interface group created successfully.
```

Set Up the Secondary Interface for Replication

This topic explains how to configure the Cohesity cluster's secondary interface for remote replication. You must perform all steps in this topic on both the local and remote Clusters.

1. Verify connectivity to the IPMI interface of all Nodes.

   This is needed for cases where misconfiguration might cause loss of connectivity to the bond0 IP address.

2. Connect the secondary interface on all Cohesity cluster Nodes.

   On each Node, connect the interface from the same port group.

   The following diagram illustrates cabling for Haswell with SFP+ ports. When cabling for replication over the secondary interfaces, it is the same for all hardware types (Haswell with SFP+ ports, Haswell with RJ45 ports and Ivy Bridge).

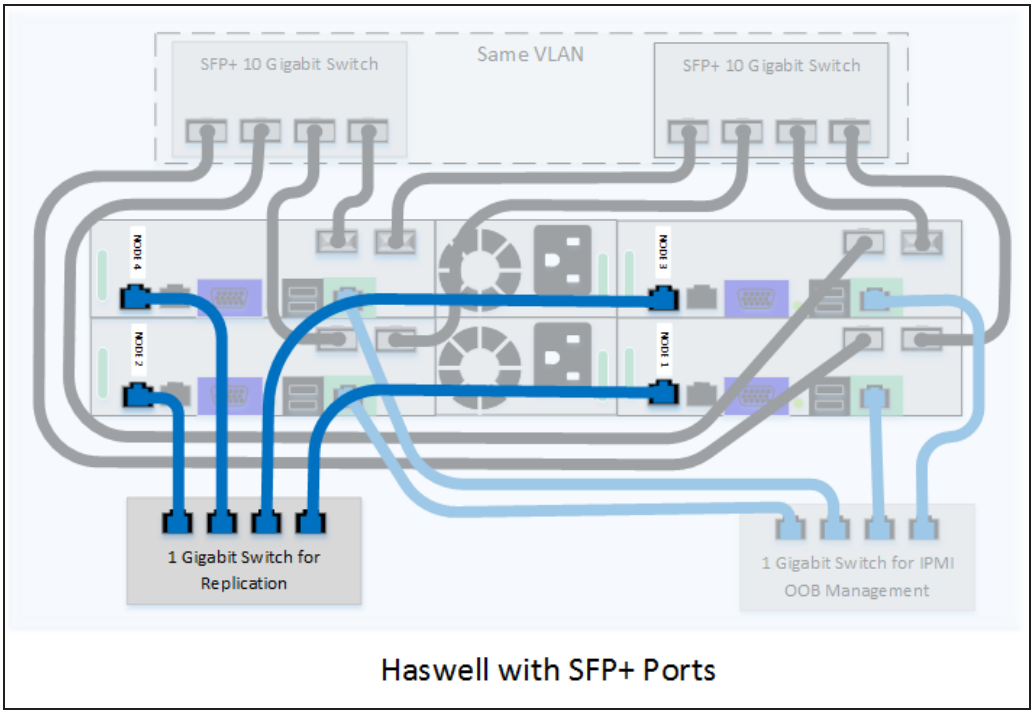

Haswell with SFP+ Ports

3. Verify secondary interfaces are up by using the following command:

```
iris_cli interface list
```

Example:

```
[cohesity@mycluster-node-1 ~]$ iris_cli interface list
CHASSIS NAME                  : BQKP63000000
NODE SLOT                     : 4
    INTERFACE ID              : 130592912804
    INTERFACE NAME            : bond0
    INTERFACE GROUP           : intf_group1
    BONDING MODE              : 4
    BOND SLAVE                : ens802f0, ens802f1
    BOND SLAVE SLOT           : 3, 4
    INTERFACE MAC ADDRESS     : 00:1e:67:f1:b5:a4
    INTERFACE IP              : 10.2.2.15
    INTERFACE SUBNET          : 10.2.32.0/20
    INTERFACE MTU             : 1500
    INTERFACE GATEWAY         : 10.2.32.1
    DEFAULT ROUTE             : true
    INTERFACE ROLE            : primary
    INTERFACE SPEED           : 10Gbit/s
    INTERFACE LINK UP         : true
    INTERFACE SERVICES        : ClusterDataService
```

```
    INTERFACE ID              : 130590912508
    INTERFACE NAME            : bond1
    INTERFACE GROUP           : intf_group2
    BONDING MODE              : 1
    BOND SLAVE                : enp8s0f0, enp8s0f1
    BOND SLAVE SLOT           : 1, 2
    INTERFACE MAC ADDRESS     : 00:1e:67:d3:2f:fc
    INTERFACE IP              :
    INTERFACE SUBNET          :
    INTERFACE MTU             : 1500
    INTERFACE GATEWAY         :
    INTERFACE ROLE            : secondary
    INTERFACE SPEED           : 1Gbit/s
    INTERFACE LINK UP         : true
```

> **Note:** bond1 is typically the interface in a Cohesity 6.x release. You can verify by checking “INTERFACE_SPEED”.

4. Find the Cohesity IDs for all interfaces in the group.

   Use the following command to find all interface id from interface name. Interface id is used as input for replication CLI input.

   ```
   iris_cli cluster status
   ```

5. Configure the VIP, subnet, and gateway using the following command.

   ```
   iris_cli vlan add
   ```

   Configuring mtu is optional. For more information, see Set MTU.

6. Check the configured parameters.

   Use the following command to check the VIP parameters on interface name that were configured.

   ```
   iris_cli vlan ls
   ```

7. Configure remote replication over the interface.

   a. Add a static route

   > **Note:** Adding a static route is required only if the interface for the local and remote Cohesity clusters are not in same network.

A static route enables the local Cohesity cluster and remote Cohesity cluster to reach each other. Use the following command to add a static route on both Clusters:

```
iris_cli route add
```

For more information, see Add a Static Route.

Example for adding a static route:

```
cohesity@mycluster-node-4 ~]$ iris_cli route add desc="inside
cluster route" dest-network=10.2.55.0/22 interface group=intf_
group2 next-hop=10.1.0.1
Success: Static route add successfully.
```

In above example, bond1 is the interface, you can get this information from `iris_cli cluster ls-ifaces`.

```
[cohesity@mycluster-node-4 ~]$ iris_cli route ls
DEST NETWORK : 100.2.32.0/24
NEXT HOP : 10.2.32.1
INTERFACE GROUP NAME : intf_group1
DESCRIPTION : inside cluster route
```

b. Add a remote Cluster pair using the Cohesity UI.

Select **Infrastructure > Remote Clusters > Add Cluster**.

Select **Interface Group** and choose the corresponding interface group.

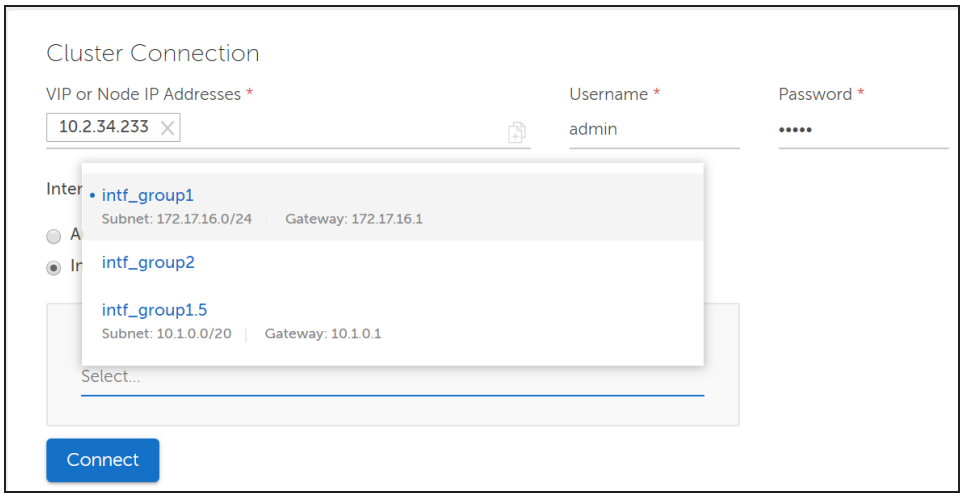

Select **Connect > Replication > Create**.

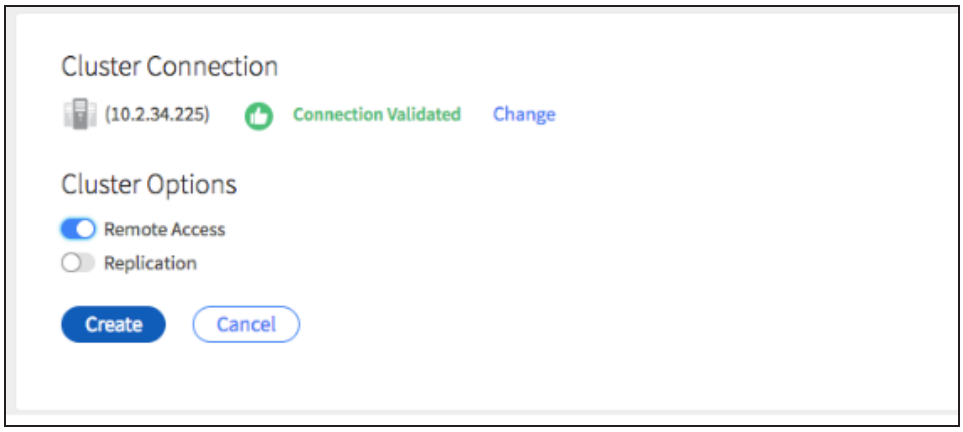

> **Note:** Perform all steps in this topic on both the local and remote Cohesity cluster.

Add a Static Route

This topic describes how to configure a static route on an interface. It also explains how to configure and move the default route from the primary interface to another configured interface on a Cohesity cluster.

Configure a Static Route using Cohesity Dashboard

To configure a static route using the Cohesity Dashboard:

1. In the Cohesity Dashboard, select **Settings > Networking** and then select the **Routes** tab.
2. Click **+Add** and provide the following details:
    a. **Interface Group**: Select one of the pre-configured interface groups from the drop-down menu.
    b. **Destination Network**: Provide the destination network's IP/prefix. For example, 10.0.0.0/24
    c. **Next Hop**: Next hop to the destination network.
    d. **Description (Optional)**: Optional description.
3. Click **Save**.

Configure a Static Route using Cohesity CLI

To configure a static route:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

```
./iris_cli -server 172.16.3.101 -username=admin
```

2. To add a static route, use the following command:

```
iris_cli route add dest-network=<subnet>/<mask> interface-group=<intf_
group> next-hop=<gateway_ip_address>
```

Example:

If the subnet is 10.0.0.0/24, the interface group associated with this subnet is `intf_group1` and the gateway address is 10.1.2.1:

```
iris_cli route add dest-network=10.0.0.0/24 interface-group=intf_group1
next-hop=10.1.2.1
```

If the subnet is 12400:0:0:6000::0/64, the interface group associated with this subnet is `intf_group1` and the gateway address is 2400:0:0:9000::1:

```
iris_cli route add dest-network=2400:0:0:6000::0/64 interface-
group=intf_group1 next-hop=2400:0:0:9000::1
```

3. To verify the configured routes, use the `route list` command.

In the following example, the `route list` command output shows that the primary interface (intf_group1) of a cluster is configured with a static route.

```
iris_cli route list

DEST NETWORK                  : 10.0.0.0/24
NEXT HOP                      : 10.1.2.1
INTERFACE GROUP NAME          : intf_group1
```

> **Note:** Removal of static entry may not work if it overlaps with the primary interface's subnet. For more information, see the Recovering from incorrect static route config KB article.

Configure the Default Route

By default, the primary interface (intf_group1) is configured with the default route. To configure and move the default route to another configured interface, perform the steps provided in this section.

> **Note:** If you configure the default route on a non-primary physical interface (for example, bond1), configure the bond1 interface on the free node in the same subnet as the non-primary interface of the cluster node before you add the free node to the cluster.

To configure the default route:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. By default, the primary interface (intf_group1) is configured with the default route. To configure and move the default route to another configured interface, run the following command:

   ```
   iris_cli route add dest-network=0.0.0.0/0 interface-group=<intf_group>
   next-hop=<gateway_ip_address>
   ```

   Example 1: Moving the default route to a configured physical interface

   If the interface group is intf_group2 and the gateway address is 10.2.32.1:

   ```
   iris_cli route add dest-network=0.0.0.0/0 interface-group=intf_group2
   next-hop=10.2.32.1
   ```

   Example 2: Moving default route to a configured VLAN interface

   If the interface group is intf_group1.100 and the gateway address is 10.1.100.1:

   ```
   iris_cli route add dest-network=0.0.0.0/0 interface-group=intf_
   group1.100 next-hop=10.1.100.1
   ```

3. To verify the configured routes, use the `route list` command.

   In the following example, the `route list` command output shows that the secondary interface (intf_group2) of a cluster is configured with the default route.

   ```
   iris_cli route list

   DEST NETWORK                  : 0.0.0.0/0
   NEXT HOP                      : 10.2.32.1
   INTERFACE GROUP NAME          : intf_group2
   ```

   In the following example, the `route list` command output shows that the VLAN interface (intf_group1.100) of a cluster is configured with the default route.

```
iris_cli route list

DEST NETWORK                  : 0.0.0.0/0
NEXT HOP                      : 10.1.100.1
INTERFACE GROUP NAME          : intf_group1.100
```

Delete a Static Route using the Cohesity Dashboard

1. In the Cohesity Dashboard, select **Settings > Networking** and then select the **Routes** tab.
2. Hover your mouse over the static route you want to delete and then click the delete icon .
3. Click **Delete**.

Delete a Static Route using the Cohesity CLI

To delete the static route:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. To see the list of configured routes, run the `route list` command:

   ```
   iris_cli route list
   ```

   Sample Output

   ```
   DEST NETWORK                  : 10.0.0.0/24
   ```

   ```
   NEXT HOP                      : 10.1.2.1
   ```

   ```
   INTERFACE GROUP NAME          : intf_group1
   ```

   In the preceding example, the output shows that the primary interface (intf_group1) has been configured with a static route (10.0.0.0/24).

3. To delete the static route, run the `route delete` command:

   ```
   iris_cli route delete dest-network=<IP address of the static route>
   interface-group=<name of the interface group>
   ```

   Example IPv4

   ```
   iris_cli route delete dest-network=10.0.0.0/24 interface-group=intf_
   group1
   ```

   Example IPv6

```
iris_cli route delete dest-network=2400:0:0:6000::0/64 interface-
group=intf_group1
```

Manage IPMI Configuration

IPMI configuration for the Cohesity cluster is optional. For ease of managing the nodes remotely, you can either configure IPMI while creating the cluster or after you have created the Cohesity cluster.

The IPMI configuration includes:

- IPMI credentials (username and password).
- IPMI subnet and subnet gateway.
- IPMI address for each node in the cluster.

After you have created the Cohesity cluster, using the Cohesity dashboard, you can:

- Configure IPMI for the Cohesity cluster, if you have not configured IPMI while creating the Cohesity cluster.
- View the IPMI address configured for the nodes in the cluster.
- Edit the IPMI configuration, such as IPMI subnet, subnet gateway, and IPMI address for each node in the cluster.

> **Note:** IPMI communication is required for certain systems to facilitate a proper call-home or dial-home functionality. Please refer to the Setup guides for more information.

Configure IPMI

If you have not configured IPMI while creating the Cohesity cluster, you can set up the IPMI credentials and IPMI addresses at the cluster level using the Cohesity Dashboard.

To configure IPMI:

1. In the Cohesity Dashboard, select **Settings > Networking** and then select the **IPMI** tab.
2. In the **IPMI** tab, click **Setup IPMI**.

   > **Note:** The **Setup IPMI** option is available only if you have not configured IPMI while creating the Cohesity cluster.

3. On the **Setup IPMI** page, in the **User** section, perform the following:
    1. In the **Username** field, enter the IPMI or iRMC username to connect to the IPMI or iRMC interface for each node in the Cohesity cluster.
        - The username length should not exceed 16 characters.
        - Requirements for IPMI credentials may differ according to the hardware platform. For more information, see the Comment column in the Knowledge Base article.
        - The Cohesity cluster supports only one IPMI user.
        - The IPMI user configured on the Cohesity cluster is enabled.

        > **Note:** If the IpmiUserUpdateFailed alert is triggered, check if the node BMC has reached the maximum number of IPMI users. If, yes, delete unwanted users.

    2. In the **New Password** and **Confirm Password** fields, enter and confirm the IPMI or iRMC password to connect to the IPMI or iRMC interface for each node in the Cohesity cluster.
        - The password must not be the same as the username.
        - The password must not be the reverse of the username.
        - The password can be 8 to 16 characters. It must include at least three characters of the following:
            - Upper case
            - Lower case
            - Numbers
            - Special characters: _, -, @, #, ^, &, !, +, ~.
4. In the **Networking** section, perform the following:
    1. In the **Subnet** field, enter the subnet for the IPMI or iRMC network interfaces.
    2. In the **Gateway** field, enter the subnet gateway for the IPMI or iRMC network interfaces.
5. In the **IPMI Address** section, enter the IPMI IP address for each node in the Cohesity cluster.
6. Click **Save**.

View the IPMI Address of the Nodes

Using the Cohesity Dashboard, you can view the IPMI address that you have configured for each node in the Cohesity cluster.

To view the IPMI addresses:

In the Cohesity Dashboard, select **Settings > Networking** and then select the **IPMI** tab. The **IPMI** tab displays the nodes, node IP addresses, and IPMI address configured for the nodes in the cluster.

Edit the IPMI Configuration

You can update the IPMI addresses that you have configured for the nodes in the Cohesity cluster. It provides you options to change:

- The IPMI address configured for the node.
- The IPMI subnet configured for the node, in case if you plan to move the node to a different subnet from the common subnet configured for the Cohesity cluster.
- The subnet gateway IP address.

To edit the IPMI addresses:

1. In the Cohesity Dashboard, select **Settings > Networking** and then select the **IPMI** tab.
2. Click the **Edit** icon next to the node.
3. In the **Edit IPMI Settings** page, update the IPMI address, IPMI subnet, and subnet gateway address for the node as per your requirement.

Manage IPMI Configuration Using the Cohesity CLI

You can update IPMI credentials, list IPMI users, and view IPMI addresses using the iris_cli.

Update IPMI Credentials at the Cluster Level

You can change IPMI credentials at the Cohesity cluster level. For more information on IPMI specific commands, see the ipmi topic in the Using the Cohesity CLI.

To update the IPMI credentials:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi cluster-update-users ipmi-password=<new_password> ipmi-username=admin
   ```

   Sample:

   ```
   cohesity@mycluster-node-1:$ iris_cli ipmi cluster-update-users ipmi-password=Admin@123 ipmi-username=admin
   Success: Updated IPMI credentials.
   ```

Update the IPMI Credentials at the Node Level

If you want a node or set of nodes to use different IPMI credentials than other nodes in the cluster, you can change the IPMI credentials for a node in the cluster. For more information on IPMI specific commands, see the ipmi topic in the Using the Cohesity CLI.

To update the IPMI credentials at node level:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi cluster-update-users node-ipmi-password=<new_password>
   node-ipmi-username=admin node-ips=<comma_separated_IP_addresses_of_the_
   nodes>
   ```

   Sample:

   ```
   cohesity@mycluster-node-1:$ iris_cli ipmi cluster-update-users node-
   ipmi-passwords=Admin@123,Admin@1234 node-ipmi-usernames=admin,admin1
   node-ips=10.1.1.9,10.1.1.11
   Success: Updated IPMI credentials.
   ```

Update IPMI Subnet and Gateway Configuration

You can change IPMI subnet and gateway configuration at the Cohesity cluster level. For more information on IPMI specific commands, see the ipmi topic in the Using the Cohesity CLI.

To update the IPMI subnet and gateway configuration:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi cluster-update-lan-info ipmi-subnet-gateway=<gateway_IP>
   ipmi-subnet-mask=<subnet_IP>
   ```

   Sample:

```
cohesity@mycluster-node-1:$ iris_cli ipmi cluster-update-lan-info ipmi-
subnet-gateway=172.17.16.1 ipmi-subnet-mask=255.255.240.0
```

Update IPMI Subnet and Gateway Configuration and Node IP addresses

Simultaneously, you can change IPMI subnet, gateway configuration, and IP addresses of the nodes in the Cohesity cluster at the cluster or node level. For more information on IPMI specific commands, see the ipmi topic in the Using the Cohesity CLI.

Update IPMI Subnet and Gateway Configuration and Node IP address at Cluster Level

To update the IPMI subnet, gateway, and node IP addresses:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi cluster-update-lan-info ipmi-subnet-gateway=<gateway_IP>
   ipmi-subnet-mask=<subnet_IP> node-ips=<comma_separated_IP_addresses_of_
   the_nodes> node-ipmi-ips=<comma_separated_IPMI_IP_addresses_of_the_
   nodes>
   ```

   > **Note:** IPMI addresses are from the same subnet.

   Sample:

   ```
   cohesity@mycluster-node-1:$ iris_cli ipmi cluster-update-lan-info ipmi-
   subnet-gateway=172.17.16.1 ipmi-subnet-mask=255.255.240.0 node-
   ips=10.1.2.235,10.1.2.236 node-ipmi-ips=172.17.16.200,172.17.16.201
   ```

Update IPMI Subnet and Gateway Configuration and Node IP addresses at Node Level

To update the IPMI subnet, gateway, and node IP addresses:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

```
iris_cli ipmi cluster-update-lan-info node-ipmi-subnet-
gateways=<gateway_IP> node-ipmi-subnet-masks=<subnet_IP> node-
ips=<comma_separated_IP_addresses_of_the_nodes> node-ipmi-ips=<comma_
separated_IPMI_IP_addresses_of_the_nodes>
```

Sample:

```
cohesity@mycluster-node-1:$ iris_cli ipmi cluster-update-lan-info node-
ipmi-subnet-gateways=172.17.16.1,172.17.16.1 node-ipmi-subnet-
masks=255.255.240.0,255.255.240.0 node-ips=10.1.2.235,10.1.2.236 node-
ipmi-ips=172.17.16.200,172.17.16.201
```

View Network Configuration of the Cluster

You view the IPMI subnet, gateway and the IP addresses of the nodes in the Cohesity cluster. For more information on IPMI specific commands, see the ipmi topic in the Using the Cohesity CLI.

To view the network configuration:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi cluster-get-lan-info
   ```

   Sample:

   ```
   cohesity@mycluster-node-1:$ iris_cli ipmi cluster-get-lan-info
   ipmiSubnetMask               : 255.255.240.0
   ipmiGateway                  : 10.2.160.1
   nodesIpmiInfo
   nodeIp                     : 10.2.162.53
   nodeIpmiIp                 : 10.2.162.54
   nodeIp                     : 10.2.162.51
   nodeIpmiIp                 : 10.2.162.52
   nodeIp                     : 10.2.162.49
   nodeIpmiIp                 : 10.2.162.50
   nodeIp                     : 10.2.162.55
   nodeIpmiIp                 : 10.2.162.56
   ```

View Network configuration of the Node

You view the network configuration of the node or a particular node in the Cohesity cluster. For more information on IPMI specific commands, see the ipmi topic in the Using the

Cohesity CLI.

To view the network configuration:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi get-lan-info
   ```

   Sample:

   ```
   cohesity@mycluster-node-1:$ iris_cli ipmi get-lan-info
   ipmiLanInfo
   Channel                     : 3
   IpSrc                       : Static Address
   ip                          : 10.2.162.56
   subnetMask                  : 255.255.240.0
   defaultGatewayIp            : 10.2.160.1
   DefaultGatewayMac
   ```

View the IPMI Users in the Cluster

You can view the IPMI users in the Cohesity cluster. For more information on IPMI specific commands, see the ipmi topic in the Using the Cohesity CLI.

To view IPMI users in the cluster:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi cluster-list-users
   ```

   Sample:

   ```
   cohesity@mycluster-node-1:$ iris_cli ipmi cluster-list-users
   clusterIpmiUser              : admin
   ```

View the IPMI Users in the Node

You can view detailed information about the IPMI users and IPMI channel used in the node or a particular node in the Cohesity cluster. For more information on IPMI specific commands, see the ipmi topic in the Using the Cohesity CLI.

To view IPMI users in the node:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Run the following command on the node:

   ```
   iris_cli ipmi list-users
   ```

   Sample:

   ```
   cohesity@mycluster-node-1:$ iris_cli ipmi list-users
   ipmiUserList
   Channel                    : 3
   MaxUserId                  : 4
   UserList
   Id                       : 2
   name                     : root
   CallIn                   : false
   LinkAuth                 : true
   IpmiMsg                  : true
   Priv                     : 4
   Id                       : 3
   name                     : admin
   CallIn                   : true
   LinkAuth                 : true
   IpmiMsg                  : true
   Priv                     : 4
   Id                       : 4
   name                     : ADMIN
   CallIn                   : true
   LinkAuth                 : false
   IpmiMsg                  : false
   Priv                     : 4
   ```

Reset Networking Configuration Settings

Cohesity supports resetting the network configuration (IP configuration, bonding modes, VLANs, etc) from the Networking Dashboard. This is particularly helpful during the movement of the Cohesity cluster from one network to another network. In other words, when you want to move your data center to a different location or when you are reconfiguring the whole network settings. There will be no effect on the data in the cluster.

When you reset the network configuration settings, ensure the following:

- The Cohesity cluster is not running backup or restore jobs.
- No new nodes are added to the Cohesity cluster.
- No nodes are removed from the Cohesity cluster.

> **Note:** This functionality is not supported on the virtual ROBO, physical ROBO, and cloud clusters.

> **Note:** This functionality does a cluster-wide reset only.

To reset the networking configuration settings:

1. Log in to the Cohesity cluster.
2. Navigate to **Settings > Networking**.
3. Click **Reset**.

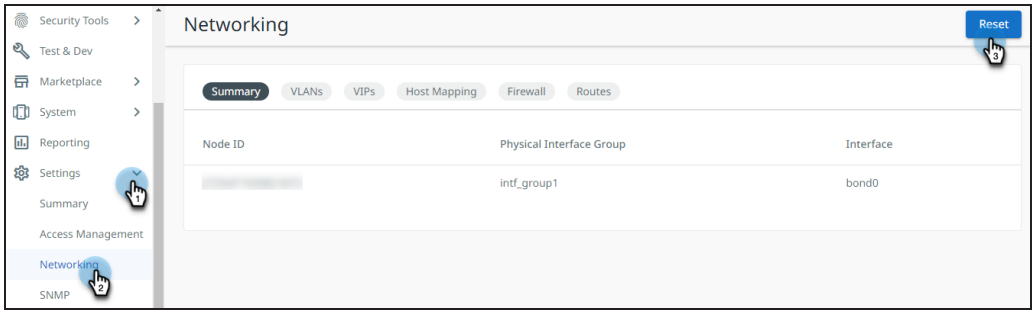

4. On the **Reset Networking** pop-up, type **YES** to confirm and click **Reset**.

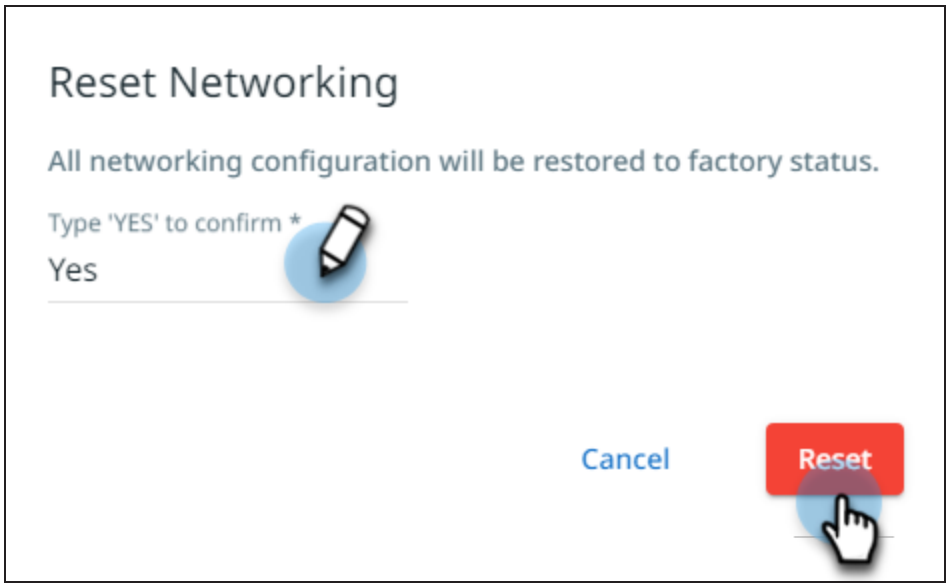

You can see the progress bar in the **Resetting Network Configurations** pop-up.

The progress will automatically be updated. Optionally, you can click the **Refresh**( ) icon to see the real-time progress of resetting the network configurations.

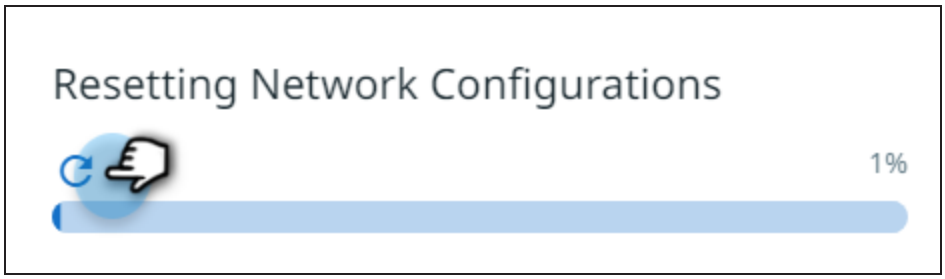

5. After the network configurations reset completes, select one of the following options in the **Reset Completes** pop-up:
   - **Shut Down Cluster**: Select this option to shut down the cluster. You may choose this option when moving the Cohesity cluster to a new location. For more information on powering on a Cohesity cluster, see Power Off and Power On a Cohesity Cluster.
   - **Restart Cluster**: Select this option to reconfigure the network settings on the cluster. For more information on configuring cluster settings, see Configure Settings.
   - **Resume Current Network Configuration**: Select this option to restore your previous network settings.
6. Type **YES** to confirm your selection. Based on your section in step 4, one of the following buttons will be displayed.

- **Shut Down Cluster**: Click this button to shut down the cluster.
- **Restart Cluster**: Click this button to restart the cluster.

  Resume: Click this button to restore the previous network settings.

> **Note:** It is recommended to click the **Resume** button, instead of **Cancel**. If you want to restore your previous network settings. If you click the **Cancel** button after you reset the network settings, this can leave the Cohesity cluster in an unstable state.

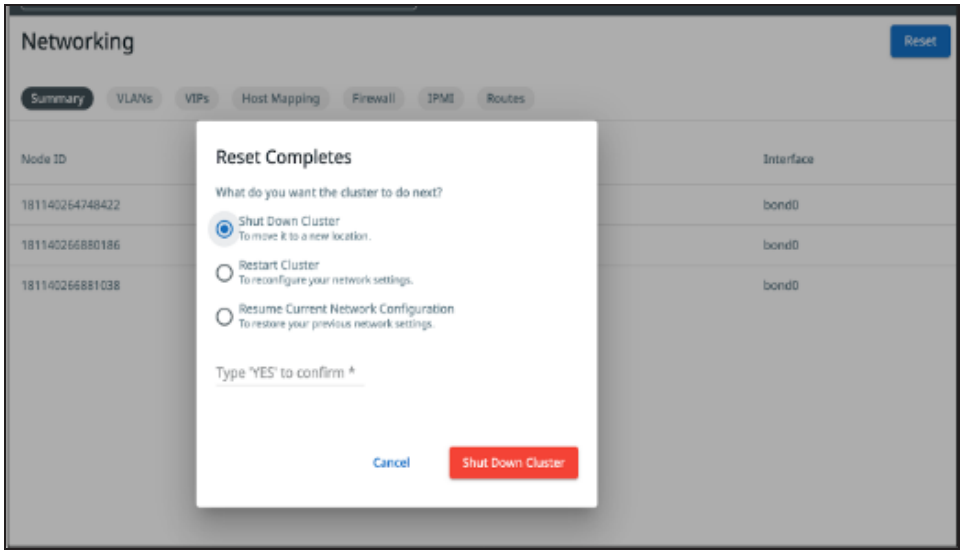

When you select **Shutdown** or **Restart Cluster** in step 5, you need to go through the cluster creation process again and recreate your cluster. You have to re-configure the network configuration on at least one node, update the bonding mode to match the new network, or create a non-native VLAN primary interface. For more information, see Initial Cluster Setup.

Refer to the Setup guide, depending on the hardware model from the Initial Cluster Setup Help.

> **Note:** After you update the node IP addresses of Cohesity Virtual Edition clusters protecting VMware VMs, you must enable the Promiscuous mode in all the ESXi hosts before you rebuild the cluster. For information on how to enable the Promiscuous mode in the vCenter, see the VMware article.

Configure a Non-routable Primary Network for Cluster Creation and Expansion

A few customer environments have a dedicated storage network where the traffic is not routed outside of that network. In such a scenario, you can use the instructions in this topic to deploy Cohesity cluster data traffic on an isolated (non-routable) primary network and use the secondary interface for replication traffic.

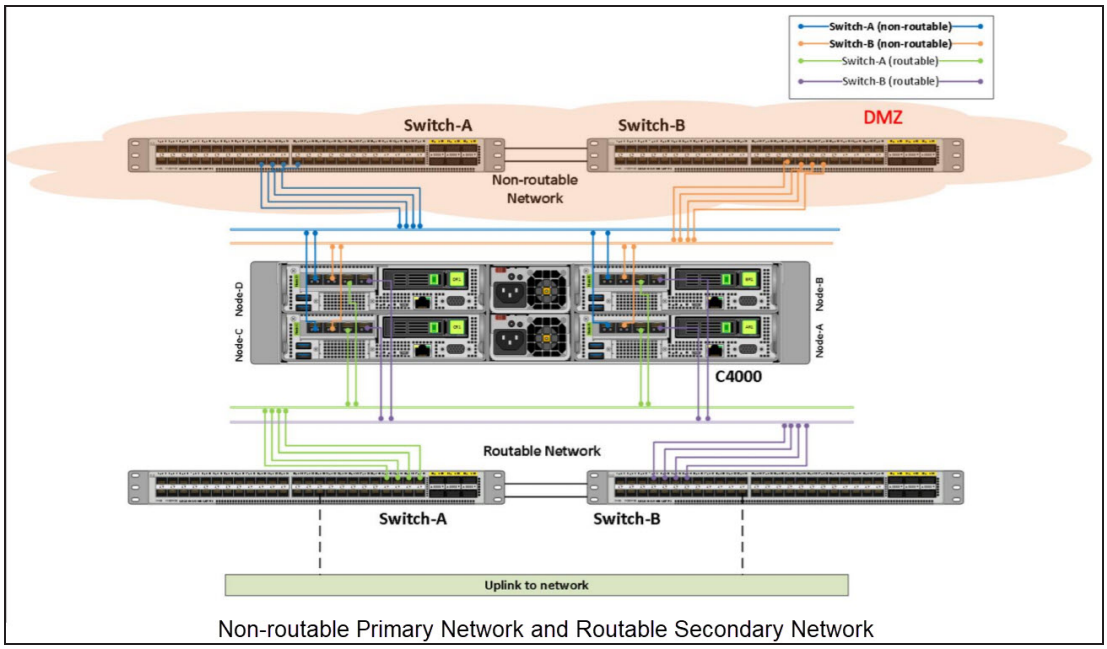

Non-routable Primary Network and Routable Secondary Network

> **Important:** In cases where data traffic is isolated on a primary interface, the gateway is configured on the secondary interface. Ensure that NTP and other ports are reachable via the secondary interface.

Creating a Cohesity cluster

1. Install 6.6 on Cohesity cluster nodes.
2. Launch all cluster nodes' IPMI remote console. (ssh to nodes is not possible.)
3. To create a cluster, you can use one of the following steps on each node:
    a. Use the following command to configure the IP on a routable secondary interface.

    ```
    iris_cli ip config interface-name=<primary_interface_name> iface-ips=xx subnet-gateway=yy subnet-mask-bits=zz mtu=qq(optional)
    ```

    Example:

    ```
    Iris_cli ip config interface-name=bond1 interface-ips=10.1.1.19 subnet-gateway=10.1.0.1 subnet-mask-bits=20
    ```

b. Use the following command to set the secondary interface used as the default route:

```
Iris_cli route add dest-network=0.0.0.0/0 interface-name=xx next-hop=yy
```

Example

```
Iris_cli route add dest-network=0.0.0.0/0 interface-name=bond1 next-hop=10.1.0.1
```

4. Use secondary interface IP or the Cohesity Dashboard and start cluster setup.

5. Assign node IP after the node selection (non-routable primary interface IP).

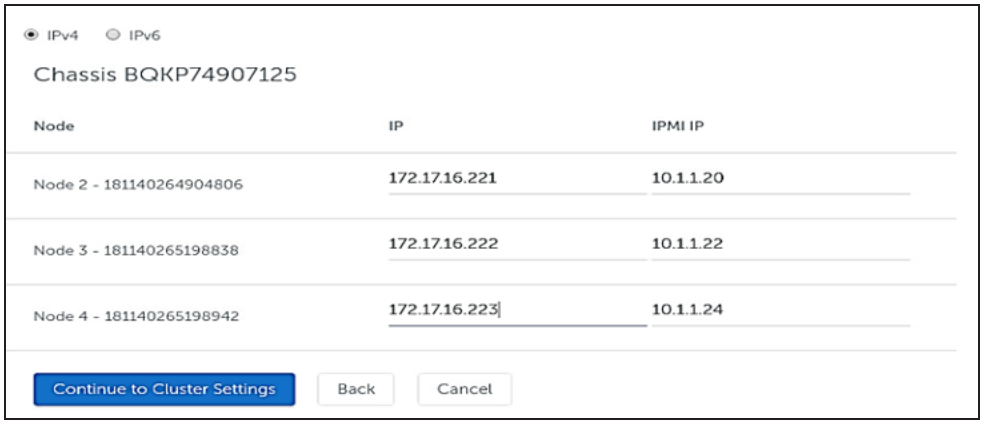

6. Continue with cluster setup and leave the **Cluster Subnet Gateway** field blank as the gateway is not available on the primary node IP network.

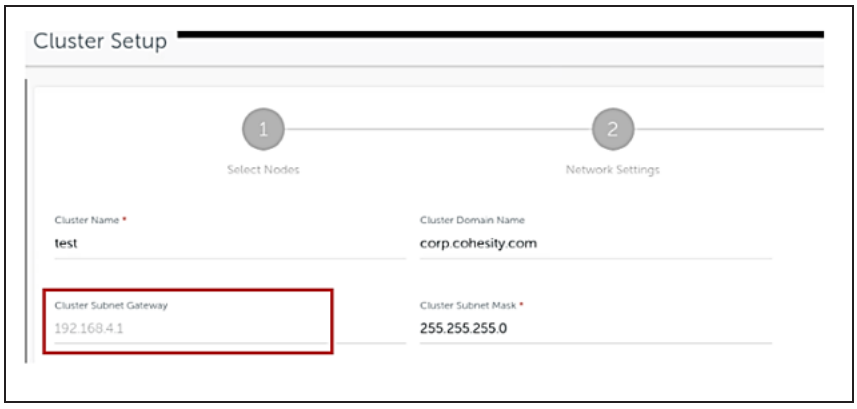

Expanding a Cohesity cluster

1. Install 6.6 on Cohesity free nodes.
2. Launch all free nodes' IPMI remote console. (ssh to nodes is not possible.)
3. You can use one of the following steps on each free node:
    a. Use the following command to configure the IP address on a routable secondary interface.

    ```
    iris_cli ip config interface-name=<primary_interface_name> iface-ips=xx subnet-gateway=yy subnet-mask-bits=zz mtu=qq(optional)
    ```

    Example:

    ```
    Iris_cli ip config interface-name=bond1 interface-ips=10.1.1.19 subnet-gateway=10.1.0.1 subnet-mask-bits=20
    ```

    b. Use the following command to set the secondary interface used as the default route:

    ```
    Iris_cli route add dest-network=0.0.0.0/0 interface-name=xx next-hop=yy
    ```

    Example:

    ```
    Iris_cli route add dest-network=0.0.0.0/0 interface-name=bond1 next-hop=10.1.0.1
    ```

Use the Cohesity Dashboard to add new nodes to the cluster. For more information, see Manage Hardware Nodes.

Set MTU

A maximum transmission unit, also called MTU, defines the largest size of the packet that can be transmitted as a single entity in a network connection. The size of the MTU determines the amount of data that can be transmitted in bytes over a network.

> **Note:** Ensure that the MTU is set with the same value end to end and there is no mismatch.

To enable jumbo frames, ensure the following:

- Switch ports have the MTU set to 9000
- Every network hop between client hosts, such as any cluster nodes, hypervisor hosts, Windows or Linux hosts, and the Cohesity cluster have the MTU set to 9000

After setting the MTU on the path between all hosts and the Cohesity cluster, follow these steps to manually set the MTU size to 9000 on each node of every required network interface:

1. Log in to the *iris_cli* as an Admin user.
2. Stop the cluster services by running the following command:

   ```
   admin> cluster stop
   ```

3. Check the cluster status by running the following command:

   ```
   admin> cluster status
   ```

   > **Note:**
   > Before proceeding, ensure the cluster status is as follows:
   > - SERVICE STATE SYNC: DONE
   > - CLUSTER HEAL STATUS: NORMAL

4. Update the MTU for a node in the Cohesity cluster, using the `interface-name` parameter in the CLI interface update command, where the `interface-name` is the name of the interface group:

   ```
   iris_cli -username=<username> interface update interface-name=<name of
   the interface-group> mtu=<new value of mtu>
   ```

   Example:

   ```
   iris_cli -username=admin interface update interface-name=intf_group1
   mtu=9000
   ```

5. Enter the password.
6. Start the cluster services by running the following command:

   ```
   iris_cli -username=<user> cluster start
   ```

7. Enter the password.

8. Check the cluster status by running the following command:

```
admin> cluster status
```

> **Note:**
> Before proceeding, ensure the cluster status is as follows:
>
> - SERVICE STATE SYNC: DONE
> - CLUSTER HEAL STATUS: NORMAL

9. Run the following command to confirm that the MTU setting is on the physical interfaces and VIPs are set to 9000:

```
netstat -ian
```

The netstat command output should show MTU of 9000 for all bond0* interfaces.

```
Iface      MTU    RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-
OVR Flg
bond0     9000 41096477      0 3037 0      32947757      0      0
0 BMmRU
bond0:2   9000       - no statistics available -
BMmRU
bond1   1500      5189087    37 3036 0     6953          0      0
0 BMmRU
bond0:avahi  9000       - no statistics available -
  BMmRU
eno1      1500     5189087   37 3036 0     6953          0      0
0 BMsRU
eno2      1500     0          0 0    0     0             0      0
0 BMsU
ens802f0  9000     0          0 0    0     0             0      0
0 BMsRU
ens802f1  9000     41096486   0 3037 0     32947757      0      0
0 BMsRU
lo       65536     8959595    0 0    0     8959595       0      0
0 LRU
virbr0   1500      0          0 0    0     0             0      0
0 BMU
virbr0-nic 1500    0          0 0    0     0             0      0
0 BM
```

> **Note:** To update the MTU for a free node that is not part of a Cohesity cluster, use the `interface-name` parameter in the CLI interface update command, where `interface-name` is the interface's name and then enter the password:
> `iris_cli -username=<user> interface update interface-name=<name of the interface> mtu=<new-mtu>`
> Example:
> `iris_cli -username=admin interface update interface-name=bond0 mtu=9000`

Dual-Stack Network Configuration

Dual-stack network configuration enables Cohesity clusters to support both IPv4 and IPv6 addresses on the same network interfaces. This setup allows for a smooth transition between IP protocols and ensures compatibility with different network environments.

Process Overview

A quick overview of the process is as follows:

- Dual stack configuration for New clusters
  - Configure IPv6 and IPv4 on Free Nodes
  - Create Cluster
- Dual stack configuration for Existing clusters
  - Configure dual-stack networking with changing IP family preference
  - Configure dual-stack networking without changing IP family preference
- Create IPv6 and IPv4 subnet for Marketplace Apps
- Enable Apps Management on the Cohesity Cluster
- Install NoSQL and Hadoop Service

Dual stack configuration for New clusters

This topic provides detailed instructions for configuring dual-stack networking on new Cohesity clusters. It includes guidance on setting up IPv4 and IPv6 addresses and managing network settings.

Configure IPv6 and IPv4 on Free Nodes

You need to configure the IPv6 and IPv4 command on the primary interface of the free nodes using `iris_cli` command.

> **Note:** You must configure IPv6 address first and then followed by IPv4.

Run the following command to configure IPv6:

```
iris_cli ip config interface-name=<primary_interface_name> family=2 inter-
face-ips=<IPv6_address> subnet-mask-bits=<subnet_mask> subnet-gate-
way=<gateway>
```

Example

```
admin@127.0.0.1> ip config interface-name=bond0 family=2 interface-
ips=fc00:10:95:160::18 subnet-mask-bits=64 subnet-gateway=fc00:10:95:160::1
Success: Ip configured successfully
```

Run the following command to configure IPv4:

> **Note:** You must configure a private IPv4 address for internal communication within Cohesity cluster nodes.

```
iris_cli ip config interface-name=<primary_interface_name> family=1 inter-
face-ips=<IPv4_address> subnet-mask-bits=<subnet_mask>
```

Example

```
admin@127.0.0.1> ip config interface-name=bond0 family=1 interface-ips-
s=10.95.160.18 subnet-mask-bits=20
Success: Ip configured successfully.
```

Create Cluster

You can create a Cohesity cluster using either an IPv4 or IPv6 address. The IP preference setting determines which IP address is used for inter-node communication.

To complete the initial creation of the Cohesity cluster:

1. Enter the initial Cohesity cluster settings:

> **Note:**
> You can access the IPv6 cluster only from the IPv6 enabled host.
>
> **Syntax**: https://[IPv6]
>
> **Example**: https://[2400::5000:0:0:100:1001]/

Components	Description
Cluster Name	Specify a unique name for the Cohesity cluster. Only alphanumeric characters and hyphens are allowed. A hyphen cannot be the first or last character. Length cannot exceed 32 characters. No other characters are allowed.
Cluster Domain Name	The fully qualified domain name for the Cohesity cluster.
Cluster Subnet Gateway	Specify the IPv6 address of the Subnet Gateway for the Cohesity cluster.
Cluster Subnet Mask	Specify the Subnet Mask for the Subnet that the Cohesity cluster is a part of.
IPMI Subnet Gateway (optional)	Specify the IPv4 address of the subnet Gateway for the IPMI or iRMC network interfaces **Note:** You can either specify the IPMI configuration while creating the Cohesity cluster or after creating the cluster.
IPMI Subnet Mask (optional)	Specify the IPv4 Subnet Mask for the IPMI or iRMC Subnet.
IPMI Username (optional)	Specify the IPMI username to connect to the IPMI interface for each of the nodes in the Cohesity cluster. The cluster uses the IPMI username to get system health information about the nodes in the cluster. All nodes in the cluster must use the same IPMI username and IPMI password. Only alphanumeric characters and hyphens are allowed, but a hyphen cannot be the first character. The length cannot exceed 32 characters. The default IPMI username is **admin**. The Cohesity cluster does not depend on the IPMI configuration to get system health information about the nodes. For ease of managing the nodes remotely, you can configure IPMI.

Components	Description
IPMI Password (optional)	Specify the IPMI password to connect to the IPMI interface for each node in the Cohesity cluster. All nodes in the cluster must use the same IPMI username and IPMI password. The password can be 8 to 16 characters. It cannot include the following characters: dollar sign ($), asterisk (*), quote ("), single quote (`) or backslash (\). The default IPMI password is **admin**. Cohesity recommends changing the default password after creating the cluster. The IPMI password can be set for each node in the cluster, and the password can be different from the IPMI password configured for the cluster. The password can be 8 to 16 characters. It can include at least the following three characters: • Upper case • Lower case • Numbers • Special character, such as: _, -, @, #, ^, &, !, +, ~.
Search Domains	Specify a domain search list for hostname lookup.
DNS Servers	The IPv6 addresses of the Domain Name System (DNS) servers that the Cohesity cluster should use. Separate multiple IPs with commas. Ensure that the Active Directory DNS IP address (if applicable) is listed first. Verify that the NTP servers and other entities in the system can be resolved by the specified DNS server.

Components	Description
NTP Servers	Cohesity recommends using the external Google Public Network Time Protocol (NTP) server and specifying multiple servers (time1.google.com, time2.google.com, time3.google.com, time4.google.com). Avoid using the pool.ntp.org or time.nist.org NTP servers, as they are sometimes unavailable and their IP addresses tend to change. If using an internal NTP server, use only one server (and no external servers). Specify the IP address or the Fully Qualified Domain Name of the NTP server(s). The Cohesity cluster uses the specified NTP server to synchronize the time on all nodes in the cluster. **Note:** To use a Windows NTP server, follow the instructions in this Knowledge Base article. Add IPv6 NTP server if IPv6 DNS server is not added. Also, toggle on **Use Authentication Key** to secure the communication between the NTP server and the Cohesity cluster. In the **Key ID** field, enter the Key id associated with the SHA-1 key and in the **Key** field, enter the SHA-1 key **Note:** Only SHA-1 Keys are supported.
Configure Apps management network	Specify the private IPv4 or IPv6 address for the app subnets. The default IP 192.168.0.0/16 is used for app subnets. If the default IP 192.168.0.0/16 is allocated to a node network, cluster network, or any other network, provide any other private network IP range. Only a private IP range with a minimum subnet size /24 and a maximum subnet size /12 is supported.
Network Settings	Choose either IPv4 or IPv6 under "IP Assignment for Selected Nodes" to set the preferred IP address type for the cluster.

2. If desired, toggle **Encryption** on. Turning on encryption for a Cohesity cluster encrypts all data that will be stored on the cluster.

> **Note:** To encrypt an entire Cohesity cluster, you must specify the encryption option while creating the Cohesity cluster. After a cluster has been created, cluster encryption is not editable. If encryption is not enabled for a Cohesity cluster, you can enable encryption at the Storage Domain level.

The Cohesity cluster supports AES256 software encryption. The Cohesity cluster provides a built-in Key Management Service (KMS) that automatically generates keys.

After Encryption is turned on, the following option is available:

- The Rotation Period is how often the Cohesity cluster's encryption key is rotated. After the time period is reached, the old encryption key is replaced by a new key and the data on the Cohesity cluster remains as it was originally encrypted. The Rotation Period default value is 30 days. You can change this to the desired value.

  > **Note:** Use of Federal Information Processing Standard (FIPS) 140-2 certified modules is enabled by default for data at rest and cannot be disabled.

3. A partition and default Storage Domain are created automatically. Cohesity adds all currently selected nodes to the partition.
   - Specify a fully qualified domain name (**FQDN**). For a Cohesity cluster hosted directly on Cohesity hardware, specify an FQDN that DNS round robin resolves to the specified VIPs.

     > **Note:** You must use an IPv6 address if you are configuring a VIP.

     In situations where you have not yet added the FQDN with VIPs to DNS, enter the FQDN but do not add the VIPs.

     > **Note:** Best practice warrants a DNS entry for the cluster's FQDN and VIPs to achieve optimum cluster performance.

   - Complete **VIP Address** fields. Specify individual virtual IP addresses or ranges of virtual IP addresses for the Cluster. Specifying a VIP range means network traffic to the cluster can be routed to a range of IP addresses instead of a single IP address. For better load balancing, Cohesity recommends specifying the same number of VIPs as there are nodes in the Cohesity cluster. Click **Add VIP or VIP Range**.

     If you do not have VIP addresses yet, leave these fields empty. After the cluster is created, you can add VIP addresses in the Cohesity Platform. Select **Settings** > **Networking** and select the **VIPs** tab.

> **Note:** When the preferred network is IPv6, specify individual IPv6 virtual IP addresses or ranges in the VIP Address fields.

4. Click **Create Cluster**.

   The page displays the Cohesity cluster creation progress.

5. Wait several minutes to allow services to restart, and then click the displayed URL and log in to the Cohesity Platform using the default System Admin account called `admin`, and the default password, `admin`.
6. Read the EULA License Agreement and click **Agree**.
7. Select **SaaS** or **On Prem** configurations and click **Connect**.
8. In the **Change Password** dialog box, enter and confirm the new password for the System Admin account. The minimum length of the password must be 8 characters.

   An Overview Dashboard page displays.

   See Node Detection and Cluster Creation Issues if the cluster creation process is stuck or succeeds with warnings. If Cohesity detects an issue, make corrections before continuing.

You now need to configure the Cohesity cluster. For more information, see Configure Settings.

Dual stack configuration for Existing clusters

This topic provides instructions for configuring dual-stack networking on existing Cohesity clusters. It includes guidance on setting up IPv4 and IPv6 addresses and managing network settings.

Configure dual-stack networking with changing IP family preference

Follow these steps to configure dual-stack networking on existing Cohesity clusters when the IP family preference needs to be changed,

> **Note:** To configure dual-stack networking on CE clusters, please refer to Configure dual-stack networking without changing IP family preference, as changing IP family preference and Reset Networking are not supported on CE clusters.

1. Reset the network configuration settings from the Cohesity dashboard. For more information, see Reset Networking Configuration Settings.

> **Note:** Reset networking is applicable only for physical environments. For VE, a new cluster is needed to change the preferred network.

2. Complete the initial cluster creation process. For more information, see Create Cluster.
3. Configure the Cohesity cluster. For more information, see Configure Settings.

> **Important:** Resetting the network configuration settings, completing the initial cluster creation process, and configuring the Cohesity cluster will not result in data loss. These steps are designed to preserve existing data on the cluster.

Configure dual-stack networking without changing IP family preference

Follow these steps to configure dual-stack networking on existing Cohesity clusters without changing the IP family preference,

Run the following command to set the IP address (IPv4 or IPv6) for the respective Secondary Interface.

For IPv4

```
iris_cli ip config interface-name=<interface_name> interface-ips=<ip_
addresses> node-ids=<node_ids> subnet-gateway=<gateway> subnet-mask-bit-
s=<subnet_mask>
```

Example:

```
iris_cli ip config interface-name=intf_group2 interface-ips=10.1.1.1 node-
ids=1234 subnet-gateway=10.1.1.100 subnet-mask-bits=24
```

For IPv6

```
iris_cli ip config interface-name=<interface_name> interface-ips=<ip_
addresses> node-ids=<node_ids> subnet-gateway=<gateway> subnet-mask-bit-
s=<subnet_mask> family=<ip address_family>
```

Example:

```
iris_cli ip config interface-name=intf_group2 interface-ips=2001::100 node-
ids=1234 subnet-gateway=2001::1 subnet-mask-bits=64 family=
```

> **Note:** The preceding commands cannot be used to update the IPv4 or IPv6 addresses of a Primary Interface.

Create IPv6 and IPv4 subnet for Marketplace Apps

You need to create the IPv6 and IPv4 subnet for MarketPlace apps using `iris_cli` command.

Run the following command for IPv6:

> **Note:**
> - Make sure the subnet is unused to avoid routing conflicts.
> - Supported IPv6 subnet mask bits range is 8-118.

```
iris_cli cluster update-subnet subnet-ip="<ipv6_address>" subnet-mask-bit-
s=<subnet_mask> desc="<description>" service-name="<service_name>"
```

Example

```
admin@127.0.0.1> cluster update-subnet subnet-ip="fc00:192:168::0" subnet-
mask-bits=64 desc="testing ipv6" service-name="apps"
SUBNET ID                     : 0
SUBNET IP                     : fc00:192:168::0
SUBNET MASK BITS              : 64
SUBNET MASK                   : 64
SERVICE NAME                  : apps
DESCRIPTION                   : testing ipv6
```

Run the following command for IPv4:

> **Note:**
> - Make sure the subnet is unused to avoid routing conflicts.
> - Supported IPv4 subnet mask bits range is 12-24.

```
iris_cli cluster update-subnet subnet-ip=<ipv4_address> subnet-mask-ip4-
4=<subnet_mask> service-name=<service_name>
```

Example

```
admin@127.0.0.1> cluster update-subnet subnet-ip=192.168.0.0 subnet-mask-ip4-
4=255.255.0.0 service-name=apps
SUBNET ID                     : 0
SUBNET IP                     : 192.168.0.0
SUBNET MASK BITS              : 16
SUBNET MASK                   : 255.255.0.0
SERVICE NAME                  : apps
```

Enable Apps Management on the Cohesity Cluster

Before you can install and run apps, the cluster must be enabled to support app management.

Perform the following steps to enable app management:

1. Log in to the Cohesity cluster with valid credentials.
2. Navigate to **Marketplace** > **My Apps**.
3. Click the **My Apps** tab and then click **App Settings** located at the top right of the page.
4. Turn on the **Enable Apps Management** toggle.

You can also **Enable Apps Management** using the iris_cli command. Run the following command:

```
iris_cli apps update-apps-settings marketplace-apps-mode=kBareMetal
```

Example

```
admin@127.0.0.1> apps update-apps-settings marketplace-apps-mode=kBareMetal
APPS MODE                       : BareMetal
MARKETPLACE APPS MODE           : BareMetal
```

Install NoSQL and Hadoop Service

You need to download the NoSQL and Hadoop Service and install it on the cluster.

Perform the following steps:

1. Obtain the NoSQL and Hadoop Service package from your Cohesity Accounts contact. The app package format is `app1011-700xxxx.pkg` where `700xxxx` is the version number.
2. Download the package to a location from where you are planning to connect to the Cohesity Platform.
3. Log in to the Cohesity Platform with valid credentials.
4. Navigate to **Marketplace** > **Apps**.
5. Under the **My Apps** folder, click **Upload App** and navigate to the location where you have downloaded the app package, and select the app package.
6. Click **Upload and Install**.
7. Once the upload is complete, click **Run App** to launch the app.

Node Detection and Cluster Creation Issues

If you encounter issues during initial Cohesity cluster creation, read the following sections for possible resolutions:

Verify the following:

All of the following are reachable from the Cohesity cluster:

- Default gateway
- DNS server
- NTP server

Use a Proxy Server

The Cohesity cluster supports routing network traffic to a proxy server for:

- getting upgrade packages from Cohesity
- sending proactive monitoring (Helios) diagnostics from the Cluster to Cohesity
- connecting to the Cluster via the Cohesity Support Channel (if a proxy server is configured for the Cluster as described below, the Support Channel will use the proxy server)

This functionality has been qualified with the Squid3 proxy server.

Configure this functionality using the following CLI commands: `cluster update-proxy-server`, `cluster rm-proxy`, `cluster ls-proxies` and `cluster edit-proxies`. To learn more, log in to the Cohesity CLI and enter the name of the command followed by `help` such as: `cluster update-proxy-server help`.

To configure a proxy for outbound Web traffic using Squid, see the How to configure a proxy for outbound web traffic using Squid KB article.

Update the Proxy Server for the Cluster

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Use the `cluster update-proxy-server` command, for example:

   ```
   admin>cluster update-proxy-server address=10.10.10.10 name=myproxy port-
   t=443
   ```

The command options are shown below:

CLI Option	Description
`address=`	Required. The hostname or IP address of the proxy server.
`name=`	Required. The name of the proxy server.

CLI Option	Description
`port=`	Required. The port to use for the proxy server.
`password=`	Optional. The password to use for the proxy server.
`schemes=`	Optional. A comma separated list of schemes supported by the proxy server. Choose either http or https.
`username=`	Optional. The username to use for the proxy server.

To configure Web Proxy from the Cohesity cluster, see Web Proxy settings.

Verify NTP Setup and Detect Issues

Starting with Cohesity cluster version 7.2.2, Cohesity will transition from Red Hat Enterprise Linux 7.9 OS to Red Hat Enterprise Linux 9.2 as the Host Operating System. Chronyd is now the default time daemon, replacing ntpd.

When a Cohesity cluster is first created, it must have connectivity to an NTP server. Common public NTP servers that can be used:

- time.nist.gov
- time.google.com

Verify the Cluster Can Query the NTP Server

Perform the following from the Cohesity node that is used to create the Cohesity cluster.

`ssh` to a Cohesity Node with the support user and run the `chronyd -Q` command:

```
#chronyd -Q 'server <hostname_of_ntp_server> iburst'
```

For example:

```
#chronyd -Q 'server time.nist.gov iburst'
```

If the NTP server is reachable, expect the following output:

```
[root@f7e746e6 ~]# chronyd -Q 'server time.nist.gov iburst'
2025-08-21T03:26:23Z chronyd version 4.5 starting (+CMDMON +NTP +REFCLOCK
+RTC +PRIVDROP +SCFILTER +SIGND +ASYNCDNS +NTS +SECHASH +IPV6 +DEBUG)
2025-08-21T03:26:23Z Disabled control of system clock
2025-08-21T03:26:27Z System clock wrong by 0.000159 seconds (ignored)
2025-08-21T03:26:27Z chronyd exiting
[root@f7e746e6 ~]#
```

Detect and Resolve NTP Synchronization Issues

NTP time synchronization is critical to Cohesity clusters. A cluster's opclock, which is a vector clock, relies on strict NTP synchronization. If time services become out of sync, they can shut down all I/O operations if a metadata opclock out-of-order condition occurs.

NTP out of sync condition symptoms include:

- Cluster creation does not complete. An error of "kClusterBringup" may be reported.
- After entering Node and IPMI IP addresses in the initial setup wizard, the screen begins to flicker and ultimately returns to the login prompt.
- Cluster creation fails with "WARNING | CE00411001 - FrequentProcessRestarts | Process bridge_exec is restarting frequently".
- "An internal component may be down" message is reported.
- The Cohesity Bridge service experiences FATAL errors with the following message:

```
F1013 01:21:08.107616 120432 client_base_op.cc:247] Server asked the
client to crash. Timestamp in the request is 10001 msecs in the future.
Please check the time difference between nodes.
logs/bridge_exec.FATAL (END)
```

Change operations may not be allowed within the Cohesity cluster. For example, the time cannot be set earlier because of the opclock dependency.

There are multiple ways to check the NTP server configuration and current state. This following information describes how to check an chronyd configuration and status.

- Verify the time difference on multiple Nodes:

```
[cohesity@9166193e ~]$ allssh.sh "date"
=========== 10.10.10.10 ===========
Mon Oct 12 16:27:48 PDT 2015     |-- NTP master is faster than other
nodes
=========== 10.10.10.11 ===========
Mon Oct 12 16:25:51 PDT 2015
=========== 10.10.10.12 ===========
Mon Oct 12 16:25:50 PDT 2015
=========== 10.10.10.13 ===========
Mon Oct 12 16:25:50 PDT 2015
```

- Verify the chrony.conf on all Nodes:

```
[cohesity@9166193e ~]$ allssh.sh "less /etc/chrony.conf"
=========== 10.10.10.10 ===========
# /etc/chrony.conf, configuration for chronyd

#-----------------------------------------------
#  DO NOT EDIT THIS FILE - it is machine generated
#  Don't change config without validation for security holes
```

```
#-------------------------------------------------

server 10.10.10.107 iburst maxpoll 16 |--- Configured to local NTP
server

# Record the rate at which the system clock gains/losses time.
driftfile /var/lib/chrony/drift

# Server(s)
=========== 10.10.10.11 ===========
# /etc/chrony.conf, configuration for chronyd

#-------------------------------------------------
#  DO NOT EDIT THIS FILE - it is machine generated
#  Don't change config without validation for security holes
#-------------------------------------------------

server 10.10.10.10 iburst maxpoll 16 |--- Configured to local cluster
NTP master

# Record the rate at which the system clock gains/losses time.
driftfile /var/lib/chrony/drift

=========== 10.10.10.12 ===========
# /etc/chrony.conf, configuration for chronyd

#-------------------------------------------------
#  DO NOT EDIT THIS FILE - it is machine generated
#  Don't change config without validation for security holes
#-------------------------------------------------

server 10.10.10.10 iburst maxpoll 16 |--- Configured to local cluster
NTP master

# Record the rate at which the system clock gains/losses time.
driftfile /var/lib/chrony/drift

=========== 10.10.10.13 ===========
#-------------------------------------------------
#  DO NOT EDIT THIS FILE - it is machine generated
#  Don't change config without validation for security holes
#-------------------------------------------------

server 10.10.10.10 iburst maxpoll 16 |--- Configured to local cluster
```

```
NTP master

# Record the rate at which the system clock gains/losses time.
driftfile /var/lib/chrony/drift
```

- Query the time from NTP server for internal and external:

```
[cohesity@9166193e ~]$ ifconfig | grep inet
inet 10.10.10.10  netmask 255.255.255.0  broadcast 10.10.10.255

[cohesity@9166193e ~]$ chronyd -Q 10.10.10.107
2025-08-21T03:26:23Z chronyd version 4.5 starting (+CMDMON +NTP
+REFCLOCK +RTC +PRIVDROP +SCFILTER +SIGND +ASYNCDNS +NTS +SECHASH +IPV6
+DEBUG)
2025-08-21T03:26:23Z Disabled control of system clock
2025-08-21T03:26:27Z System clock wrong by 0.000159 seconds (ignored)
2025-08-21T03:26:27Z chronyd exiting |-- NTP master node is 0.000159
seconds behind the local NTP server, 10.10.10.107
[cohesity@9166193e ~]$ chronyd -Q time.nist.gov
2025-08-21T03:26:23Z chronyd version 4.5 starting (+CMDMON +NTP
+REFCLOCK +RTC +PRIVDROP +SCFILTER +SIGND +ASYNCDNS +NTS +SECHASH +IPV6
+DEBUG)
2025-08-21T03:26:23Z Disabled control of system clock
2025-08-21T03:26:27Z System clock wrong by -0.000459 seconds (ignored)
2025-08-21T03:26:27Z chronyd exiting    |-- NTP master node is 0.000459
seconds faster than public time.nist.gov server
```

- Check detailed NTP status:

```
[cohesity@9166193e ~]$ sudo chronyc sources
MS Name/IP address         Stratum Poll Reach LastRx Last sample
===============================================================================
=========
^* 10.2.1.120                    2   6    77    31    +57us[ +242us]
+/-   21ms

[cohesity@9166193e ~]$ sudo ntpstat

synchronised to NTP server (10.2.1.120) at stratum 3
time correct to within 19 ms
polling server every 128 s
```

Once you identify a time inconsistency, correct it immediately. Cohesity recommends synchronizing the local time server time to a public time server. To do so, run the following command which stops chronyd daemon, resynchronizes to a public time server and starts chronyd daemon:

```
allssh.sh sudo systemctl stop chronyd && sudo chronyd -q time.google.com
```

> **Note:** The Cohesity cluster time cannot be changed to an earlier time. If such a change is needed, the Cluster or a specific Node must be shutdown to allow for synchronization to occur.

Manage VLANs

After the initial setup of a Cohesity cluster, you can extend your existing on-premise VLANs into the Cohesity cluster to logically separate traffic and improve performance. After you add a VLAN to the Cohesity cluster, users can select the VLAN when recovering VMs, files, folders, Pure Storage volumes, NAS volumes, or when cloning a vCenter Source. Some recovery and clone tasks can auto-select the appropriate VLAN for the task, but users can override the selection.

The Cohesity cluster supports a maximum of 4094 VLAN IDs irrespective of the cluster size.

The following diagram shows an example Cohesity environment using VLANs.

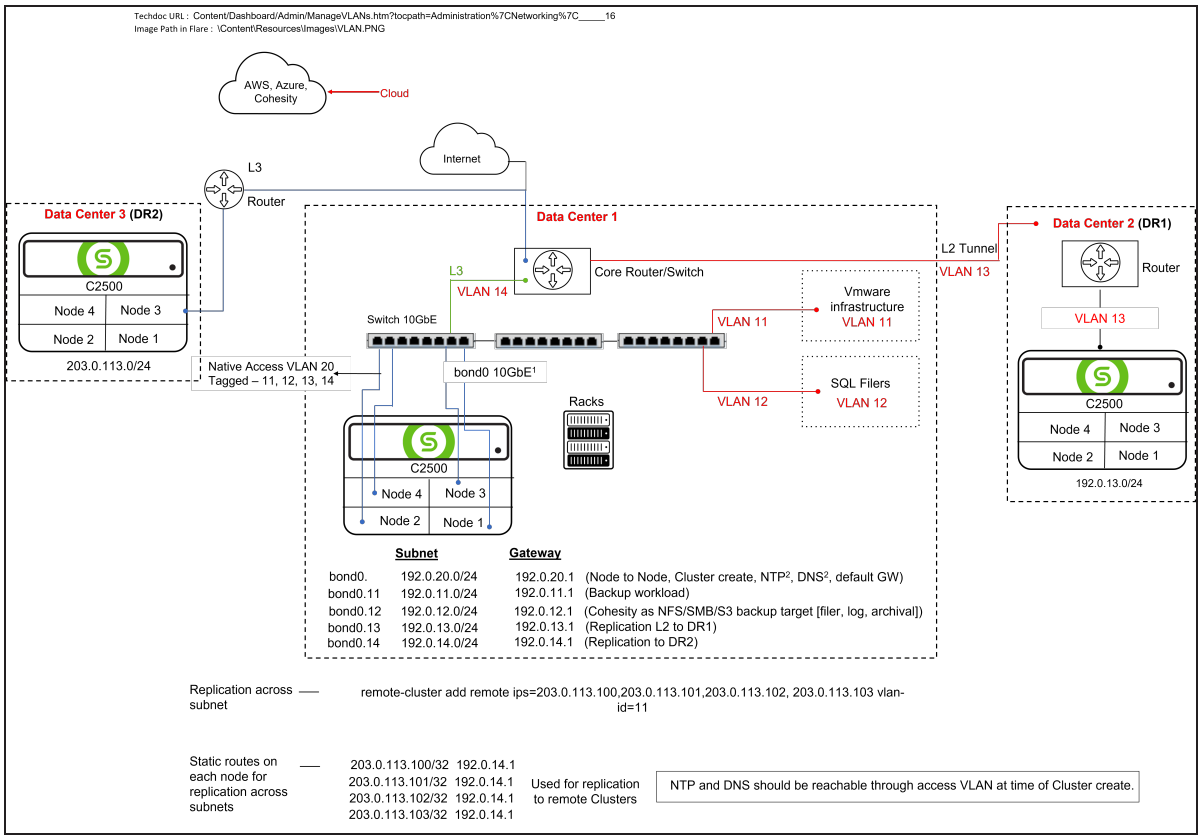

Add a VLAN

Perform the steps in this section to add a VLAN.

> **Note:**
> Cohesity recommends the following:
>
> - Do not configure a VLAN gateway if you want to use global static route for outgoing traffic from the Cohesity cluster.
> - Enable a separate VLAN gateway (different from the default VLAN gateway) to route incoming traffic to the Cohesity cluster via a VLAN IP address such as NFS or UI. Static routes are required only for outgoing traffic, and the Default Gateway setting of the VLAN is only for incoming traffic regardless of Static Routing.

1. In the Cohesity Dashboard, select **Settings > Networking** and then select the **VLANs** tab.
2. Click **Add VLAN**.
3. Provide the following information:
   - **Interface Group:** Select the interface group to use for the VLAN.
   - **VLAN ID:** Specify the ID of the VLAN you want to add.
   - **Subnet:** Specify the IPv4 or IPv6 address of the subnet mask bits for the VLAN.
   - **Gateway:** Optional. Specify the IPv4 or IPv6 address of the subnet gateway for the VLAN.
   - **Enable For All Organizations**:Optional. Toggle this setting on if you want to enable VLAN for all organizations. For more information, see Enabling VLANs for all Organizations.
4. Click **Save**.

The VLAN is added to the Cohesity cluster and available for the supported recovery and clone tasks.

To add optional virtual IP addresses (VIPs) to the VLAN, select **Settings > Networking** and then select the **VIPs** tab. From the Interface Group drop-down, select **intf_group1**.

Edit a VLAN

1. In the Cohesity cluster, select **Settings > Networking** and select the **VLANs** tab.
2. In the list of VLANs, click the actions menu (⋮) for the VLAN and then click the edit icon ( ).
3. Change the settings as needed. Note the following:

- You cannot change the VLAN ID.
- If you are editing the subnet mask of a VLAN, you must also specify a new value for the subnet, as changing the subnet mask changes the subnet and, thereby, the number of hosts and subnets allowed.
- If the number of assigned VIPs is equal to the number of nodes, you must add the new IPs before you can remove the old IPs. This ensures that all the nodes always have a VIP assigned.

4. Click **Save**.

Delete a VLAN from the Cohesity cluster

1. In the Cohesity Dashboard, select **Settings > Networking** and select the **VLANs** tab.
2. In the list of VLANs, click the actions menu (⋮) for the VLAN and then click the delete icon (🗑).
3. Click **Delete**.

Manage VIPs

When a cluster is created, every node in the cluster is assigned a Node IP (or static IP) address. However, if the cluster goes offline, the Node IP addresses become inaccessible and go down as well.

Virtual IP (VIP) is assigned to the cluster nodes after the cluster is created. Even if the cluster goes down, the assigned virtual IP remains available.

Add a VIP

You can add IPv4 or IPv6 virtual IP addresses (VIPs) to an interface group.

To add a VIP:

1. In the Cohesity Dashboard, select **Settings > Networking** and then select the **VIPs** tab.
2. Select the VIP type you want to add, **IPv4** or **IPv6**.
3. **Interface Group:** Select an interface group from the drop-down.
4. Provide the following information:
   - **Subnet:** Specify the IP address of the subnet mask bits for the interface group.
   - **FQDN:** Specify the fully qualified domain name associated with the interface group.
   - **Gateway:** Optional. Specify the IP address of the subnet gateway for the interface group.

- **VIP Address or Range:** Optional. Specify a single virtual IP (VIP) address or a range of addresses assigned to the interface group and then click **Add VIP or Range**. If providing a range, enter a complete address in the first field and just the upper limit of the range (such as 5) in the second field.

  For the non-primary interface group, the number of VIPs must be equal to or greater than the number of Nodes in the Cluster.

  If you provide a number of VIPs equal to or greater than the number of Nodes, the VLAN is automatically added to the Cluster.

  For the primary interface group, the number of VIPs can be lesser than the number of Nodes in the Cluster.

  > **Note:** VIPs cannot be assigned to two interfaces within the same subnet.

5. **Inbound DNS (Optional)**Type Zone name and Zone Server VIPs. Click the (+) icon to add multiple zones.

   > **Note:** Use this option to use the Cohesity cluster as an internal DNS server. For more information, see Internal DNS for Load Balancing.

   Configuring it helps in load balancing by selecting available VIPs for balancing network traffic.

   A Zone name consists of an internal DNS name and a client-side DNS name. For example, Testintdns.storage.eng. In the example, storage.eng is the DNS zone in the client environment and Testintdns is the child zone hosted on Cohesity cluster.

   > **Note:** You should know your DNS server name to add the zone name and can use one of the VIPs as Zone Server VIP. You cannot delete a VIP that is used as a Zone Server VIP.

6. Click **Update**.

   > **Note:** If you have updated any of the above mentioned fields, you must click **Update** before proceeding.

Delete a VIP

You can delete IPv4 or IPv6 virtual IP addresses (VIPs) from an interface group.

> **Important:** Ensure that no backup traffic using the subnet before deleting the VIPs.

To delete a VIP:

1. In the Cohesity Dashboard, select **Settings > Networking** and select the **VIPs** tab.
2. In the list of VIPs, click the delete icon ( ) to delete the VIP.

Internal DNS for Load Balancing

The Domain Name System (DNS) in the Cohesity cluster helps in load balancing by selecting available VIPs for balancing network traffic on VLANs. From release 6.6 onwards, your Cohesity cluster can act as a DNS server. You can configure your cluster as a DNS server by configuring **In Bound DNS** option and creating DNS load balancing zones in the Cohesity Cluster. The configured DNS load balancing zones are valid with a VLAN.

You can configure the **In Bound DNS** when creating a VIP.

Cohesity uses the round-robin technique for managing the client's requests.

You should register one cluster serving VIP on a DNS server (client-side DNS) in your environment. The Cohesity cluster acts as a subdomain for your DNS server. The configured DNS server inside the cluster addresses the issue of updating the client-side DNS when you add or remove a node in the Cohesity cluster or add a new VIP.

Advantages of DNS Load Balancing

The advantages of DNS load balancing are:

- You do not need to add all VIPs on the client-side DNS server for mapping.
- Map a single VIP with the client-side DNS.

Prerequisites

- Access to external DNS server

Considerations

- You can’t delete a VIP configured with an external DNS server.
- A single FQDN can resolve only with one zone.
    - The View mount paths for SMB, NFS, and S3 do not show paths for other DNS zones when you have a single FQDN and. the cluster is not registered with AD.
- In a multitenancy environment, DNS load balancing is supported only for VLAN-based organizations (tenants).

Create DNS Load Balancing Zone

DNS load balancing zone assigns available VIPs to clients using the round-robin method and manages network traffic.

When you configure the DNS load balancing zone on your Cohesity cluster, available VIPs are linked to a load balancing zone in a VLAN. You need to configure the client-side DNS server to delegate all DNS requests to the DNS load balancing zone on the Cohesity cluster.

You can create a DNS load balancing zone when creating a VIP. For more information, see Manage VIPs.

Configuration on Client-Side DNS Server

Create a new delegation zone in the DNS Server for the child domain that you have created on the Cohesity cluster. When you create a delegation, any request to your DNS server is forwarded to the DNS load balancing zone created on your Cohesity cluster.

Add A-record and Delegation Zone on Client-Side Windows DNS Server

You need to add an address mapping record (‘A’ Record for IPv4) or ‘AAAA’ (for IPv6) record in the client-side DNS server and add a delegation.

For example, we have this client-side DNS server zone.com and you added an internal DNS zone on your Cohesity cluster: cohesityV4dns.zone.com with VIP 10.1.1.1.

Now you need to make an 'A' record in the zone.com DNS for the VIP 10.1.1.1 and add a delegation using the 'A' record and the internal DNS zone name: cohesityV4dns.zone.com.

Prerequisites

- Cohesity DNS zone name.
- VIP address used for creating the DNS zone on the cluster.

To add a delegation zone on the client-side Windows DNS Server:

1. Log in to your Windows Server.
2. Navigate to **Windows Administrator Tools > DNS Manager > Forward Lookup Zones**.
3. Select the required domain and add an A or AAAA entry for one of the VIPs from your Cohesity cluster.
4. Enter a name for the Cohesity internal DNS server.
5. After adding the A or AAAA record entry, add a DNS delegation.
6. Right-click the domain and select **New Delegation**.
7. In the **New Delegation** wizard, enter the DNS zone name that you added in the cluster. Enter the A or AAAA record name that you added to the DNS server.
8. Complete the instructions in the wizard to add the delegation.

Once the A or AAAA record is added and delegation is created, any requests to the client-side DNS Server are forwarded to the DNS zone created on the Cohesity internal DNS server.

Add or Remove Inbound DNS Zone

You can add or remove a load balancing zone.

To add or remove DNS zone:

1. In the Cohesity cluster Dashboard, navigate to **Settings > Networking** and then click the **VIPs** tab.
2. In the **Inbound DNS (Optional)** section of the page, click the (+) icon to add multiple zones or click (-) to delete a zone.
3. Click **Update**.

Manage Host Mappings

Use the Cohesity cluster to provide Cluster-wide /etc/hosts entries that map IP addresses to domain names. Using the dashboard ensures that the entries are centrally stored in the Cluster and propagated to all existing Nodes and to any new Nodes added to the Cohesity cluster. The entries are also preserved during Cohesity cluster upgrades.

On a new or newly upgraded Cohesity cluster running release 6.5.1 or later, the Cluster-wide file is initialized automatically. The Cluster then monitors the Nodes at regular intervals for changes to /etc/hosts and propagates the changes to the other Nodes within a few seconds. For an upgraded Cohesity cluster, any existing /etc/hosts entries are retained and uploaded to the Cluster-wide file.

> **Note:** Do not directly edit the /etc/hosts files on the Nodes in the Cohesity cluster.

Manage /etc/hosts Entries

You can manually enter the host names and IP addresses or upload a host file with the mappings between them.

To manually enter the host names:

1. Select **Settings > Networking** and then select the **Host Mapping** tab.
2. Click the Edit icon (✎) next to **Host Mapping**.
3. Enter the IP addresses and host names.
4. Click **Save**.

To upload entries from a host file (in *.txt format):

> **Note:** Once you upload a host file, existing entries are overwritten. If the existing entries are critical, Cohesity recommends adding those entries to the host file.

1. Select **Settings > Networking** and then select the **Host Mapping** tab.
2. Click **Upload hosts file**.
3. Browse to open your text file that has the IP address vs. host name mapping. The IP address and hostnames must be separated by a space or tab, as shown below. Comma-separated values are not supported.

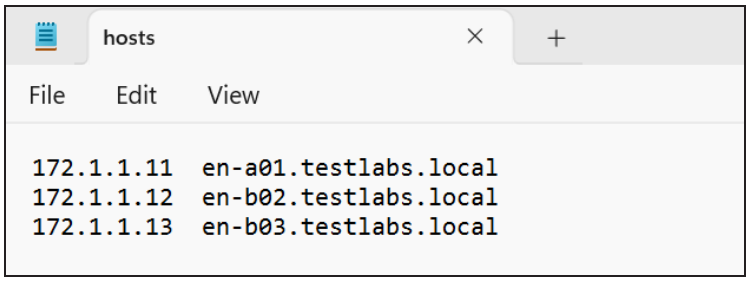

4. Click **Save**.

   Once you click **Save**, the contents of your text file are uploaded and displayed on the same page in the Cohesity cluster UI.

Configure a Non-native VLAN for Cohesity cluster Creation and Expansion

Create a Cohesity cluster

1. Install 7.4.1 on Cohesity cluster nodes.
2. Launch all Cohesity cluster nodes' IPMI remote console. (ssh to nodes is not possible.)
3. To create a Cohesity cluster, you can use one of the following methods:
   1. Through iris_cli. Perform these steps on all nodes:
      a. Use the `iris_cli add vlan` command to set up the non-native VLAN to be used for the Cohesity cluster create workflow.

      ```
      iris_cli vlan add interface-group=<name of the interface group> id=<VLAN-ID> subnet-mask-bits=<Subnet-Mask-Prefix>
      ```

      Example:

      ```
      iris_cli vlan add interface-group=intf_group1 id=5 subnet-mask-bits=20
      ```

b. Use the following command to set the non-native VLAN logical bond interface as primary. Replace vlan_id with the ID of the VLAN you added.

```
iris_cli ip config interface-name=<bond0.vlan_id> interface-role=primary
```

Example:

```
iris_cli ip config interface-name=bond0.5 interface-role=primary
```

Alternatively, to configure the IP on each node and access the node using the IP (not required if using Avahi to discover all nodes), use this command:

```
iris_cli ip config interface-name=<bond0.vlan_id> interface-ips=xx subnet-gateway=yy subnet-mask-bits=zz interface-role=primary
```

Example:

```
iris_cli ip config interface-name=bond0.5 interface-ips=10.5.1.5 subnet-gateway=10.5.0.1 subnet-mask-bits=20 interface-role=primary
```

2. Through configure_network.sh script: Perform these steps on all nodes.

   Use `configure_network.sh option 10` to add VLAN and follow the prompts to complete VLAN addition on the desired network bond.

   Location: /home/cohesity/bin/network/configure_network.sh

4. Restart the Cluster Manager Services:

```
sudo service nexus restart
```

5. Run `ifconfig` and ensure Avahi runs on the non-native VLAN bonded interface. Cohesity relies on Avahi for node discovery.

   Sample output from a node running on the primary non-native VLAN 5

```
bond0.5:avahi: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 169.254.7.187 netmask 255.255.0.0 broadcast 169.254.255.255
ether a4:bf:01:2d:7e:b2 txqueuelen 1000 (Ethernet)
```

6. Perform these steps on any one node:

   If an IP was assigned on the VLAN interface, access the UI using the assigned IP on the VLAN interface. Otherwise, complete these steps:

   a. Connect the Mac laptop to the node for Avahi discovery for Cohesity cluster create.

b. Perform discovery through the browser, using this URL format:

   https://chassis-<chassis_serial_number>-<node_number>.local

7. On any node, start the Cohesity cluster create workflow and provide Cohesity cluster IPs from the configured non-native VLAN.

> **Note:** If necessary, the user can configure Cohesity cluster IPs and VIPs from the non-native VLAN and retain the IPMI in the native VLAN or some other subnet.

Expand a Cohesity cluster

To expand a Cohesity cluster, perform the following steps before adding new nodes to the Cohesity cluster.

1. You can use one of the following methods to achieve it:
   i. Through iris_cli
      a. Use the `iris_cli add vlan` command to set up the non-native VLAN to be used for the node add workflow.

```
iris_cli vlan add interface-group=<name of the interface group> id=<VLAN-ID> subnet-mask-bits=<Subnet-Mask-Prefix>
```

Example:

```
iris_cli vlan add interface-group=bond0 id=5 subnet-mask-bits=20
```

      b. Use the following command to set the non-native VLAN logical bond interface as primary. Replace vland_id with the ID of the VLAN you added.

```
iris_cli ip config interface-name=<bond0.vland_id> interface-role=primary
```

Example:

```
iris_cli ip config interface-name=bond0.5 interface-role=primary
```

Alternatively, to configure the IP on a new node and access the node using the IP (not required if using Avahi to discover all nodes), use this command:

```
iris_cli ip config interface-name=<bond0.vlan_id> interface-ips=xx subnet-gateway=yy subnet-mask-bits=zz interface-role=primary
```

Example:

```
iris_cli ip config interface-name=bond0.5 interface-
ips=10.5.1.5 subnet-gateway=10.5.0.1 subnet-mask-bits=20
interface-role=primary
```

ii. Through configure_network.sh script

Use `configure_network.sh option 10` to add VLAN and follow the prompts to complete VLAN addition on the desired network bond.

Location: /home/cohesity/bin/network/configure_network.sh

2. Restart the Cluster Manager Services:

```
sudo service nexus restart
```

3. Run `ifconfig` and ensure Avahi runs on the non-native VLAN bonded interface. Cohesity relies on Avahi for node discovery.

Sample output from a node running on the primary non-native VLAN 5

```
bond0.5:avahi: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 169.254.7.187 netmask 255.255.0.0 broadcast 169.254.255.255
ether a4:bf:01:2d:7e:b2 txqueuelen 1000 (Ethernet)
```

4. On any node in the existing Cohesity cluster, start the node add workflow from the UI and provide Cohesity cluster IPs from the configured non-native VLAN.

> **Note:** If necessary, the user can configure Cohesity cluster IPs and VIPs from the non-native VLAN and keep the IPMI in the native VLAN or some other subnet.

Manage Firewall Rules and Ports

Using the Cohesity Dashboard or CLI, you can configure application-based firewall rules and manage incoming traffic on a Cohesity cluster. For more details, see Manage Application-Based Firewall Rules.

To allow the Cohesity cluster to transmit and receive data, you must open specific ports. For more details, see Manage Firewall Ports.

Manage Application-Based Firewall Rules

Firewall restricts network access and manages inbound traffic on Cohesity cluster. It allows you to secure your cluster from unauthorized access and malicious traffic. Cohesity lets you configure application-based firewall rules using firewall profiles. Firewall profiles are rules that define whether data from specific applications, ports, protocols, and ipsets (source IP address sets) are allowed or denied through one or more interface groups. You can

customize the supported firewall profiles based on your requirements and apply them to one or more interface groups.

> **Note:** Currently, you can apply firewall rules only on the incoming traffic on the Cohesity cluster.

The following table lists the firewall profiles supported on a Cohesity cluster.

Port (Protocol)	Source	Target	Profile	Description
80 (TCP)	Web UI	Cohesity cluster	Management	External management (HTTP/HTTPS) traffic. Configure the firewall settings to allow the VIPs.
443 (TCP)	Web UI	Cohesity cluster	Management	External management (HTTP/HTTPS) traffic. Configure the firewall settings to allow the VIPs.
	Cohesity cluster	Disaster Recovery Cluster	Replication	Required for replication Configure the firewall settings to allow the Cohesity IPs (VIPs or node IPs).

Port (Protocol)	Source	Target	Profile	Description
22 (TCP)	SSH Client	Cohesity cluster	SSH	SSH client. Support Channel used by Cohesity Support for troubleshooting. Configure the firewall settings to allow the Cohesity IPs (VIPs or node IPs).
68 (UDP)	Cohesity cluster	DHCP server	DHCP	DHCP traffic. Configure the firewall settings to allow the VIPs.
11111 (TCP)	Cohesity cluster	Disaster Recovery Cluster	Internal Traffic	I/O Operations Service. Configure the firewall settings to allow the Cohesity node IP addresses and not the VIPs.
111 (TCP/UDP)	Mount target	Cohesity cluster	Data Protection SmartFiles Recovery	NFS restores on Linux physical servers. Configure the firewall settings to allow the VIPs.

Port (Protocol)	Source	Target	Profile	Description
2049 (TCP/UDP)	ESXi host	VIPs and VLAN IPs on Cohesity cluster	Data Protection Recovery	VM restores. Configure the firewall settings to allow the VIPs.
	NFS Clients	Cohesity NFS target	Data Protection SmartFiles	Required for NFS Configure the firewall settings to allow the VIPs.
445 (TCP)	SMB Clients	Cohesity SMB target	Data Protection SmartFiles	Traffic for SMB mount. Configure the firewall settings to allow the VIPs.
3000 (TCP)	S3 Clients	Cohesity cluster	Data Access	Access to Cohesity views using S3 protocol. Configure the firewall settings to allow the VIPs.
11113 (TCP)	Oracle Source (Cohesity Agent)	Cohesity cluster	Backup and Recovery	Source side deduplication for Oracle when using the Remote Adapter RMAN script. Configure the firewall settings to allow the VIPs.

Port (Protocol)	Source	Target	Profile	Description
11117 (TCP)	Backup Target (any port)	Cohesity cluster	Backup	Secure gRPC connections on the cluster from the backup target. Configure the firewall settings to allow the VIPs.
20004 (TCP)	Cohesity cluster	Disaster Recovery Cluster	Replication	Required for replication. Configure the firewall settings to allow the Cohesity node IP addresses and not the VIPs.
20000 (TCP)	Cohesity cluster	Disaster Recovery Cluster	Replication	Interaction handling between Cohesity and target backup systems. Configure the firewall settings to allow the Cohesity node IP addresses and not the VIPs.
161 (UDP), 162 (UDP)	SNMP	Cohesity cluster	SNMP	For SNMP traffic, if the cluster is configured to use SNMP. Configure the firewall settings to allow the VIPs.

Port (Protocol)	Source	Target	Profile	Description
No specific ports	Archival target	Cohesity cluster	Archival	Cohesity cluster uses HTTP/HTTPS to archive data to the target. The connection is not initiated from the archive target to the Cohesity cluster.
27999	Cohesity cluster	Third-Party tools (for example Tableau, SSRS etc.)	Reporting Database	Used for connections to Custom Reporting Database.

Configure Firewall Rules

Using the Cohesity Dashboard or CLI, you can configure firewall rules and apply them to one or more interface groups on a cluster.

> **Note:** If you configure multiple bond interfaces on a cluster for various traffic types, ensure that you configure the same bond interfaces on the free nodes before adding them to the cluster.

To configure firewall rules using the supported profiles:

1. In the Cohesity Dashboard, select **Settings > Networking** and then click the **Firewall** tab.

   The Rules page appears and displays all the supported firewall profiles. By default, all the incoming application traffic listed in the firewall profiles are allowed on all the available interface groups.

2. Click the edit icon ( ).
3. Provide the details in the applicable fields for the application traffic you plan to manage:
   - **Application:** Predefined name of the firewall profile. Specifies an external service or a source of the incoming traffic.
   - **Interface Group:** Select one or more interface groups from the drop-down menu. By default, all the interface groups are selected.

- **Client's IPSet:** Select **Custom IP/Subnet** from the drop-down menu and specify the IP address and subnet mask bits of the source network. You can enter multiple IP addresses and subnet mask sets. For example, 10.1.1.2/24 and 10.2.3.2/16. By default, the **All IP Addresses** option is selected.
- **Action:** Select **Allow** or **Deny** from the drop-down menu to allow or restrict the traffic. By default, the **Allow** option is selected.

4. Optional. Repeat **Step 3** to customize the firewall profiles and manage other network traffic types based on your requirements.
5. Click **Apply** to save the configuration and apply the firewall rules.

Restore Firewall Rules to the Default Settings

If you misconfigured firewall rules or if you want to reset all your firewall settings, you can restore the configured firewall rules to their default settings.

> **Important:** If you restore firewall rules to default, all the configured firewall rules will be lost.

To restore firewall rules to the default settings:

1. In the Cohesity Dashboard, select **Settings > Networking** and then click the **Firewall** tab.

   The Rules page appears and displays all the supported firewall profiles.

2. Click the reset icon (↺).
3. Click **Restore** to confirm the reset operation.

You can also use the Cohesity CLI to configure the firewall rules or restore the rules to their default settings. For more details, see Cohesity CLI Commands Used to Manage Firewall Ports.

Manage Firewall Ports

You must open certain ports in the firewall to allow the Cohesity cluster to transmit and receive data. The cluster sends the following types of traffic over the network. You can isolate traffic on physical or logical networks to improve performance and security.

- IPMI traffic for admins to access nodes, typically for pre-boot BIOS access and post-boot console access.
- Management traffic:
  - from Nodes to external services such as Active Directory, DNS, NTP, and SNMP
  - for admins to access the Cohesity cluster
- Backup and restore traffic for moving data between protected sources and the cluster.

- Replication traffic for replicating data from one cluster to another cluster.
- Data access (NAS) traffic for accessing data on the cluster by using protocols such as NFS, SMB, or S3.
- Archive and tiering traffic for moving data to external targets such as a cloud or tape system.
- Support Channel traffic between nodes and the Cohesity Support servers.
- Internal traffic between nodes on a cluster.

> **Note:**
> Note the following when applying firewall rules on a Cohesity cluster:
>
> - Cohesity uses Virtual IPs (VIPs) to manage incoming traffic to the cluster.
> - Cohesity uses Host IPs (node IP addresses) for outgoing traffic and node-to-node communication within the cluster.
> - When setting up replication on a Cohesity cluster, you can specify either the VIPs or node IP addresses of the remote cluster. Cohesity recommends specifying VIPs, as they can failover to other nodes if a node becomes unreachable or goes down.
> - Do not enable Deep Packet Inspection (DPI) for replication workflows and Teleport support traffic, as it can degrade replication performance and impact Teleport functionality.
> - Ensure to have network connectivity between the source and the cluster.

The following diagram shows the port requirements and directional flow for some of the traffic in a typical cluster. More information is provided following the diagram.

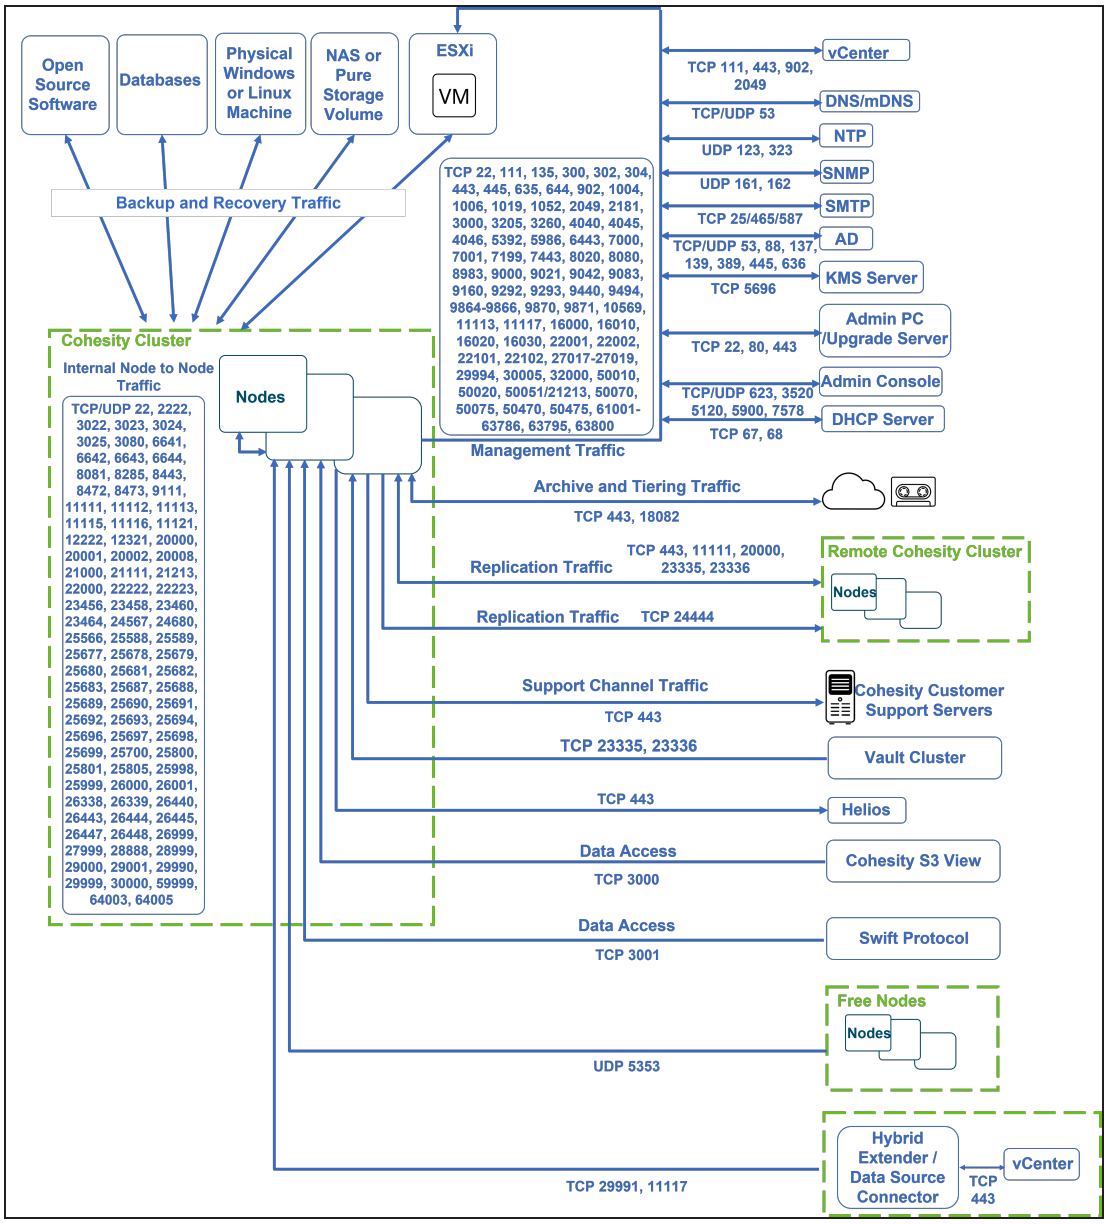

Using the Cohesity UI or CLI, you can configure application-based firewall rules and manage incoming traffic on a Cohesity cluster. For more details, see Manage Application-Based Firewall Rules.

Identify Source IP of Outgoing Traffic

The source IP of outgoing traffic depends on the matching static route's outgoing interface. If there is no static route, the source IP depends on the default route's outgoing interface. Use the route list command to view the list of routes configured on an interface in a Cohesity cluster and the IP address of the destination network.

In the following example, the route list command output shows the VLAN interface (intf_group1.100) of a cluster configured with the default route.

```
iris_cli route list
```

```
DEST NETWORK                    : 0.0.0.0/0
NEXT HOP                        : 10.1.100.1
INTERFACE GROUP NAME            : intf_group1.100
```

For more details, see Add a Static Route.

Cluster Management

Cluster Upgrade

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	• https://downloads.cohesity.com/oauth2/login • https://s3-us-west-2.amazonaws.com • Amazon CloudFront CDN (Content Delivery Network) Facility: https://dipowd173xxgy.cloudfront.net/	443	TCP	For cluster upgrade packages. You can select and download the version packages from downloads.cohesity.com, AWS location, or CloudFront site. The cluster upgrade packages are hosted on the Amazon S3 CloudFront server to increase the download speed of release packages.

Agent Upgrade

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Client	Cohesity cluster	80,443	TCP	For Agent upgrade from UI.

> **Note:** The ports listed here are designated solely for the initial upgrade of agents from versions older than 7.2.2. For all subsequent upgrades on newer agent versions, only the agent port will be utilized.

Troubleshooting

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
SSH client	Cohesity cluster	22	TCP	Support Channel used by Cohesity Support for troubleshooting

Networking

Agent Communication

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Cohesity Agent	50051*	TCP	Used for Cohesity Agent and agentless data transfer to the cluster. Needs to be open on all nodes on the cluster and any backup target using the Cohesity Agent, such as SQL servers, physical servers, and VMs running the Cohesity-installed Agent. Cohesity clusters pick any available ephemeral port to initiate communication with the agent service via port 50051. The firewall needs to allow the connection request from the Cohesity nodes if the destination port is 50051. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Backup and Recovery

Network Services

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	DNS	53	TCP/UDP	DNS, if an external DNS server is configured.	Management
Cohesity cluster	DNS server	5353	UDP	mDNS multicasting.	

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	External NTP Server (time.google.com)	123	UDP	Requests to external NTP servers are sent via port 123.	Management
		323		• Required for cluster setup. • Chrony is the default implementation of NTP used by recent versions of CentOS, RHEL, and AlmaLinux. Open port 323 if you want to use the chronyc tool to monitor the synchronization status of Chrony and make changes if necessary.	

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Cohesity cluster	123	UDP	Time synchronization among Cohesity nodes during node-to-node communication is implemented via port 123.	Management
Cohesity cluster	DNS	53	TCP/UDP	Required if an internal DNS server is configured.	
Cohesity cluster	Cohesity cluster	5353	UDP	mDNS multicasting.	

Cohesity IPMI Network

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Admin (Browser)	Cohesity IPMI address on each node	80, 443	TCP	Required for HTTP and HTTPS	Management and IPMI
		3520, 5900		Required for IKVM web browser access	

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster (bond0 on each node)	Cohesity IPMI address on each node	623	UDP	IPMI, Intel standard manageability port.	IPMI
		7578	TCP		
Admin		5120	UDP/TCP	Required for virtual media mount.	

SMB and NFS Media Servers

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
NFS Clients	Cohesity NFS target	2049	TCP	Required for NFS	Data Protection
SMB Clients	Cohesity SMB target	445		SMB filer functionality. SMB and SMB2 restore.	Data Access

NFS For Cohesity Views

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
NFS Client	Cohesity cluster (with NFS server)	2049 (NFS server), 111 (port mapper)	NFS	Port 111 needs to be open on the NFS client and server. NFS Linux clients do not need any additional configuration on the Cohesity cluster	Required for mounting the Backup and Recovery views.

DHCP Server

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	DHCP Server	67, 68	TCP	Required for DHCP. Required for SMB and ICMP.	Management

SNMP

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
SNMP	Cohesity cluster	161	UDP	SNMP, if the cluster is configured to use SNMP.	Management

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Network Management System (NMS)	162	UDP	SNMP, if the cluster is configured to use SNMP traps.	Management

SMTP

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	SMTP server	Frequently opened ports on the SMTP server are either 25 or 465 or 587	TCP	Depends on the port setting on the SMTP server. If SMTP opens a random port, the same port must be opened on the Cohesity cluster.	Management

Virtualization

VMware

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
ESXi host	VIPs and VLAN IPs on the Cohesity cluster	111 2049	TCP	Required for VMDK restore with Instant Recovery Required for VMDK Restore Required for Test & Dev clone Required for File Recovery and Instant Volume Mount for Windows VMs when using agent-based recovery.	Recovery

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	ESXi Host	902	TCP	Ensure that SSL communication for TCP port 902 to the ESXi host is enabled. Otherwise, backups will fail.	Backup and Recovery

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	VMware vCenter ESXi Host VMware vCloud Director	443	TCP	The Cohesity cluster uses HTTPS by default. Required for VMware API. Required for VMware Tools-based file and folder recoveries. Required for HTTPS/HTTPS (TLS). Required for VMware vRealize Automation (vRA) and Orchestration (vRO) integration.	Backup and Recovery

Microsoft SCVMM and Hyper-V

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Windows machine with Cohesity Agent installed	SCVMM	5986	TCP	SCVMM machines can be LOCALHOST or remote targets.	Management
	Hyper-V			The Hyper-V host machine is LOCALHOST.	Backup and Recovery
Cohesity cluster	Host	50051*	TCP	Required for both SCVMM machines and Hyper-V host machines. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Backup and Recovery

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Destination or target Windows machine	445 (SMB)	TCP	Required for agentless file recovery	Backup and Recovery
Windows machine with Cohesity Agent installed	Destination or target Windows machine	135 (WMI)	TCP	Required for agentless file recovery	Backup and Recovery

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Host	Cohesity cluster	445	TCP	Required by the Hyper-V host to access SMB data.	Recovery

Nutanix Acropolis

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Data Services IP and IP of all CVMs on the Nutanix cluster and Prism Central.	VIPs and VLAN IPs on the Cohesity cluster	111 2049	TCP	Required for Instant Recovery	Backup and Recovery

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Data Services IP and IP of all CVMs on the Nutanix cluster and Prism Central.	9440	TCP	Used for REST API calls	Backup and Recovery
		3205, 3260		iSCSI port for AHV backups.	

RHV/OLVM

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Host	Cohesity cluster	50051*	TCP	Required for RHV/OLVM host machines. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Backup and Recovery

Instant Volume Mount

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
ESXi host	VIPs and VLAN IPs on Cohesity cluster	111, 2049	TCP	Instant Volume Mount	Recovery

File Restore to Windows VM

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
ESXi host	VIPs and VLAN IPs on Cohesity cluster	111, 2049	TCP	File restore to Windows VM	Recovery

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Host	Cohesity cluster	50051*	TCP	Required for persistent agent workflow for file-level recovery and IVM. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Recovery

File Restore to Linux VM

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Linux VM	VIPs and VLAN IPs on the Cohesity cluster	111, 2049	TCP	Required for file restore to Linux VM.	Recovery

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Host	Cohesity cluster	50051*	TCP	Required for Linux VM file restores. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Backup and Recovery

NAS

Isilon

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Isilon	Cohesity cluster	111	TCP/UDP	Required for RPC connection.	Backup and Recovery
		300, 302, 304, 2049		Required for NFS.	
		8080		Required for HTTP connection.	
		443	TCP	Required for HTTPS connection with Isilon.	
		445		Required for SMB.	

NetApp

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
NetApp	Cohesity cluster	111	TCP/UDP	Required for RPC connection.	Backup and Recovery
		443		Required for HTTPS connection with NetApp.	
		635, 2049, 4045, 4046		Required for NFS.	
		445	TCP	Required for SMB.	
		1052, 10569	TCP	Required for SnapDiff-based backup.	
		3000	TCP	Required by NetApp clusters to write to Cohesity S3 storage during SnapDiff backups.	

Nutanix Files

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Nutanix Files	Cohesity cluster	111	TCP/UDP	NFS Portmapper /RPCBIND	Backup and Recovery
		2049		NFS file access (v3/v4)	
		635, 4046, 2049, 4045		Required for NFS.	
		445	TCP	SMB Direct Host/ Microsoft-DS	
		9440		REST API Calls.	

Generic NAS

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Generic NAS	Cohesity cluster	111	TCP/UDP	Required for RPC connection.	Backup and Recovery
		635 2049		Required for NFS.	
		445	TCP	Required for SMB.	

Pure FlashBlade

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity	Pure Storage	3260, 3205	TCP	iSCSI and FC port for Pure Storage.	Backup and Recovery

Elastifile

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Target	Destination Port	Protocol	Usage Notes	Type of Traffic
Elastifile	Cohesity cluster	443	TCP	Required for EMS/API operations.	Backup and Recovery
		111, 2049, 644, 4040, 4045	TCP/UDP	Required for NFS Operations.	

Databases

MS SQL Servers

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Server	Cohesity cluster	445	TCP	Required for SQL servers.	Backup and Recovery
MS SQL server (any port)	Cohesity cluster	11113, 11117		Required for VDI-based backup and restore of the MS SQL database (using RemoteWANSnapFS).	
MS SQL server (any port)	Cohesity cluster	11117		Required for log backup and restore of MS SQL database.	

Oracle

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Oracle Source (Cohesity Agent)	Cohesity cluster	11113	TCP	Source-side deduplication for Oracle when using the Remote Adapter RMAN scriptX. Secure gRPC used by Oracle SBT.	Backup and Recovery
Oracle Source	Cohesity cluster	11117	TCP	Required from 6.8.1_u3 onwards with the introduction of secure gRPC for Oracle SBT.	Backup and Recovery

PostgreSQL

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
PostgreSQL	Cohesity cluster	11113	TCP/IP	Required for physical backups and recovery at the cluster level.	Backup and Recovery

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	PostgreSQL	50051*	TCP/IP	Required for physical backups and recovery at the cluster level. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Backup and Recovery

SAP HANA

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	DestinatiAon Port	Protocol	Usage Notes	Type of Traffic
SAP HANA	Cohesity cluster	11113	TCP	Access to Cohesity views.	Backup and Recovery

Cassandra

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Cassandra nodes	22	TCP	Required for Cassandra backup. Required for Cassandra optimized restore/DSE 6.x restore.	Backup and restore
		7000		Required for Cassandra bulk restore.	Restore
		7001		Required for Cassandra bulk restore in case of an SSL-enabled cluster.	
		7199		Required for Cassandra repository registration. Required for Cassandra backup (snapshot operation). Required for Fetching Cassandra configurations (list of nodes, token map, etc).	Management
		8983		Required for Cassandra solr index management.	
		9042		Required for Cassandra backup. Required for Cassandra restore. Required for Cassandra source registration.	
		9160		Required for Cassandra bulk restore (DSE 4.x / Cassandra 2.x).	

MongoDB

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	MongoDB Nodes	27017, 27018, 27019	TCP	Required for MongoDB backups and recovery.	Backup and Recovery

Open-Source Software Platform

Hadoop - HDFS

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

The following ports must be open to enable communication between Cohesity and **Cloudera Data Platform 7.1.x**. These are the default ports used. Refer the configuration in Cloudera Manager if you have overridden these ports.

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Hadoop Primary node	8020	TCP	Required to connect to Namenode through native protocol	Backup and Recovery
Cohesity cluster	Hadoop Primary node	9870	TCP	Required to connect to Namenode through REST APIs (Non-SSL)	Backup and Recovery
Cohesity cluster	Hadoop Primary node	9871	TCP	Required to connect to Namenode through REST APIs (SSL)	Backup and Recovery
Cohesity cluster	Hadoop Primary node	9864	TCP	Required to connect to datanode through REST APIs (Non-SSL)	Backup and Recovery
Cohesity cluster	Hadoop Primary node	9865	TCP	Required to connect to datanode through REST APIs (SSL)	Backup and Recovery
Cohesity cluster	Hadoop Primary node	9866	TCP	Required to connect to datanode through native protocol	Backup and Recovery
Cohesity cluster	Hadoop Primary node	9292	TCP	HTTP port required to connect to KMS	Backup and Recovery
Cohesity cluster	Hadoop Primary node	9494	TCP	HTTPS port required to connect to KMS (SSL)	Backup and Recovery

Hadoop - Hive

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

The following ports must be open to enable communication between Cohesity and **Cloudera Data Platform 7.1.x**.

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Hadoop Primary node	9083	TCP	Required to connect to Hive Metastore through the native protocol (Thrift)	Backup and Recovery

Hadoop - HBase

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

The following ports are required for **Cloudera Data Platform 7.1.x**.

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Hadoop Primary node	22001	TCP	Required to connect to HMaster through native protocol	Backup and Recovery
Cohesity cluster	Hadoop Primary node	22002	TCP	Required to connect to HMaster through HTTP protocol	Backup and Recovery
Cohesity cluster	Hadoop Primary node	22101	TCP	Required to connect to Region server through native protocol	Backup and Recovery
Cohesity cluster	Hadoop Primary node	22102	TCP	Required to connect to Region server through HTTP protocol	Backup and Recovery
Cohesity cluster	Hadoop Primary node	2181	TCP	Required to connect to Zookeeper service through native protocol	Backup and Recovery

gRPC

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Backup Target (any port)	Cohesity cluster	11117	TCP	Required for secure gRPC connections on the cluster from the backup target.	Backup

Kubernetes

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Kubernetes cluster	VIPs on Cohesity cluster	3000	TCP	Velero S3 communication for metadata backup and restore.	Backup and Recovery

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Kubernetes cluster api-address	6443	TCP	Kubernetes cluster and Velero API Access. For more details on Velero support, see the Considerations section.
Cohesity cluster	Kubernetes cluster worker nodes	32000	TCP	Cohesity cluster to Cohesity Datamover communication

S3-Compatible Storage

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Note	Type of Traffic
Cohesity cluster	Dell EMC ECS	9021	TCP	Required for backup	Backup and Recovery
Cohesity cluster	Pure Storage Flashblade	443	TCP	Required for backup	Backup and Recovery
Cohesity cluster	NetApp StorageGrid	443	TCP	Required for backup	Backup and Recovery
Cohesity cluster	IBM Cloud Object Storage	443	TCP	Required for backup	Backup and Recovery

Cloud VMs

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic

Cohesity cluster	• **AWS**
	  • *amazonaws.com
	• **Azure**
	  • Login.windows.net
	  • management.azure.com
	  • *.blob.core.windows.net
	• **GCP**
	  • *.googleapis.com
		443	TCP	Required for CloudSpin, CloudRetrieve, and Native Snapshot
				XL Cohesity Cloud Edition and Cohesity Cloud Edition deployment.
				Required for HTTPS connections.
					Backup and Recovery

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
	Fleet instances	50051*		Required for CloudSpin to AWS EC2. Required for Native Snapshot backup for AWS/GCP. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both.	

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
		Required for Native Snapshot backup for AWS/GCP. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment			

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
		uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.			

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cloud Source VMs		Required for file-level recovery from Native Snapshot backup. Required for Agent Protection and Restore for AWS EC2. Required for Agent Protection and Restore for Azure VM. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port			

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
	Required for Agent Protection and Restore for AWS EC2.				

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cloud Source VMs	Cohesity Cluster	50051*	TCP	Required for file-level recovery from Native Snapshot backup. Required for Agent Protection and Restore for AWS EC2. Required for Agent Protection and Restore for Azure VM. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Backup and Recovery

Amazon Elastic Block Store

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS EBS Direct API	443	TCP	AWS API – Backup and Recovery

AWS Cloud Database Workloads

> **Note:** Protection (ingest-based protection) of AWS cloud database workloads is supported only on Cohesity Cloud Edition on AWS.

Amazon Aurora MySQL

Bidirectional

Amazon Elastic Block Store	Amazon Elastic Block Store	Amazon Elastic Block Store	Amazon Elastic Block Store	Amazon Elastic Block Store
Amazon Elastic Block Store	Amazon Elastic Block Store	Amazon Elastic Block Store	Amazon Elastic Block Store	Amazon Elastic Block Store
Cohesity Cluster	AWS Aurora MySQL Endpoint	3306 or any custom port	TCP	Backup and Recovery (direct DB connection)

Amazon Aurora PostgreSQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS Aurora PostgreSQL Endpoint	443	TCP	AWS API access, source registration
Cohesity Cluster	AWS Aurora PostgreSQL Endpoint	5432 or any custom port	TCP	Backup and Recovery (direct DB connection)

Amazon RDS for MySQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS RDS MySQL Endpoint	443	TCP	AWS API access, source registration
Cohesity Cluster	AWS RDS MySQL Endpoint	3306 or any custom port	TCP	Backup and Recovery (direct DB connection)

Amazon RDS for PostgreSQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS RDS PostgreSQL Endpoint	443	TCP	AWS API access, source registration
Cohesity Cluster	AWS RDS PostgreSQL Endpoint	5432 or any custom port	TCP	Backup and Recovery (direct DB connection)

Amazon RDS for Oracle

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS RDS Oracle Endpoint	443	TCP	AWS API access, source registration
Cohesity Cluster	AWS RDS Oracle Endpoint	1521 or any custom port	TCP	Backup and Recovery (direct DB connection)

Amazon RDS for MS SQL Server

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS RDS MSSQL Endpoint	443	TCP	AWS API access, source registration
Cohesity Cluster	AWS RDS MSSQL Endpoint	1433 or any custom port	TCP	Backup and Recovery (direct DB connection)

Amazon DynamoDB

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS DynamoDB API	443	TCP	AWS API – Backup and Recovery

Amazon DocumentDB

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS DocumentDB Endpoint	443	TCP	AWS API access, source registration
Cohesity Cluster	AWS DocumentDB Endpoint	27017 or any custom port	TCP	Backup and Recovery (direct DB connection)

Amazon Redshift

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS Redshift Endpoint	443	TCP	AWS API access, source registration

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AWS Redshift Endpoint	5439 or any custom por	TCP	Backup and Recovery (direct cluster connection)

RDS Custom for MS SQL Server

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	MS SQL server	445	TCP	Required for SQL servers.
Cohesity Cluster	MS SQL server	11113, 11117	TCP	Required for VDI-based backup and restore of the MS SQL database (using RemoteWANSnapFS)
Cohesity Cluster	MS SQL serverr	11117	TCP	Required for log backup and restore of MS SQL database.

RDS Custom for Oracle

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	Oracle Host	50051*	TCP	Backup and Recovery
Oracle Host	Cohesity cluster	111	TCP/UDP	Backup and Recovery
Oracle Host	Cohesity cluster	2049	TCP/UDP	Backup and Recovery

Azure Cloud Database Workloads

> **Note:** Protection of Azure cloud database workloads is supported only on Cohesity Cloud Edition on Azure.

Azure Blob Storage and Data Lake

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure Blob / Data Lake Endpoint	443	TCP	Backup and Recovery

Azure MySQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure MySQL	3306	TCP	Backup and Recovery

Azure Cosmos DB – Cassandra

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure Cosmos DB	10350	TCP	Backup and Recovery

Azure Cosmos DB – MongoDB

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure Cosmos DB for MongoDB	10255	TCP	Backup and Recovery

Azure Cosmos DB – NoSQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure Cosmos DB for NoSQL	443	TCP	Backup and Recovery

Azure Cosmos DB – Table

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration

Azure SQL Database

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure SQL Server	1433	TCP	Backup and Recovery

Azure SQL Managed Instance

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure SQL Managed Instance	1433	TCP	Backup and Recovery

Azure Table Storage

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure Blob / Data Lake Endpoint	443	TCP	Backup and Recovery

Azure PostgreSQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Azure Management API	443	TCP	Azure API access, source registration
Cohesity Cluster	Azure PostgreSQL	5432	TCP	Backup and Recovery

Google Cloud Database Workloads

> **Note:** Protection of Google Cloud database workloads is supported only on Cohesity Cloud Edition on GCP.

BigQuery

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	BigQuery	443	TCP/IP	Backup and Recovery

Cloud SQL for MySQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Cloud SQL for MySQL	443	TCP/IP	Backup and Recovery
Cohesity Cluster	Cloud SQL for MySQL	3306	TCP/IP	Backup and Recovery

Cloud SQL for PostgreSQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Cloud SQL for PostgreSQL	443	TCP/IP	Backup and Recovery
Cohesity Cluster	Cloud SQL for PostgreSQL	5432	TCP/IP	Backup and Recovery

AlloyDB for PostgreSQL

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	AlloyDB for PostgreSQL	443	TCP/IP	Backup and Recovery
Cohesity Cluster	AlloyDB for PostgreSQL	5432	TCP/IP	Backup and Recovery

Cloud SQL for SQL Server

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Cloud SQL for SQL Server	443	TCP/IP	Backup and Recovery
Cohesity Cluster	Cloud SQL for SQL Server	1433	TCP/IP	Backup and Recovery

Firestore

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Firestore	443	TCP/IP	Backup and Recovery

Spanner

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Spanner	443	TCP/IP	Backup and Recovery
Cohesity Cluster	Spanner	9010	TCP/IP	Backup and Recovery (Dataflow)

Bigtable

Bidirectional

Source	Destination	Port	Protocol	Type of Traffic
Cohesity Cluster	Bigtable	443	TCP/IP	Backup and Recovery

External Targets

Cloud Workflow

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Cohesity cluster	29999	TCP	Required for Archive Services.	Backup and Recover

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic

Cohesity cluster	• **AWS**
	  • *amazonaws.com
	• **Azure**
	  • *.blob.core.windows.net
	  • *vault.azure.net
	• **GCP**
	  • *.googleapis.com
		443	TCP	Required for CloudArchive, CloudTier.	Backup and Recovery

Oracle Cloud Infrastructure

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	• OCI ◦ *.oraclecloud.com	443	TCP	Required for CloudArchive, CloudRetrieve, and CloudTier to Oracle Cloud Infrastructure.	Backup and Recovery

Cohesity S3 View

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity S3 View	Cohesity cluster	3000	TCP	Access to Cohesity views using the S3 protocol.	Data Access

QStar

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	QStar ASM host	18082	TCP	QStar External Target - Port number 18082 is set on the QStar host by default. If your QStar host is not using the default Web Services Port, update the port number. SOAP over HTTPS.	Backup and Recovery

Storage Arrays

Nimble Storage

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Nimble Storage	5392	TCP	REST API	Backup and Recovery
		3260, 3205		iSCSI port for Nimble Storage.	

Pure Storage

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity	Pure Storage	3260, 3205	TCP	iSCSI port for Pure Storage.	Backup and Recovery

IBM Storage FlashSystem

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity	IBM Storage FlashSystem	3260, 7443	TCP	3260 port for iSCSI 7443 port for API	Backup and Recovery

Active Directory

Microsoft Active Directory (RBAC)

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Client	Cohesity	53	TCP/UDP	Serve DNS requests from an external source.	Management
Cohesity cluster	Active Directory	445	TCP	Required only when initially joining the Cluster to Active Directory.	

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	Kerberos Key Distribution Center (AD)	88	TCP/UDP	Required for Kerberos if the cluster is configured to use Active Directory.	Management
	LDAP	389		Required if the cluster is configured to use Active Directory or LDAP.	Management
	Active Directory	137	UDP	User and Computer Authentication and NetLogon, NetBIOS Name Resolution. Required only when initially joining the Cluster to Active Directory.	Management
		139	TCP	Required only when initially joining the Cluster to Active Directory (for the NetBIOS session service).	Management

Microsoft Exchange Servers

No specific ports need to be opened to protect Microsoft Exchange Servers, provided the following requirements have been met as part of your cluster configuration:

- The ports listed in the Physical Servers section are open to allow communication between your Physical Servers and the Cohesity cluster for backup operations.
- The ports listed in the SMB and NFS Media Servers section are open to allow communication between the Cohesity cluster and SMB clients for recovery operations.
- The ports listed in the Active Directorysection are open to allow communication between the Cohesity cluster and the Active Directory for recovery operations.

Physical Servers

Physical Server Backup

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Mount Target	Cohesity cluster	111	TCP/UDP	NFS restores on Linux physical servers.
Local Host (Physical Windows or Linux Server)	Local Host (Physical Windows or Linux Server)	59999	TCP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by your firewall.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Physical Linux Server	Cohesity cluster	111, 2049	TCP	Required for NFS block-based file recovery.
Physical Windows Server	Cohesity cluster	445	TCP	Required for SMB block-based file recovery.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Physical Server with Cohesity Agent installed	50051*	TCP	Linux and Windows physical server backups. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.

Bare Machine Recovery

CBMR

Port	Protocol	Source	Target	Usage Notes
50051*	TCP	Cohesity cluster	Physical server	Backup operations. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.
111, 2049	TCP	Cohesity cluster	Physical server	NFS restore operations on a Linux server.
445	TCP	Cohesity cluster	Physical server	SMB restore operations on a Windows server.
443	TCP	Recovery environment	Cohesity cluster	To connect the CBMR recovery environment with the Cohesity cluster.

Port	Protocol	Source	Target	Usage Notes
50051*	TCP	Cohesity cluster	Recovery environment	To connect the CBMR recovery environment with the Cohesity cluster. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.

CoBMR

Port	Protocol	Source	Target	Usage Notes
50051*	TCP	Cohesity cluster	Physical server	Backup operations. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.
111, 2049	TCP	Cohesity cluster	Physical server	NFS restore operations on a Linux server.

Port	Protocol	Source	Target	Usage Notes
445	TCP	Cohesity cluster	Physical server	SMB restore operations on a Windows server.
443	TCP	Recovery environment	Cohesity cluster	Port 443 is only required to perform a CoBMR recovery. It is not required for backup. You need this port open on the target system to which you are performing the recovery.
50051*	TCP	Cohesity cluster	Recovery environment	To connect the CBMR recovery environment with the Cohesity cluster. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.

Key Management Service

KMS (Key Management Service)

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster	SafeNet KeySecure server	5696	TCP	Required if the cluster is configured to use an external Key Management Service (KMS) such as SafeNet KeySecure. 5696 is the default KMIP port, and can be changed when configuring the external KMS. For more information, see External Key Management Service.	Management

OpenStack Integration

Swift Protocol

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Swift Clients	Cohesity cluster	3001	TCP	Access to Cohesity views using the Swift protocol.	Data Access

Helios Ports

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	https://helios.cohesity.com https://helios-data.cohesity.com	443	TCP	SSL connection for Cohesity Support telemetry upload. Standard Helios traffic is sent to helios.cohesity.com. Query and API call traffic is sent to helios-data.cohesity.com.

DataHawk App

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	helios-data2.cohesity.com	443	TCP	Heartbeats and telemetry from the DataHawk app to Cohesity Helios.
Cohesity cluster	helios-dc3p.cohesity.com			Classification of heartbeats, telemetry, and results from the DataHawk app to Cohesity Helios.

MarketPlace Apps

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
HTTP client	Cohesity cluster	61001 through 63786	TCP	Used for Cohesity MarketPlace Apps, which are available from the MarketPlace and Helios. One port is dynamically assigned from this range for the user interface of each app instance.	Management

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
VCF Component	Cohesity cluster	63795	TCP	SFTP port required to transfer backup data.	Backup
Windows or Linux client				Port to access the data backup during recovery.	Recovery
HTTP client	Cohesity cluster	63800	TCP	Port used for SecureView UI	Management

Multitenancy

Hybrid Extender

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Hybrid Extender	DNS	53	TCP/UDP	DNS, if an external DNS server is configured.	Management
Client	Hybrid Extender	53		Serve DNS requests from an external source.	Management

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
SSH client	Hybrid Extender	22	TCP	Required for SSH server on Bifrost VM for Bifrost-RT	Management
Server	Hybrid Extender	445		Required for SQL servers.	Backup and Recovery
MS SQL server (any port)	Hybrid Extender	11113, 11117		VDI-based backup and restore of MS SQL database (using RemoteWAN SnapFS)	Backup and Recovery
MS SQL server (any port)	Hybrid Extender	11117		Required for log backup and restore of the MS SQL database.	Backup and Recovery
Hybrid Extender	Organization (Tenant)	29994		The internal port used on the organization's network for debugging and uploading the configuration file on the Hybrid Extender VM.	Management
NFS Clients	Hybrid Extender NFS target	2049		Required for NFS	Data Protection
SMB Clients	Hybrid Extender SMB target	445		SMB filer functionality. SMB and SMB2 restore.	Data Access
Linux VM	Hybrid Extender	111, 2049		File restore to Linux VM	Recovery
Hybrid Extender	Cohesity Cluster	11117		Required for Hybrid Extender with Oracle SBT	Backup and Recovery
ESXi host	Hybrid Extender	111, 2049		File restore to Windows VM	Recovery

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Hybrid Extender	NTP Server	123	UDP	NTP, if an external NTP server is configured.	Management
HTTP client	Hybrid Extender	80	TCP	HTTP	Management
Hybrid Extender	ESXi Host	902		Ensure that SSL communication for TCP port 902 to the ESXi host is enabled, otherwise, backups will fail.	Backup and Recovery
Hybrid Extender	Active Directory	137	UDP	User and Computer Authentication and NetLogon, NetBIOS Name Resolution. Required only when initially joining the Cluster to Active Directory.	Management
Hybrid Extender	Active Directory	139	TCP	Required only when initially joining the Cluster to Active Directory (for the NetBIOS session service).	Management
Hybrid Extender	Active Directory	445		Required only when initially joining the Cluster to Active Directory.	Management
Hybrid Extender	Kerberos Key Distribution Center (AD)	88	TCP/UDP	Required for Kerberos if the cluster is configured to use Active Directory.	Management
Hybrid Extender	LDAP	389		Required if the cluster is configured to use Active Directory or LDAP.	Management
	LDAPS	636			

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Hybrid Extender	DHCP Server	67, 68	TCP	DHCP. SMB/ICMP	Management
Hybrid Extender	VMware vCenter VMware Standalone ESXi Host VMware vCloud Director	443	TCP	The Cohesity cluster uses HTTPS by default. Required for Pure FlashBlade Required for HTTPS/HTTPS(TLS).	Backup and Recovery
Hybrid Extender	Cohesity cluster	29991	TCP	Used in a multi-tenant environment for interaction between the Hybrid Extender and Cohesity cluster.	Management
Hybrid Extender	Isilon	111	TCP/UDP	Required for RPC connection.	Backup and Recovery
		300, 302, 304, 2049		Required for NFS.	
		443	TCP	Required for HTTPS connection with Isilon.	
		445	TCP	Required for SMB.	
		8080	TCP/UDP	Required for HTTP connection.	
Generic NAS	Hybrid Extender	111	TCP/UDP	Required for RPC connection.	Backup and Recovery
		445	TCP	Required for SMB.	
		635, 2049	TCP/UDP	Required for NFS.	

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Mount Target	Hybrid Extender	111	TCP/UDP	NFS restores on Linux physical servers.	Backup and Recovery
Host	Hybrid Extender	50051*	TCP	Linux and Windows physical server backups. Use Node IP. If VLANs are configured, use VLAN VIPs. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Backup and Recovery

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Hybrid Extender	Cohesity cluster	29991	TCP	For the Oracle Adapter, need this port for communication, and NFS ports for mounting the Backup/Recovery Views.	Backup and Recovery.
Host	Hybrid Extender	50051*	TCP	Required for persistent agent workflow for file-level recovery and IVM. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.	Recovery

Data Source Connector

Ports Between Data Source Connector and Cohesity Cluster

Ensure that the following ports are open to allow communication between the Cohesity cluster and Data Source Connectors:

Source	Destination	Port	Protocol	Purpose
Data Source Connector	DNS	53	TCP/UDP	DNS, if an external DNS server is configured.
Data Source Connector	Cohesity cluster	443	TCP	Required for claiming and obtaining certificates for the secure grpc connection.
Data Source Connector	Cohesity cluster	29991	TCP	Required for Data path traffic.
Data Source Connector	Cohesity clustert)	11117	TCP	Required for Data path traffic.

Ports Between Data Source Connector and Sources

VMware

Ensure that the following ports are open to allow communication between the Data Source Connector(s) and VMware environment:

Source	Destination	Port	Protocol	Purpose
Data Source Connector	VMware vCenter	443	TCP	Required for making VMware API calls for backup and recovery over HTTPS/HTTPS (TLS).
Data Source Connector	VMware vCenter	443	TCP	Required for VMware Tools-based file and folder recoveries.
Data Source Connector	VMware vCenter	902	TCP	Data movement from VMware vCenter.

Physical Servers

Ensure that the following ports are open to allow communication between the Data Source Connector(s) and Physical Servers:

Source	Destination	Port	Protocol	Usage Notes
Data Source Connector	Physical Windows or Linux Server	50051*	TCP	Required for Backup and Recovery operations. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting

Source	Destination	Port	Protocol	Usage Notes
				unnecessary firewall rules from your network, firewall, or security teams.
Local Host (Physical Windows or Linux Server)	Local Host (Physical Windows or Linux Server)	59999	TCP	Required for local-to-local communication for self-monitoring and debugging purposes.

Agent Upgrade

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Client	Data Source Connector	80,443	TCP	For Agent upgrade from UI.

Oracle Servers

Ensure that the following ports are open to allow communication between the Data Source Connector(s) and Oracle Server:

Source	Destination	Port	Protocol	Usage Notes
Data Source Connector	Oracle Server	50051*	TCP	Required for Backup and Recovery operations. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting

Source	Destination	Port	Protocol	Usage Notes
				unnecessary firewall rules from your network, firewall, or security teams.
Oracle Server	Data Source Connector	111, 2049	TCP	Required for mount-based workflows.
Oracle Server	Data Source Connector	11113, 11117	TCP	Required for secure gRPC for Oracle SBT.
Local Host (Physical Windows or Linux Server)	Local Host (Physical Windows or Linux Server)	59999	TCP	Required for local-to-local communication for self-monitoring and debugging purposes.

Microsoft SQL Servers

Ensure that the following ports are open to allow communication between the Data Source Connector(s) and Microsoft SQL Server:

Source	Destination	Port	Protocol	Usage Notes
Data Source Connector	MS SQL Host	50051*	TCP	Required for Backup and Recovery operations. *By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both. Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting

Source	Destination	Port	Protocol	Usage Notes
				unnecessary firewall rules from your network, firewall, or security teams.
MS SQL Host	Data Source Connector	11113, 11117	TCP	Required for Backup and Recovery operations.
Data Source Connector	MS SQL Host where the CohesityAgent is running.	1433	TCP	Default TCP port for MS SQL instances. Ensure port is open to allow communication between the Data Source Connector and the MS SQL instance where the Cohesity Agent is running.

Cohesity Node Detection

Cohesity Node Detection

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes	Type of Traffic
Cohesity cluster node or free node	Cohesity cluster	5353	UDP	Used for node discovery during cluster creation and code addition.	Management

Cohesity Support

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Outgoing Traffic

> **Note:** Make sure your network team has unblocked and is not monitoring the Cohesity support channel sites or ports for seamless connectivity and effective communication.

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	support-connect-v1.cohesity.com	443	TCP	Used by the Cohesity cluster to initiate secure outbound connectivity to the Cohesity support channel server. Ensure to disable the SSL/TLS certificate inspection on the firewall.
	support-connect-trust-v1.cohesity.com	443	TCP	Used by the Cohesity cluster to establish and maintain a persistent secure reverse tunnel connection to the Cohesity support channel server. Ensure to disable the SSL/TLS certificate inspection on the firewall.
	Cohesity cluster	3022, 3023, 3024, 3025, 3080	TCP	The primary node of the Cohesity cluster connects to the Support Channel server, and the other nodes use these ports to communicate with the primary node. For more information, see Manage the Support Channel.

Internal Traffic

Internal Traffic Between Nodes

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Incoming Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	2222	TCP	As the support user, access the Host Shell via port 2222. Enabling Host Shell access requires generating and sharing the Support Token with Cohesity Support. For details, see Using the Secure Shell.
Cohesity cluster	Cohesity cluster	8081	TCP	SPIRE Server Service
Cohesity cluster	Cohesity cluster	9111	TCP	Serviceability Service
Cohesity cluster	Cohesity cluster	11112	TCP	I/O Operations Service (Low Latency)
Cohesity cluster	Cohesity cluster	11115	TCP	Archive Helper Service
Cohesity cluster	Cohesity cluster	11116	TCP	I/O Operations Helper Service
Cohesity cluster	Cohesity cluster	11121	TCP	For Fusion FS server
Cohesity cluster	Cohesity cluster	12222	TCP	Distributed Key-Value Store Service
Cohesity cluster	Cohesity cluster	12321	TCP	Distributed Key-Value Store Helper Service
Cohesity cluster	Cohesity cluster	20001	TCP	NAS Backup Helper Service
Cohesity cluster	Cohesity cluster	20002	TCP	SMB Backup Helper Service v1
Cohesity cluster	Cohesity cluster	20008	TCP	Bandwidth Throttling for Archive Services

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	21111	TCP	Alert Service
Cohesity cluster	Cohesity cluster	22222	TCP	Cluster Configuration Manager Service
Cohesity cluster	Cohesity cluster	22223	TCP	Cluster Configuration Manager Service
Cohesity cluster	Cohesity cluster	23458	TCP	Cluster Health Monitor Service
Cohesity cluster	Cohesity cluster	23460	TCP	Support Automation & Proactive Wellness Service
Cohesity cluster	Cohesity cluster	24567	TCP	UI and REST API Helper Service
Cohesity cluster	Cohesity cluster	24680	TCP	Healer and Space Reclamation Service
Cohesity cluster	Cohesity cluster	25566	TCP	Cluster Metrics Service
Cohesity cluster	Cohesity cluster	25588	TCP	To view the status of throttler master internally
Cohesity cluster	Cohesity cluster	25678	TCP	Hardware and Health Monitor Service
Cohesity cluster	Cohesity cluster	25679	TCP	Hardware and Health Monitor Helper Service
Cohesity cluster	Cohesity cluster	25680	TCP	Cluster Metrics Store Service
Cohesity cluster	Cohesity cluster	25687	TCP	etcd Client HTTPS Service

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	25690	TCP	Kubelet cAdvisor Service
Cohesity cluster	Cohesity cluster	25691	TCP	Kubelet HTTPS Service
Client	Cohesity cluster	25700	TCP	Indexing Service
Cohesity cluster	Cohesity cluster	25800	TCP	Indexing Service
Cohesity cluster	Cohesity cluster	25801	TCP/UDP	Analytics & Indexing Service
Cohesity cluster	Cohesity cluster	25805	TCP	etcd Client HTTP Service
Cohesity cluster	Cohesity cluster	25998	TCP	Marketplace Apps Service
Cohesity cluster	Cohesity cluster	25999	TCP	Indexing Helper Service
Cohesity cluster	Cohesity cluster	26000	TCP	Index Catalog Service v2
Cohesity cluster	Cohesity cluster	26001	TCP	Index Catalog Service v1
Cohesity cluster	Cohesity cluster	26440	TCP	Kubernetes Infra Management API Server
Cohesity cluster	Cohesity cluster	26443	TCP	Kubernetes API Server Load Balancer
Cohesity cluster	Cohesity cluster	26444	TCP	Kube Lifecycle Manager Service Peer Communication

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	26445	TCP	Kube Lifecycle Manager Service Client Communication
Cohesity cluster	Cohesity cluster	26447	TCP	Kubernetes API Server
Cohesity cluster	Cohesity cluster	26448	TCP	Kubelet API
Cohesity cluster	Cohesity cluster	26999	TCP	Reporting Service
Cohesity cluster	Cohesity cluster	27999	TCP	Reporting Service PostgreSQL
Cohesity cluster	Third-Party tools (for example: Tableau, SSRS etc.)		TCP	Used for connections to the Custom Reporting Database.
Cohesity cluster	Cohesity cluster	29999	TCP	Required for Archive Services
Cohesity cluster	Cohesity cluster	30000	TCP	Required for applying a patch.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	3022, 3023, 3024, 3025, 3080	TCP	For the Cohesity cluster to communicate with the Teleport support server. Make sure the Support Channel is enabled for the cluster (**Settings** > **Summary**).
Cohesity cluster	Cohesity cluster	22222	TCP	For Cohesity cluster Cloud Edition operations.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	22	TCP	Support Channel used by Cohesity Support for troubleshooting.
Cohesity cluster	Cohesity cluster	6641	TCP	OVN Northbound DB Service
Cohesity cluster	Cohesity cluster	6642	TCP	OVN Southbound DB Service
Cohesity cluster	Cohesity cluster	6643	TCP	OVN NB Raft Service
Cohesity cluster	Cohesity cluster	6644	TCP	OVN SB Raft Service
Cohesity cluster	Cohesity cluster	8285	UDP	Flannel Backend Service
Cohesity cluster	Cohesity cluster	8443	TCP	SPIRE Federation Service
Cohesity cluster	Cohesity cluster	8472	UDP	Flannel VXLAN Service
Cohesity cluster	Cohesity cluster	8473	UDP	Flannel2 VXLAN Service
Cohesity cluster	Cohesity cluster	11113	TCP	Source-side Deduplication Service
Cohesity cluster	Cohesity cluster	21000	TCP	gRPC Communication Service
Cohesity cluster	Cohesity cluster	22000	TCP	Key Management Service

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	23456	TCP	Cluster configuration and cluster management traffic
Cohesity cluster	Cohesity cluster	23464	TCP	
Cohesity cluster	Cohesity cluster	25677	TCP	Envoy Admin Service
Cohesity cluster	Cohesity cluster	25678	TCP	Service Identity Outbound Service
Cohesity cluster	Cohesity cluster	25679	TCP	Service Identity Inbound Service
Cohesity cluster	Cohesity cluster	25681	TCP	Marketplace Apps Service
Cohesity cluster	Cohesity cluster	25682	TCP	Docker Registry Service
Cohesity cluster	Cohesity cluster	25683	TCP	Optimus Service
Cohesity cluster	Cohesity cluster	25688	TCP	etcd Peer (Raft) Service
Cohesity cluster	Cohesity cluster	25689	TCP	Kube API Server HTTP Service
Cohesity cluster	Cohesity cluster	25692	TCP	CoreDNS Service
Cohesity cluster	Cohesity cluster	25693	TCP	CoreDNS Health Service
Cohesity cluster	Cohesity cluster	25694	TCP	Marketplace Apps Helper Service

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	25696	TCP	Kube API Server HTTPS Service
Cohesity cluster	Cohesity cluster	25697	TCP	Marketplace Apps Monitor Service
Cohesity cluster	Cohesity cluster	25698	TCP	ETL Server gRPC Service
Cohesity cluster	Cohesity cluster	25699	TCP	Sys Srvc Docker Registry Service
Cohesity cluster	Cohesity cluster	26338	TCP	OCI Registry Lifecycle Mgr Service
Cohesity cluster	Cohesity cluster	26339	TCP	OCI Registry API Service
Cohesity cluster	Cohesity cluster	26440	TCP	Bootstrap CAPI Service
Cohesity cluster	Cohesity cluster	26443	TCP	Kube API Cluster Endpoint Service
Cohesity cluster	Cohesity cluster	26444	TCP	Cluster Operator Service
Cohesity cluster	Cohesity cluster	26447	TCP	KLM API Server Service
Cohesity cluster	Cohesity cluster	28888	TCP	Cohesity CA Service
Cohesity cluster	Cohesity cluster	25589	TCP	HTTP Throttler Service via gRPC for centralized rate limiting and retry handling
Cohesity cluster	Cohesity cluster	28999	TCP	Reporting Service

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Cohesity cluster	29000 29001	TCP	Used for rolloutflag_evaluator service, which communicates with Helios to provide harness-based feature rollout updates.
Cohesity cluster	Cohesity cluster	29990	TCP	Hybrid Extender Service
Cohesity cluster	Cohesity cluster	64003	TCP	Cluster Fault Monitoring Service
Cohesity cluster	Cohesity cluster	64005	TCP	Cluster Resource Management Service

Cluster-to-Cluster Traffic

> **Note:** If the Cohesity cluster uses an interface with the node IP address to reach the external server or client, configure the firewall settings to allow the node IP addresses. If the Cohesity cluster uses an interface with only virtual IP addresses, configure the firewall settings to allow the virtual IP addresses. To identify whether the source IP is a Static IP or Virtual IP for outgoing traffic, see Identify Source IP of Outgoing Traffic.

Bidirectional Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Remote Cohesity cluster (target or receiving)	443	TCP	Required for replication pairing and ongoing secure communications between replication pairs.

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Remote Cohesity cluster (target or receiving)	11111	TCP	Required for I/O Service between all nodes and replication pairs when replication is enabled.
Cohesity cluster	Remote Cohesity cluster (target or receiving)	23335	TCP	If both the local and remote Cohesity clusters support TLS, port number 23335 is utilized for TLS communication during replication.
Cohesity cluster	Remote Cohesity cluster (target or receiving)	23336	TCP	If both the local and remote Cohesity clusters support TLS, port number 23336 is utilized for TLS communication during replication.
Cohesity cluster	Remote Cohesity cluster (target or receiving)	20000	TCP	• Required for replication. • Required for interaction handling between Cohesity and target backup systems.

Outgoing Traffic

Source	Destination	Destination Port	Protocol	Usage Notes
Cohesity cluster	Remote Cohesity cluster (target or receiving)	24444	TCP	Required for Health Check when replication is enabled on the Cohesity cluster

*By default, Cohesity uses port 50051. This port isn't reserved exclusively for Cohesity, so if another application on your network already uses it, you can use port 21213 instead. IANA reserves port 21213 for Cohesity in its service registry. You can select the port during installation, and only one port is active at a time. Open only the port that your deployment uses. You don't need to open both.

Port 21213 is an early access feature. To enable it, contact your Cohesity account team. Add port 21213/TCP to your firewall only after you enable this feature, to avoid requesting unnecessary firewall rules from your network, firewall, or security teams.

Cohesity CLI Commands Used to Manage Firewall Ports

For on-premises clusters, you can use the following Cohesity CLI commands to manage the cluster firewall configuration. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

The firewall commands use the following terminology: Active Cohesity rules are called *attachment*s. An attachment includes a *profile* and multiple *ipsets*. A profile is a group of TCP or UDP ports. An ipset is a subnet with the format of <IP/prefix>. An attachment can be on one or more specific network interfaces, or it can apply to all interfaces when no interface is specified.

CLI Command	Description
firewall-profile activate	Activates a profile on a list of interfaces with a list of ipsets.
firewall-profile add	Adds a collection of service ports with the protocol of each port for the firewall.
firewall-profile deactivate	Deactivates a list of active firewall profiles.
firewall-profile delete	Removes a list of firewall profiles.

CLI Command	Description
firewall-profile list	Lists the groups of firewall profiles.
firewall-profile list-active	Lists all the active firewall profiles. Each attachment activates a profile on specific network interfaces and ipsets.
firewall-ipset list	Lists the IP sets for a firewall.
firewall-profile activate revert-to-default	Reverts all the firewall rules to the default settings.

To list the parameters for each command, type `help` after the `firewall-profile <option>` or `firewall-ipset <option>`. For example:

```
admin@node-1> firewall-profile activate help
```

Access Management

Cohesity Access Management ensures secure and efficient access control by allowing only authorized users to interact with specific resources. It is designed to enforce User Authentication, Authorization, Access Control, and Security Settings, ensuring users have the necessary permissions to perform their roles.

The Users & Groups tab in the Access Management page lists all local users, Active Directory users, and SSO users and groups that can access the Cohesity cluster. Administrators can assign roles to users and groups to grant specific privilege levels. For detailed instructions, see Users and Groups. You can also edit or delete existing users.

Key Components of Access Management

User Authentication and Authorization

User and Group Management: Administrators can create and manage user accounts and groups, specifying roles and permissions that determine access levels to various Cohesity resources.

Role-Based Access Control (RBAC): Organization admins can add and manage users through role-based access controls (RBAC). Applying principles of least privilege and separation of duties is simple with fine-grained control over standard and custom defined roles.

Multi-Factor Authentication (MFA): Multi-factor Authentication (MFA) is an additional layer of security used to verify the identity of a user. Cohesity strongly recommends enabling Multi-factor Authentication for Local and Active Directory Users for all users and groups for enhanced security. After the administrator enables MFA, users are prompted for

an additional authentication factor along with the Cohesity cluster user credentials during the sign-in process.

Single Sign-On (SSO): Streamlines the login process by allowing access to all the Cohesity resources with a single set of credentials, reducing the risk of password fatigue and enhancing user convenience. For more information, see Single Sign-On (SSO).

Users and Groups

Cohesity allows adding users with different privileges by assigning them various roles within the cluster. These users and groups can manage data and view statistics based on their assigned roles.

Cohesity supports the following types of users and groups. Except for the **support** user, all user types can log in to the Cohesity Dashboard to manage data on the Cohesity cluster and view statistics, depending on their role.

- **Local users:** Local users exist on the Cohesity cluster only.
- **Local groups:** You can define groups by adding local users and Active Directory principals.
- **SSO users:** If the Cohesity cluster is configured to use an identity provider for single sign-on access, you can add SSO users and assign them roles. For more information, see SSO Users.
- **Active Directory principals:** If the Cohesity cluster is joined to an AD domain, you can assign roles to principals in the AD or the AD's trusted domains, which determine the Cohesity cluster access. For more information, see Active Directory (AD).
- **Support user:** Used to access the bash shell of the Cohesity cluster using SSH. The support user account is disabled by default. For more information, see Manage Support User.

Types of Users and Groups

Local Users

Local users are exclusive to the Cohesity cluster. They can be created and managed directly within the system.

Add or Edit a Local User

1. Navigate to **Settings > Access Management**.

   In the **Users & Groups** tab, from the **Add AD Users & Groups** drop-down list, select **Add Local User**, or click the action menu and select **Edit** to edit a local user.

2. Enter the following information:

Add Local User	Enter the user name. Do not enter whitespace characters such as spaces. Do not specify any of these characters: `/ < > : #`
Email Address	Enter the email address associated with this user.
Password	Enter a password that is at least eight characters.
Confirm Password	Enter the same password again to confirm.
Groups	Add the user to one or more user groups by selecting the groups from the drop-down list
Roles	(Optional) Select the appropriate roles. For a detailed description of each role, see Manage Roles.
Effective Date	Specify the date that the user can begin using the account.
Expiration Date	Specify the date when the user's access will be revoked.

Restrict access to specific Objects	Optionally restrict access to specific Objects and Views only by toggling this option on. By default, the allowed actions apply to all Source Objects registered with the Cohesity cluster and all Views. After restricting access, you must specifically add Objects and Views through **Add Registered Object(s)** and **Add Cohesity View(s)**. The user or group does not have access to any Objects or Views until you add them. When a user has access to an Object, that user also has access to all previous Snapshots of that Object and any Objects added to the selected Source folder in the future. For example, if a user has access to an ESX host, they will also have access to VMs added to the host later. Users that have access restricted to specific Objects cannot modify users or Sources even if their role grants them these privileges. Object permissions are replicated for Virtual Machine Objects but not currently for physical Server and View Objects. The Object permissions for a local user or group are not replicated to the remote Cohesity cluster. Local users inherit the permissions of the Restricted Objects provided to the Restricted Groups to which the users belong. Users that have access restricted to specific objects cannot update the Protection Group created by an admin user even if the restricted user has access to all the objects in the Protection Group.
Set Primary User Group for File Access	Enable this option for the user to access file services. From the **Primary Group** drop-down list, select a group as a Primary group. The user can be a part of one or more groups and the SMB/NFS protocol requires a primary group to access the files.

3. Click **Add** or **Save Changes** if editing a user.

Local Groups

Local groups consist of local users and Active Directory (AD) principals. These groups can be defined and managed within the Cohesity cluster.

Add or Edit a Local Group

1. Navigate to **Settings > Access Management**.
2. In the **Users & Groups** tab, from the **Add Users & Groups** drop-down list, select **Add Local Group**, or click the action menu and select **Edit** for a local group.

3. Enter the following information:

Group Name	Enter the group name. Do not enter whitespace characters such as spaces. Only letters, numbers and the special characters given below are allowed: _ - .
Local User	Add one or more users to the group by selecting the users from the drop-down list.
Roles	Select the appropriate roles from the drop-down list. For more information on roles, see Manage Roles.
Effective date	Specify the date when the user group can begin using the account
Restrict access to specific Object	Optionally restrict access to specific Objects and Views only by toggling this option on. By default, the allowed actions apply to all Source Objects registered with the Cohesity cluster and all Views. After restricting access, you must specifically add Objects and Views through **Add Registered Object(s)** and **Add Cohesity View(s)**. The user or group does not have access to any Objects or Views until you add them. When a user has access to an Object, that user also has access to all previous Snapshots of that Object and any Objects added to the selected Source folder in the future. For example, if a user has access to an ESX host, they will also have access to VMs added to the host later. Users that have access restricted to specific Objects cannot modify users or Sources even if their role grants them these privileges. Object permissions are replicated for Virtual Machine Objects but not currently for physical Server and View Objects.
Set Primary User Group for File Access	Enable this option for the user to access file services. From the **Primary Group** drop-down list, select a group as a Primary group. The user can be a part of one or more groups and the SMB/NFS protocol requires a primary group to access the files.

4. Click **Add**, or **Save Changes** if editing a group.

Add a Local Administrator Group Using the CLI

You can use local administrator groups to grant SMB access to views and/or administrative access to the Cohesity cluster.

The Cohesity cluster has a built-in local administrators group called Administrators. By default, the group has only SMB Security privileges and the standard Windows administrator Security ID (SID) S-1-5-32-544. You can use and modify the Administrators group, or create your own local administrator groups and assign roles and SIDs by using the Cohesity CLI.

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Use the following command syntax to add a local administrator group.

   ```
   group add name=<group-name> smb-only=true roles=<cohesity_roles>
   ```

   (You can use the `role list` command to list the roles defined in the Cohesity cluster.)

   For example:

   ```
   admin@node1> group add name=admin2 smb-only=true roles=cohesity_admin
   GROUP NAME                    : admin2
   DOMAIN                        : LOCAL
   ROLES                         : COHESITY_ADMIN
   SMB PRINCIPAL ENABLED         : true
   ```

Use the following command syntax to add a user or group.

```
group add-smb-principal name=<group_name> type=<user or group> domain=<AD_
domain> principal-name=<user or group> sid=<SID>
```

For example:

```
admin@node1> group add-smb-principal name=admin2 type=group
domain=myorg.ad.domain.com principal-name=PIMS sid=S-1-5-21-4152169185-
4221176486-4244366766-1001
GROUP NAME                    : admin2
DOMAIN                        : LOCAL
ROLES                         : COHESITY_ADMIN
SMB PRINCIPAL ENABLED         : true
SMB PRINCIPAL                 :
NAME                          : PIMS
DOMAIN NAME                   : myorg.ad.domain.com
SID                           : S-1-5-21-4152169185-4221176486-4244366766-
1001
TYPE                          : group
```

> **Tip:** Use the CLI Help by running the following command: `group add-smb-principal help`.

CLI Option	Example	Description
`domain=`	`org.ad.domain.com`	Optional. The AD domain name that includes the principal you are adding.
`name=`	`admin2`	Required. Your choice of a name for the local administrator group.
`principal-name=`	`group1` `jsmith`	Optional. The name of the SMB principal user or group. The name must exist in the AD domain.
`sid=`	`S-1-100-21-8372`	Optional. A unique Security ID for the principal.
`type=`	`group` or `user`	Optional. Specify `user` or `group` to indicate the type of principal you are adding.

Enable User Account Lockout

The admin user can enable account lockout on the user account. User accounts can get locked after a threshold number of failed login attempts. The user cannot log in to the Cohesity Dashboard after three consecutive invalid login attempts for 15 minutes.

> **Note:** You can enable account lockout only for non-admin local users.

Perform the following steps to enable user account lockout:

1. Navigate to **Settings > Access Management**.
2. Select the **Users & Groups** tab.
3. Select the user, click the actions menu (⋮) and select **Lock**.

Disable User Account Lockout

The admin user can disable the account lockout on a user account.

Perform the following steps to disable user account lockout:

1. Navigate to **Settings > Access Management**.
2. Select the **Users & Groups** tab.
3. Select the user, click the actions menu (⋮) and select **Unlock**.

Delete Users and Groups

You can delete users and groups.

> **Note:** You cannot delete default users and groups.

Perform the following steps to delete user and group:

1. Navigate to **Settings > Access Management**.
2. Select the **Users & Groups** tab.
3. Select the user, click the actions menu (⋮) and select **Delete**.
4. Click **Delete**.

Change Local User Password

Using the Cohesity Dashboard, you can change your local user password. By default, the new password must replace at least eight characters with the current password.

Using the following Cohesity CLI command, you can define the number of characters you have to replace in your existing local user password:

```
iris_cli security-config update num-diff-chars=<number_of_characters>
```

For example:

If you want to restrict those six characters in your existing local user password must be replaced while changing the password, then use the following command:

```
iris_cli security-config update num-diff-chars=6
```

> **Note:** The restriction is not applicable if the admin user wants to change the password of other local users in the Cohesity cluster.

To change the local user password:

1. Navigate to **Settings** > **Access Management**.
2. Click the actions menu (⋮) next to the user and click **Edit**.
3. In the **Edit Local User: <user>** page:
    1. Enable **Change Password**.
    2. Enter the **Current Password** and then enter and confirm the new **Password**.
    3. Click **Save Changes**.

SSO Users

Single sign-on (SSO) users are added if the Cohesity cluster is configured to use an identity provider for SSO access. These users can be assigned roles within the cluster.

Add SSO Users

You can add single sign-on (SSO) users only if the Cohesity cluster is configured to use an identity provider. For more information, see Single Sign-On (SSO).

> **Note:**
> - SSO users can authenticate to the cluster either through the authorization of groups, default roles, or roles assigned to them. But they cannot access the apps on the cluster since their identity cannot be verified by the IDP.
> - When a user tries to access the app, a token is obtained for that user.
> - The identity can be verified only for LOCAL and AD domain users. It cannot be done for SSO users since IDP supports only SAML. For more information on how to make the apps accessible to SSO users, see the Giving App access to SSO users KB article.

1. Navigate to **Settings > Access Management**.
2. In the **Users & Groups** tab, from the **Add AD Users & Groups** drop-down list, select **Add SSO Users & Groups**, or click the action menu and select **Edit** for an SSO user.

3. Enter the following information:

SSO Domain	Select the SSO domain from the drop-down list.
SSO Users	Type the SSO user names.
SSO Groups	Type the SSO group names.
Roles	Select the appropriate roles. For a detailed description of each role, see Manage Roles.
Description	Optionally add a description.
Restrict access to specific Objects	Optionally restrict access to specific Objects and Views only by toggling this option on. By default, the allowed actions apply to all Source Objects registered with the Cohesity cluster and all Views. After restricting access, you must specifically add Objects and Views through **Add Registered Object(s)** and **Add Cohesity View(s)**. The user or group does not have access to any Objects or Views until you add them. When a user has access to an Object, that user also has access to all previous Snapshots of that Object and any Objects added to the selected Source folder in the future. For example, if a user has access to an ESX host, they will also have access to VMs added to the host later. Users that have access restricted to specific Objects cannot modify users or Sources even if their role grants them these privileges. Object permissions are replicated for Virtual Machine Objects but not currently for physical Server and View Objects. The Object permissions for a local user or group are not replicated to the remote Cohesity cluster.

4. Click **Add**, or **Save Changes** if editing a group.

Active Directory Users and Groups

An Active Directory User is an account that is stored and managed within an Active Directory domain. This user can be authenticated and authorized to access various resources within the domain, including the Cohesity cluster.

Add Active Directory Users and Groups

1. Navigate to **Settings > Access Management**.
2. Click **Add AD Users & Groups** and select **Active Directory Users & Groups**.
3. Enter the following information:

Active Directory Domain	Select the Active Directory domain or a trusted domain that contains the users and groups to assign roles to.
User(s) or Group(s)	Select the AD users or groups to assign roles to. .
Roles	Select the appropriate roles. For a detailed description of each role, see Manage Roles.
Restrict access to specific Objects	Optionally restrict access to specific Objects and Views only by toggling this option on. By default, the allowed actions apply to all Source Objects registered with the Cohesity cluster and all Views. After restricting access, you must specifically add Objects and Views through **Add Registered Object(s)** and **Add Cohesity View(s)**. The user or group does not have access to any Objects or Views until you add them. When a user has access to an Object, that user also has access to all previous Snapshots of that Object and any Objects added to the selected Source folder in the future. For example, if a user has access to an ESX host, they will also have access to VMs added to the host later. Users that have access restricted to specific Objects cannot modify users or Sources even if their role grants them these privileges. Object permissions are replicated for Virtual Machine Objects but not currently for physical Server and View Objects. When Snapshots are replicated from a local to a remote Cohesity cluster, the Object permissions for AD users or groups are also replicated. However, the privileges for an AD user or group are not replicated. For example, if an AD user called Fred has permission to access the Virtual Machines called VM1, VM2 and VM3 on the local Cohesity cluster, when Snapshots of those Virtual Machines are replicated to a remote Cohesity cluster, those Object permissions are also replicated. However, an administrator on the remote Cohesity cluster must add the AD User called Fred or an AD group that includes Fred and grant him privileges so Fred can access and recover from Snapshots of these VMs on the remote Cohesity cluster. AD users inherit the permissions of the Restricted Objects provided to the Restricted AD Groups to which the users belong.
Description	Optionally add a description.

Role and privilege updates for currently logged-in users take effect after logging out and logging in again. Updates for users that are not currently logged in take effect when they log in.

Support User

The support user account is designed for accessing the bash shell of the Cohesity cluster through SSH for troubleshooting and support purposes. This account is disabled by default for security reasons.

Manage Support User

Cohesity provides a support user account for improved security and you need to use the support user account to log in to the Cohesity cluster bash shell using SSH.

By default, the support user account is disabled on the Cohesity node that is part of a Cohesity cluster. You need to enable the support user account by setting a password for the user account in the Cohesity Dashboard. Also, you can Enable Linux Sudo Access for the support user account.

Set the Support User Password for the First Time

Using the Cohesity Dashboard or Cohesity CLI, you can set the password for the support user. You cannot retrieve the password once it is set. Therefore, ensure that you set a strong password that is longer and easy to remember. In case, if you want to reset the support user password and have forgotten the existing password then contact Cohesity Support.

Perform the following steps to set the support user password:

1. In the Cohesity Dashboard, select **Settings > Access Management** and click the **Support** tab.
2. Under **Linux Shell Password**, click **Set Password** and perform the following steps:
    a. In the **New Password** field, enter the new password. The password requirements are:
        - Must be at least 15 characters long.
        - Must contain at least one digit, one special character, one uppercase & one lowercase character.

          > **Note:** Do not use `!`, `` ` ``, `"` and `\`.

        - Do not use more than two of the same characters in a row. For example, 123aa is allowed. 123aaa is not allowed.
        - Do not use more than four identical consecutive character types (digit, special character, uppercase and lowercase character). For example, 123aa is allowed. 12345aa is not allowed.
    b. In the **Confirm Password** field, enter the new password.
    c. Click **Set**.

> **Note:** By default, the support user logs in to the secure shell and access to the host shell is restricted through SSH. For more information on how to enable host shell access using the support user account, see Enable Host Shell Access.

Add Support User Email Address

Perform the following steps to add the support user email address:

1. In the Cohesity Dashboard, select **Settings > Access Management**.
2. Click the **Support** tab.

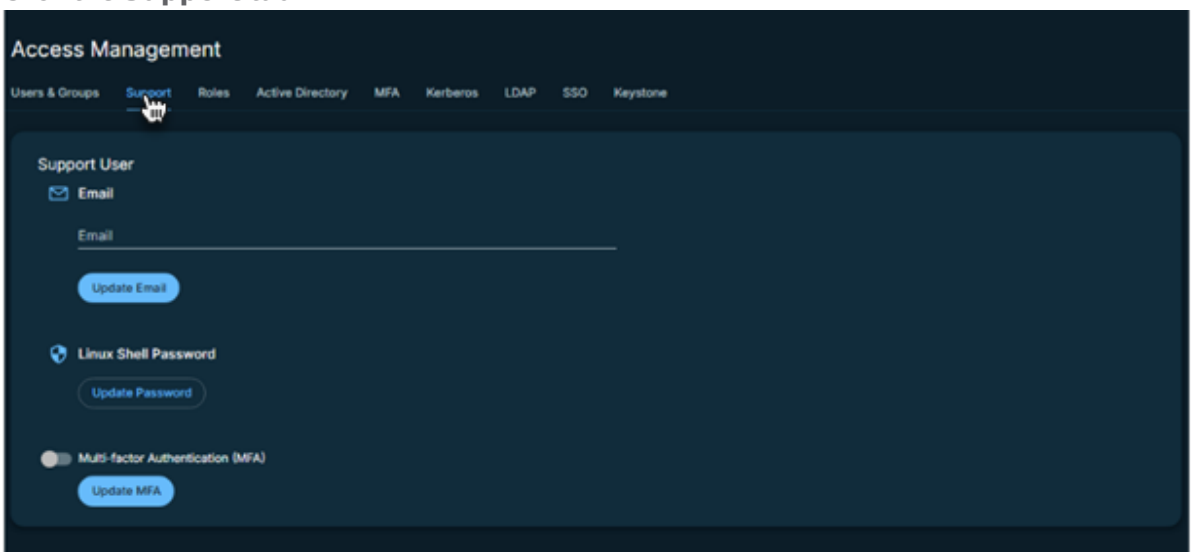

3. In the **Email** field, enter the email address and click **Update Email**.
4. In the **Update Email** pop-up window that appears, enter the support user password and click **Authorize**.

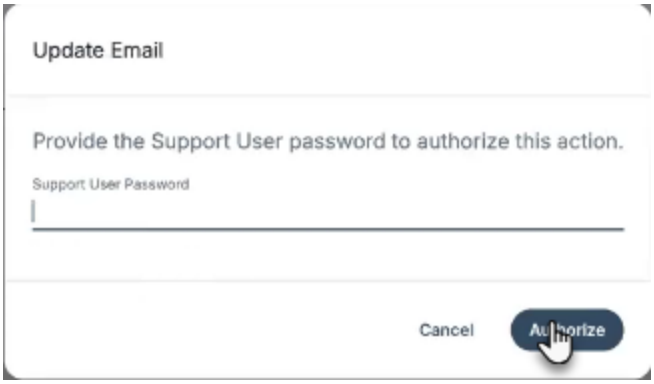

Change the Support User Password

Perform the following steps to change the support user password:

1. In the Cohesity Dashboard, select **Settings > Access Management**.
2. Click the **Support** tab.
3. Under **Linux Shell Password**, click **Update Password**.
4. Specify the current and the new password for the support user account and click **Update**.

In case, if you have forgotten the current password, contact Cohesity Support to reset the password.

Also, from the Cohesity CLI, you can use the following iris_cli command to change the support user password:

```
iris_cli user linux-user current-password=<current-password> linux-user-name=support linux-password=<new-password>
```

Enable or Disable Multi-factor Authentication

Multi-factor Authentication (MFA) is an additional layer of security used to verify the identity of a user. Cohesity supports MFA for the support user account. MFA enforces users to provide an additional authentication factor to access the Cohesity Dashboard or Cohesity CLI. For more information, see Multi-factor Authentication for the Support User Account.

Modify User Settings

Cohesity allows you to modify user settings such as email, password, and language preferences.

You can perform the following steps to edit the user settings:

1. In the Cohesity Dashboard, click the user icon ( ) > **User Settings**.
2. Enter the following information:

Field	Description
Email	Enter the email address associated with this user.
Change Password	Toggle on the setting to change the password.
Current Password	Enter the current password.
Password	Enter a new password that is at least eight characters.
Confirm Password	Enter the new password again to confirm.
Roles	Select the appropriate roles. For a detailed description of each role, see Manage Roles.
Description	Optionally enter a description

Field	Description
Language	Select the preferred language. The available options are: • English (US) • Japanese (JP)
Restrict access to specific Objects	Optionally restrict access to specific Objects and Views only by toggling this option on. By default, the allowed actions apply to all Source Objects registered with the Cohesity cluster and all Views. Object permissions are replicated for Virtual Machine Objects but not currently for physical Server and View Objects. The Object permissions for a local user or group are not replicated to the remote Cohesity cluster. Users that have access restricted to specific Objects cannot modify users or Sources even if their role grants them these privileges. Local users inherit the permissions of the Restricted Objects provided to the Restricted Groups to which the users belong.
Add Registered Object(s)	Click this option to add Objects. The user or group does not have access to any Objects until you add them. When a user has access to an Object, that user also has access to all previous Snapshots of that Object and any Objects added to the selected Source folder in the future. For example, if a user has access to an ESX host, they will also have access to VMs added to the host later
Add Cohesity View(s)	Click this option to add Views. The user or group does not have access to any Views until you add them.
Set Primary User Group for File Access	Enable this option for the user to access file services. Then from the **Primary Group** drop-down list (lists all the groups the user belongs to), select a primary group to access files using SMB/NFS protocol.
Allow SMB Access Tokens	Enable this option to grant the local user the necessary permissions to authenticate SMB mounts during the recovery workflow. > **Note:** This setting can be configured only at the time of user creation and cannot be modified later.

3. Click **Save Changes**.

Manage Roles

The Manage Roles page lists all defined roles that you can assign to users and groups. From the Manage Roles page you can add and edit custom roles. For detailed instructions, see Create or Edit a Role. You can also delete existing custom roles. For users with restricted access, their allowed actions can be performed on the explicitly selected object set only; they do not have access to other objects. This restricted access protects your data from any ransomware attack because the data is managed only by the users who have the necessary privileges on your cluster.

> **Note:** Admins can create and modify custom roles. Pre-defined system roles cannot be modified.

Pre-defined System Roles

Admin

Users with the Admin role have full access to all actions and workflow within the Cohesity UI and Cohesity CLI except for changing DataLock Views and setting their expiration dates. For example, a user assigned the Admin role can create backups, recover Snapshots and files from backups, add sources, delete sources, set up and manage the Cohesity cluster, manage Nodes, manage Views, resolve Alerts, etc. A user with the Admin role can also add and delete users.

By default, the Cohesity cluster is preconfigured with a default System Admin user called `admin` that has the same privileges as a user with the Admin role. This default System Admin account (`admin`) cannot be deleted.

Viewer

Users with the Viewer role have read-only access for all workflows within the Cohesity Dashboard.

Operator

Users with the Operator role have Viewer role privileges and the following additional privileges:

- Run existing Protection Groups.
- Create and Manage Clone tasks.
- Create and Manage Recover tasks.

A user with this privilege cannot download files from a snapshot and cannot recover from External targets.

SMB Security

Users with the SMB Security role have all SMB privileges when accessing data over SMB.

SMB Backup Operator

Users with SMB Backup Operator role have the privilege to perform SMB backup and SMB restore.

Self Service Data Protection

Users with the Self Service Data Protection role have Viewer role privileges and the following additional privileges:

- Create and Manage Protection Policies.
- Create and Manage Protection Groups.
- Create and Manage Clone tasks.
- Create and Manage Recover tasks
- Create and Manage Cohesity Views.

Data Security

Users with the Data Security role have all privileges granted by the Self Service Data Protection and Viewer roles and the following additional privileges:

- Create DataLock Views and set DataLock expiration dates.
- Create a Protection Group.
- Start an on-demand protection run and edit a protection run.
- Delete protection runs and snapshots.
- Add and remove Legal Hold to existing snapshots.

Replication

Users with the Replication role have basic privileges for replication. This role allows the user to perform replication of data to other clusters.

A user with this privilege cannot log in to the Cohesity Dashboard.

FortKnox Self Managed Vault Admin

FortKnox Self-Managed vault admin users have full access to create, update, and delete FortKnox Self-Managed cluster pairing.

Create or Edit a Role

Use this page to create or edit a custom role. For example, you can create a user with the Operator role and certain admin privileges such as Upgrade Cluster and Manage Cluster. All custom roles have Viewer role privileges.

> **Note:** Admins can create and modify custom roles. Pre-defined system roles cannot be modified.

Create or Edit a Custom Role

1. Name the role. Role names cannot be edited after you create them.

   > **Note:** Only letters, numbers, spaces and the special characters . - _ are allowed.

2. You can optionally provide a description for the role.
3. Set the role's privileges by using the toggles and check boxes. For a description of the privileges, see Privileges.

   > **Tip:** In the Cohesity Dashboard, hover your mouse over the privilege name to display its description in a pop-up tooltip.

4. Click **Create Role** or **Save Changes**.

   Role and privilege updates for currently logged-in users take effect after logging out and logging in again. Updates for users that are not currently logged in take effect when they log in.

Privileges

The privileges for all workflows are described below.

> **Tip:** To select privileges in the Source Management workflow, you must select the Manage Sources privilege in the Data Protection workflow.

Access Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
View API Key	Allow user to view other users' API Keys.	YES	NO	YES	NO
View Users	Allows viewing users, groups and other access control information.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage API Key	Allow user to modify other users' API Keys.	YES	NO	YES	NO
Manage Users	Allows modifying users, groups and other access control information. If the user's access is restricted to specific Objects, then this privilege is not granted, even if the role grants it.	YES	NO	YES	NO
Manage S3 Keys	Allow user to manage and view other users' S3 keys.	YES	NO	YES	NO

Analytics Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
View Analytics Workbench	Allows the user to view Analytics Workbench related entities.	NO	NO	YES	NO
Reporting	Allows the user to generate and view reports.	YES	NO	YES	NO
Manage Analytics Workbench	Allows the user to modify Analytics Workbench related entities.	NO	NO	YES	NO

Apps

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Launch Apps Instances	Allows viewing apps, instances and app management information. For an app instance, access is restricted to the user who launched the app instance.	NO	NO	YES	NO
Manage Apps and Instances	Allows install, uninstall and launch of the app, and pause, resume and stop of the app instance. User specific restrictions of app instance do not apply.	NO	NO	YES	NO

Clone Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
View Clone Tasks	Allows the user to view clone tasks.	YES	NO	YES	NO
Manage Clone Tasks	Allows the user to create, modify and delete clone tasks.	YES	NO	YES	NO

Cluster Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Download Hybrid Extender	Allows the user to download Hybrid Extender VM or configuration file.	YES	YES	YES	YES
View report Schedules	Allows the user to view report schedules on cluster. The user cannot make changes.	YES	NO	YES	NO
View Cluster Details	Allows the user to view Cluster settings. The user cannot make changes.	NO	NO	YES	NO
Manage report schedules	Allows the user to create or modify schedules for a report.	YES	NO	YES	NO
Manage External Targets	Allows the user to register and modify External Targets for archival.	NO	NO	YES	NO
View Alert Details	Allows the user to view alerts. The user cannot make changes.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Allow access to Cohesity UI	Allows the user to access Cohesity cluster through UI.	YES	NO	YES	NO
Manage MFA	Allow users to modify cluster level MFA settings.	NO	NO	YES	NO
Manage AD and LDAP	Allows the user to manage active directory and LDAP providers on the cluster.	YES	NO	YES	NO
Manage Keystone	Allows the user to manage Keystone configuration on the cluster.	YES	NO	YES	NO
Upgrade Cluster	Allows the user to upgrade and patch the Cluster.	NO	NO	YES	NO
View Hybrid Extender Details	Allows the user to view the Hybrid Extender details.	YES	YES	YES	YES
Manage Alerts	Allows the user to resolve alerts and manage email addresses where alert notifications should be received.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage VLANs	Allows the user to modify VLANs.	NO	NO	YES	NO
View AD and LDAP Details	Allows the user to view active directory and LDAP providers. The user cannot make changes.	YES	NO	YES	NO
Manage Linux user sudo access	Allows the user to grant sudo access to Linux support user.	YES	NO	YES	NO
View NIS	Allows the user to view NIS providers on the cluster.	NO	NO	YES	NO
Cluster Support	Allows the user to perform support related operations such as collecting data and browsing logs.	NO	NO	YES	NO
View External Targets	Allows the user to view External Targets for archival.	YES	NO	YES	NO
Manage Helios	Manage connection to Helios.	NO	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
View Remote Clusters	Allows the user to view the Remote Clusters for remote access or replication.	YES	NO	YES	NO
View VLANs	Allows the user to view available VLANs.	YES	NO	YES	NO
Manage Remote Clusters	Allows the user to register and modify Remote Clusters for replication.	YES	NO	YES	NO
View Audit Logs	Allows the user to view the Audit Logs.	YES	NO	YES	NO
Manage Patches	This privilege allows a User to apply or revert patches to the Cluster.	NO	NO	YES	NO
Manage Cluster	Allows the user to modify Cluster setup and to add and remove Nodes.	NO	NO	YES	NO
Manage NIS	Allows the user to manage NIS providers on the cluster.	NO	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
View Keystone Details	Allows the user to view Keystone configuration. The user cannot make changes.	YES	NO	YES	NO
Modify Bifrost VLANs	Allows a user to modify available Hybrid Extender VLANs.	YES	NO	YES	NO
View Bifrost VLANs	Allows the user to view available Bifrost VLANs.	YES	NO	YES	NO
View KERBEROS	Allows the user to view registered Kerberos providers on the cluster.	NO	NO	YES	NO
Manage KERBEROS	Allows the user to manage registered Kerberos providers on the cluster.	NO	NO	YES	NO
Manage Tags	Allows the user to apply/ remove tags on the cluster.	YES	NO	YES	NO
View Tags	Allows the user to view tags on the cluster.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Modify Tags	Allows the user to create/ modify tags on the cluster.	YES	NO	YES	NO
Manage Support Channel	Allows the user to modify Support Channel settings on the cluster.	NO	NO	YES	NO
View Account Security	Allows the users to view the Account Security settings on the cluster.	NO	NO	YES	NO
View data-source connections	Allows a user to view data-source connections.	NO	NO	NO	NO
Modify data-source connections	Allows a user to create or modify data-source connections.	NO	NO	NO	NO
View data-source connectors	Allows a user to view data-source connectors.	NO	NO	NO	NO
Modify data-source connectors	Allows a user to create or modify data-source connectors.	NO	NO	NO	NO
View Bifrost Connections	Allows a user to have view the available Bifrost Connections.	YES	YES	YES	YES

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Modify Bifrost Connections	Allows a user to create or modify Bifrost Connections.	YES	YES	YES	YES

Data Protection

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Protection Group Operator	Allows the user to run, cancel or pause a protection group.	YES	NO	YES	NO
View Protection Groups	Allows the user to view protection groups, protection sources and protection group runs. The user cannot make any changes.	YES	NO	YES	NO
Manage Protection Groups	Allows the user to create, modify and delete Protection Groups as well as the ability to delete runs.	YES	NO	YES	NO
View Protection Policies	Allows the user to view protection policies. The user cannot make any changes.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage Protection Policies	Allows the user to create, modify and delete Protection Policies.	NO	NO	YES	NO
Search Objects	Allows the user to perform search of objects from sources registered on the cluster.	NO	NO	YES	NO
View Runbooks	Allows the user to list and view the created Runbooks.	NO	NO	YES	NO
Manage Runbooks	Allows the user to create, modify, remove and execute Runbooks	NO	NO	YES	NO
Manage Sources	Allows the user to register or delete a Protection Source and modify its information. If the user's access is restricted to specific Objects, then this privilege is not granted, even if the role grants it.	YES	NO	YES	NO
Manage Agents	Allows the user to Deploy and Register the Cohesity Agent on servers.	NO	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Execute Runbooks	Allows the user to execute Runbooks.	NO	NO	YES	NO
Delete snapshots	Allows user to delete snapshots.	YES	NO	YES	NO
View agent upgrade tasks	Allows the user to access agent upgrade tasks	YES	NO	YES	NO
Create or Modify agent upgrade tasks	Allows the user to create or modify agent upgrade tasks	YES	NO	YES	NO

Free Node Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
View Free Nodes	Allows the user see information about free nodes.	NO	NO	YES	NO

Organization Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Switch Organization	Allows to switch to an organization's perspective and perform actions on their behalf.	NO	YES	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
View Organizations	Allows viewing organizations registered on the cluster.	NO	YES	YES	NO
Manage Organizations	Allows managing organizations registered on the cluster.	NO	YES	YES	NO

Recovery Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Download File	Allows the user to download files. The user also needs the 'Manage Recover Tasks' privilege to download files.	YES	NO	NO	NO
View Recover Tasks	Allows the user to view Recover and Clone Tasks. The user cannot make any changes.	YES	NO	YES	NO
Manage Recover Tasks	Allows the user to create Recover Tasks.	YES	NO	YES	NO
Recover from External Targets	Allows the user to search External Targets and recover using data archived from a Remote Cluster.	NO	NO	YES	NO

SMB Security

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
SMB Backup	This privilege causes the system to grant all read access control to any file, regardless of the access control list (ACL) specified for the file. Any access request other than read is still evaluated with the ACL.	YES	NO	YES	NO
Take ownership of SMB files	This privilege allows the owner value to be set only to those values that the holder may legitimately assign as the owner of an object.	YES	NO	YES	NO
Manage system access control list on SMB files	Required to perform a number of security-related functions, such as controlling and viewing audit messages. This privilege identifies its holder as a security operator.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
SMB Restore	This privilege causes the system to grant all write access control to any file, regardless of the ACL specified for the file. Any access request other than write is still evaluated with the ACL. Additionally, this privilege enables you to set any valid user or group SID as the owner of a file.	YES	NO	YES	NO

Source Access Control

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage Enterprise File DataLock	Allows the user to decrease the Enterprise File DataLock expiration period and to delete a view with Enterprise File DataLock configured.	YES	NO	YES	NO
Data Security	Allows the user to lock a View and set its lock expiration date.	YES	NO	YES	NO

Source Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage VMware Standalone Hosts	Allows a user to manage VMware Standalone Hosts.	YES	NO	YES	NO
Manage VMware VCenters	Allows a User to manage all VMware vCenter Servers.	YES	NO	YES	NO
Manage VMware VCloud Director	Allows a User to manage all VMware vCloudDirector.	YES	NO	YES	NO
Manage HyperV SCVMM Servers	Allows a User to manage HyperV SCVMM Servers.	YES	NO	YES	NO
Manage HyperV Standalone Hosts	Allows a User to manage HyperV Standalone Hosts.	YES	NO	YES	NO
Manage HyperV Failover Clusters	Allows a User to manage HyperV Failover Clusters.	YES	NO	YES	NO
Manage Acropolis Standalone Clusters	Allows a User to manage Acropolis Standalone Clusters.	YES	NO	YES	NO
Manage GCP IAM users	Allows a User to manage GCP IAM users.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage AWS IAM users	Allows a User to manage AWS IAM users.	YES	NO	YES	NO
Manage Azure Subscriptions	Allows a User to manage Azure Subscriptions.	YES	NO	YES	NO
Manage KVM OVirtualManager Sources	Allows a User to manage KVM OVirtualManager Sources.	YES	NO	YES	NO
Manage Physical Servers	Allows a User to manage Physical Servers.	YES	NO	YES	NO
Manage Office365 Domains	Allows a User to manage Office365 Domains.	YES	NO	YES	NO
Manage SQL Servers	Allows a User to manage SQL Servers.	YES	NO	YES	NO
Manage Oracle Servers	Allows a User to manage Oracle Servers.	YES	NO	YES	NO
Manage Pure Storage Arrays	Allows a User to manage Pure Storage Arrays.	YES	NO	YES	NO
Manage Nimble Storage Arrays	Allows a User to manage Nimble Storage Arrays.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage Hyperflex Servers	Allows a User to manage Hyperflex Servers.	YES	NO	YES	NO
Manage Active Directory Servers	Allows a User to manage Active Directory Servers.	YES	NO	YES	NO
Manage Isilon Clusters	Allows a User to manage Isilon Clusters.	YES	NO	YES	NO
Manage NetApp Clusters	Allows a User to manage NetApp Clusters.	YES	NO	YES	NO
Manage NetApp VServers	Allows a User to manage NetApp VServers.	YES	NO	YES	NO
Manage Generic NAS hosts	Allows a User to manage Generic NAS hosts.	YES	NO	YES	NO
Manage Pure Flashblade Storage Arrays	Allows a User to manage Pure Flashblade Storage Arrays.	YES	NO	YES	NO
Manage GPFS Cluster	Allows a User to manage GPFS Cluster.	YES	NO	YES	NO
Manage Elastifile Clusters	Allows a User to manage Elastifile Clusters.	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage Exchange Servers	Allows a User to manage Exchange Servers.	YES	NO	YES	NO
Manage Cassandra Clusters	Allows a User to manage Cassandra Clusters.	YES	NO	YES	NO
Manage All Protection Sources	This privilege allows a User to manage all Source types.	YES	NO	YES	NO
Manage Couchbase Clusters	This privilege allows a User to manage Couchbase Clusters.	YES	NO	YES	NO
Manage HBASE Sources	This privilege allows a User to manage HBASE sources of Hadoop Clusters.	YES	NO	YES	NO
Manage Acropolis Generic Standalone Clusters	This privilege allows a User to manage Acropolis Generic Standalone Clusters.	YES	NO	YES	NO
Manage SFDC Adapter Sources	Allows a user to manage SFDC Adapter Sources	YES	NO	YES	NO

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage HDFS Sources	This privilege allows a User to manage HDFS sources of Hadoop Clusters.	YES	NO	YES	NO
Manage Universal Data Adapter Sources	This privilege allows a User to manage Universal Data Adapter sources.	YES	NO	YES	NO
Manage Kubernetes Sources	This privilege allows a User to manage Kubernetes sources.	YES	NO	YES	NO
Manage MongoDB Clusters	This privilege allows a User to manage MongoDB Clusters.	YES	NO	YES	NO
Manage HIVE Sources	This privilege allows a User to manage HIVE sources of Hadoop Clusters.	YES	NO	YES	NO

Storage Management

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
Manage Cohesity Storage Domains	Allows the user to create and modify Cohesity Storage Domains.	NO	NO	YES	NO
Manage Cohesity Views	Allows the user to create, modify and delete Cohesity Views.	YES	NO	YES	NO
Read Cohesity Views	Allows the user to view Cohesity Views.	YES	NO	YES	NO
Read Cohesity Storage Domains	Allows the user to view Cohesity Storage Domains.	YES	NO	YES	NO

High Classified

Privilege	Description	Available to Organizations	Available when Multitenancy is Enabled	Available during Impersonation	Available with Hybrid Extender
High Classified	Allows the user to fetch cluster details needed for specific API calls.	NO	NO	YES	NO

Single Sign-On (SSO)

You can configure the Cohesity cluster to use an Identity Provider (IdP) for single sign-on (SSO) access to the cluster. For a multitenancy-enabled cluster, you can configure SSO for an organization defined in the cluster.

> **Important:** If you are using SSO for authentication, Cohesity recommends that you enable multi-factor authentication (MFA) for your SSO system.

First, add the cluster as an application to your IdP and get an SSO URL and certificate file. Then configure the SSO in the cluster and provide the SSO URL and certificate file. After the integration is configured, users can sign in to the cluster using the following methods:

- Use the IdP sign-in page.
- Use the Cohesity cluster sign-in page. Click the **Sign in with SSO** link at the bottom of the cluster sign-in page.

> **Note:** Cohesity supports any Identity Provider (Idp) as long as it uses SAML 2.0 based SSO. The Assertion Consumer Service (ACS) url and Entity ID for any IdP is https://cluster_fqdn/idps/authenticate.

Integrate an Identity Provider

You have to integrate an identity provider to configure SSO. The following identity providers are supported and you can follow the documentation link to learn more about how to integrate it with the Cohesity.

> **Note:** You can configure multiple Identity Providers (IDPs) when using Organizations. Each organization can have one IDP configured at the cluster level.

Identity Provider	Documentation Link
Active Directory Federation Services (ADFS)	Integration with Active Directory Federation Services for SSO
Azure Active Directory	Integration with Azure Active Directory for SSO
Duo Single Sign-on	Integration with Duo for SSO
Ping Identity	Integration with Ping Identity for SSO
Okta Single Sign-on	Integration with Okta for SSO

Add an SSO Provider to the Cohesity cluster

1. In the Cohesity Dashboard, select **Settings > Access Management** and select the **SSO** tab.
2. Click **Add SSO**.

3. Select one of the following options:
   - **SAML**: Security Assertion Markup Language (SAML) is an XML-based protocol used for SSO login.
   - **OpenID Connect**: OpenID Connect is an open authentication protocol that uses OAuth2.0 framework.
4. If you select **SAML**, then refer to the following table:

Name	Description
SSO Domain	Your choice of domain name, such as abc.com. This name must be unique within the cluster.
SSO Provider	Your choice of provider name, such as Okta.
Assign to Organization	Optional. In a multitenant-enabled cluster, you can configure SSO for an organization that has been added to the Cohesity cluster. Select an organization from the drop-down.
Single Sign-on URL	Paste the URL that you copied from your IdP. For example: https://mycompany.okta.com/app/cohesitymycompany_clusterapp/******/sso/sam
Provider Issuer ID	Paste the issuer ID that you copied from your IdP. For example: http://okta.com/*****
X.509 Certificate	Click Select File and browse to the location to select the file you downloaded and renamed previously. For example, okta.pem. Select Edit to edit the certificate.

5. Click **Add**. If you select **OpenID Connect** (OIDC), perform the following steps and then refer to the table:

   Prerequisites:

   1. Create the OIDC app within your Identity Provider (IdP). For more information, see Create OIDC app integrations.

      > **Note:** OIDC is an open standard and Single-Sign-On with Cohesity Platform is intended to work with any OpenID Connect supported Identity Provider. For setup details, refer to your Identity Provider's documentation.

   2. Map the OIDC configuration details from Okta IDP to Cohesity Platform side configurations:

1. To get the Issuer ID:
    1. Navigate to **Security** > **API**.
    2. On the **API** page, click **Authorization Servers**.

       You can find the issuer ID in the Issuer URI section.

2. To get the Client ID:
    1. Navigate to Applications.
    2. On the Helios Test OIDC App page, click **General**.
3. To generate the **JSON Web Key Set (JWKS)** URL:
    1. Construct the URL as follows:

       Format: <issuer ID>/.well-known/openid-configuration

       For example: `https://***-00000000.okta.com/oauth2/default/.well-known/openid-configuration`.
    2. Enter the constructed URL in the address bar of the browser and JSON output will be displayed.

Name	Description
OpenID Server Domain	Enter a unique domain name.
OpenID Server URL for the public (JWKS)	Enter the JSON Web Key Set (JWKS) URL. You can get this URL from your identity provider.
Polling Interval (Hours)	Specifies the caching duration for responses from the JWKS endpoint.
Audience IDs	Enter the IDs verified against the 'aud' field in the JWT. The 'aud' in the JWT must align with one of the strings listed in the Audience IDs field.

> **Important:**
> When enabling signed requests in Okta (or any SAML identity provider), you must explicitly set the following values:
>
> - Signature Algorithm: RSA-SHA256
> - Digest Algorithm: SHA256

6. Click **Add**.

> **Important:** Make sure to add an SSO user before you add the SSO provider. For more information, see Add SSO Users.

The cluster validates the connection to the IdP. If the connection succeeds, the SSO provider is added to the provider list and you can edit, delete or deactivate it as described in the following sections.

Users can start accessing the cluster via their IdP home page or the cluster sign-in page by clicking the **Sign in with SSO** link.

Edit an SSO Provider

You can change any of the SSO provider options.

1. In the Cohesity Dashboard, select **Settings > Access Management** and select the **SSO** tab.
2. Click the Actions Menu (⋮) and then click **Edit**.
3. Change the options as needed.
4. Click **Update**.

The cluster validates the connection to the IdP using the new information.

Deactivate or Activate an SSO Provider

You might want to deactivate an SSO provider for testing or investigation purposes. Deactivation does not delete the provider configuration, so you can activate it later. Once deactivated, users associated with the IdP will no longer bypass the Cohesity sign-in page.

1. In the Cohesity Dashboard, select **Settings > Access Management** and select the **SSO** tab.
2. Click the Actions Menu (⋮) and then click **Deactivate** or **Activate**.

Delete an SSO Provider

You can permanently delete an SSO provider if you no longer need it. Once deleted, users associated with the IdP provider will no longer bypass the Cohesity sign-in page.

1. In the Cohesity Dashboard, select **Settings > Access Management** and select the **SSO** tab.
2. Click the Actions Menu (⋮) and then click **Delete**.
3. Respond to the confirmation prompt.

Active Directory (AD)

Active Directory (AD) in Cohesity Access Management is a directory service implemented by Microsoft that provides several services such as LDAP (Lightweight Directory Access Protocol), Kerberos, and DNS. It is used for user and group identity management and authentication within a Cohesity cluster.

You can join the Cohesity cluster to one or more Active Directory (AD) domains. Joining to AD allows the following:

- Existing AD users and groups can be assigned Cohesity roles
- SMB shares created by the Cohesity cluster can be authenticated based on AD credentials

Cohesity Roles

Once the Cohesity cluster is joined to an AD domain, existing AD principals can be added assigned different roles so they can access the Cohesity Dashboard. For example, you can assign the View role to the AD group "ITGroup." This gives all group members read-only access to the Cohesity Dashboard. You can also assign the Admin role to the AD user "ITManager."

Cohesity SMB Share Access

The action steps in the following procedure are represented in the figure below with blue numbered time squares:

. The procedure assumes the Cohesity cluster is already joined to an Active Directory domain.

Client access to an Active Directory-joined Cohesity cluster SMB share consists of the following actions:

T1) A client requests access to a Cohesity cluster SMB share.

T2) The client attempts to perform an operation on the Cohesity cluster SMB share.

T3) Active Directory authenticates the client's membership within a group that can perform the requested operations on the Cohesity cluster SMB share.

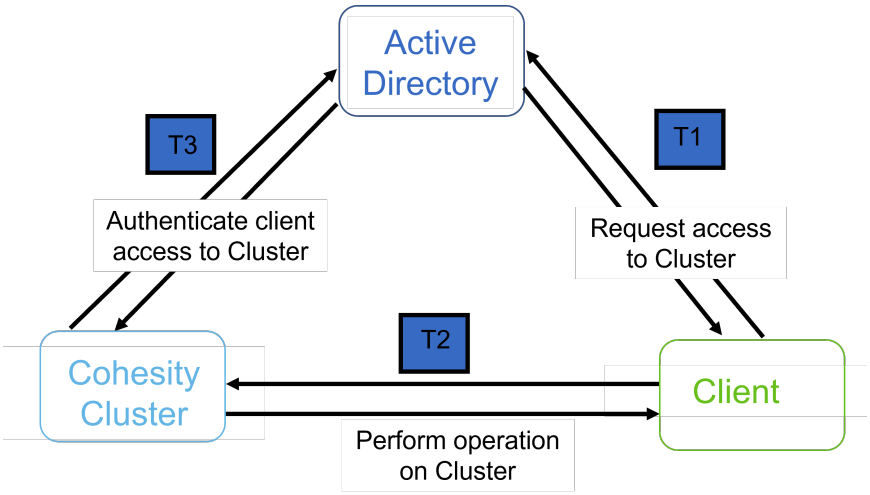

Join to Active Directory Domain

For detailed instructions, see Join to Active Directory.

Join Active Directory

This topic explains how to configure a Cohesity cluster to use Active Directory (AD) credentials for SMB access to views and to access the Cohesity Dashboard. (For information about protecting Active Directory domain controllers, see About Active Directory Data Protection.)

Joining the Cohesity cluster to an AD domain is optional. You can join the Cohesity cluster to one or more AD domains to enable:

- existing AD users can access Cohesity cluster SMB shares based on their AD credentials
- existing AD users and groups can be assigned Cohesity roles and access the Cohesity Dashboard

On Cohesity clusters that are **not** joined to AD, all views are accessible as SMB shares and there is no authentication. If the AD administrator blocks NTLM, the shares on the Cohesity cluster are not accessible.

All accounts and Cohesity cluster SMB share privileges for users in AD are managed through Active Directory or where the SMB shares are mounted. Non-admin AD users must be granted backup and restore privileges on Windows clients explicitly. This is in addition to backup and restore privileges for the Cohesity cluster.

Migrated AD Users and Groups with Historical SIDs

Cohesity supports migrated AD users and groups with historical SIDs and grants access accordingly in Active Directory domains. The SID history retains user and group membership and access rights during Active Directory domain migration.

Considerations

Review the following considerations of migrated AD users and groups with historical SIDs:

- Historical SIDs with multiprotocol access is not supported.
- Historical SIDs user quota accounting is not supported.
- Privilege assignment for Historical SIDs is not supported.

Prerequisites and Considerations

For information about the required firewall ports for AD, see Manage Firewall Ports.

Join the Cohesity cluster to an AD Domain

1. Go to **Settings > Access Management** and click the **Active Directory** tab.
2. Click **Add Active Directory**.
3. **Domain Name:** Provide the fully qualified domain name (FQDN) of the Active Directory domain to join.
4. **Username**: Provide the samAccountName.
5. **Password**: Provide the password for the username.
6. **Domain Controllers:** Add preferred domain controllers to the AD domain to ensure that the Cohesity cluster uses a local domain controller (when possible) to speed up AD queries. You can set preferred domain controllers for the following:
   - Domain controllers
   - Trusted domains of the domain controller

   > **Note:**
   > - The SRV records must exist for the domain (the list of domain controllers displayed in the Cohesity Dashboard is generated by querying the SRV records).
   > - You can add any number of preferred domain controllers.

   1. Click **Manage** to allow or deny domain controllers and change the order of allowed controllers:

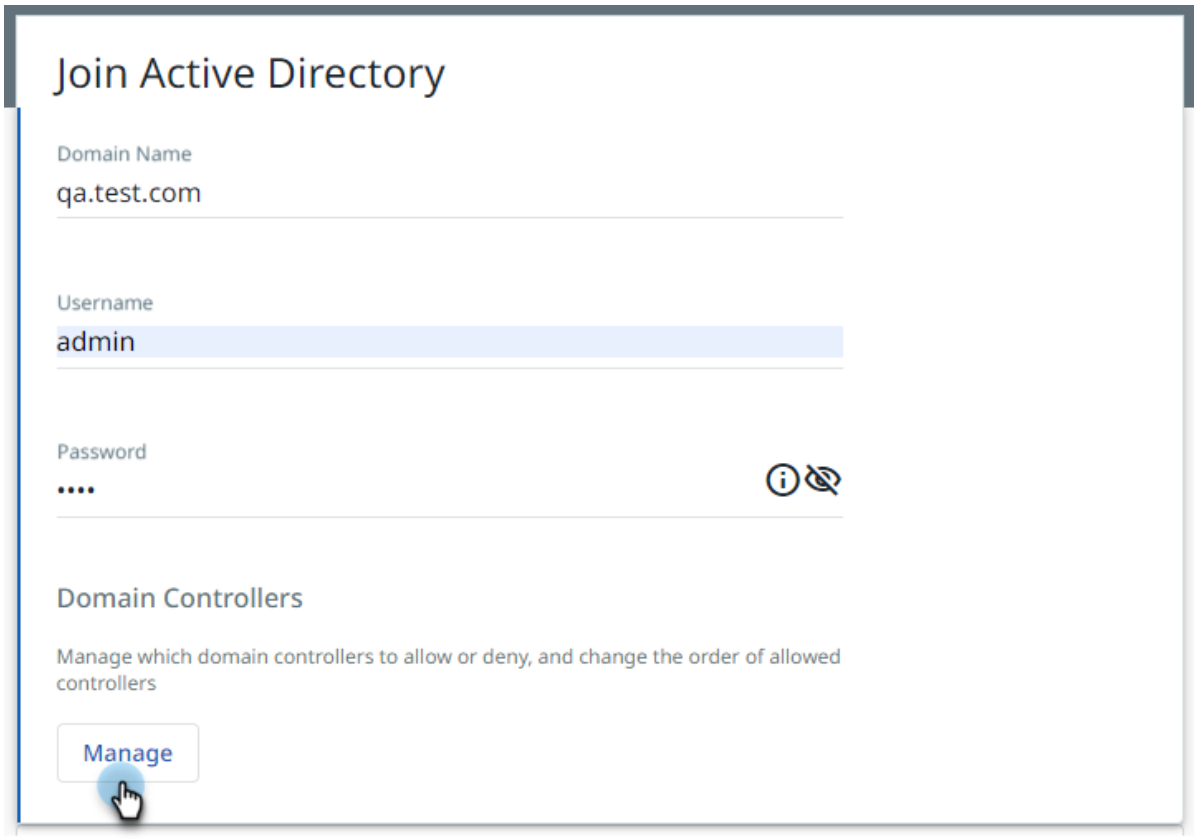

2. In the **Manage Domain Controllers for AD Domain** dialog, the domain controllers are listed in order of preference. You can move the allowed domain controller(s) up and down to put them in order of preference. To block a particular domain controller, hover on it and click **Deny:**

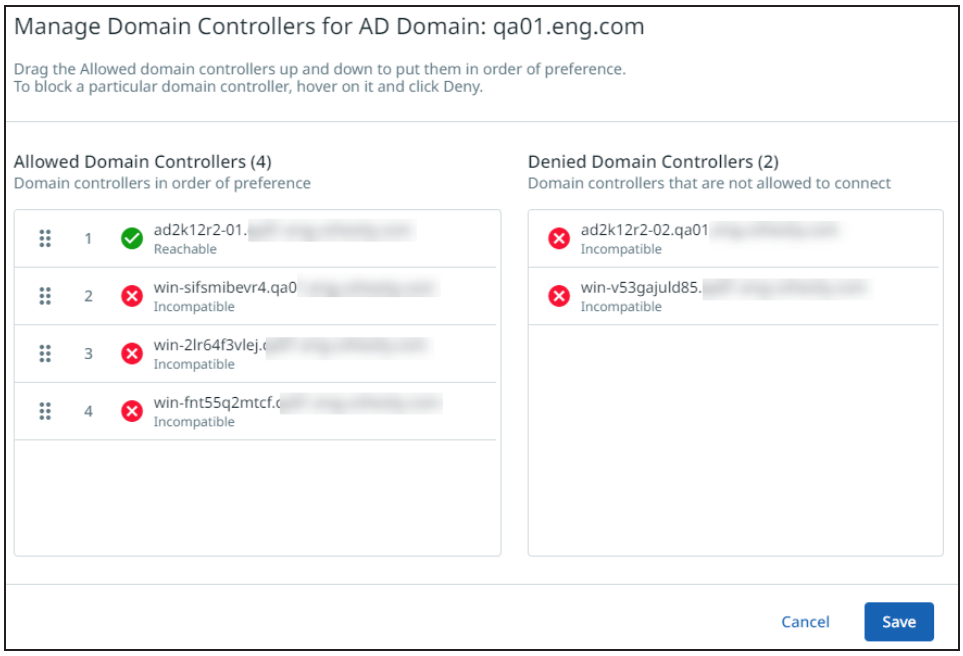

3. Click **Save**.

   The allowed domain controllers are listed in the **Domain Controllers** section:

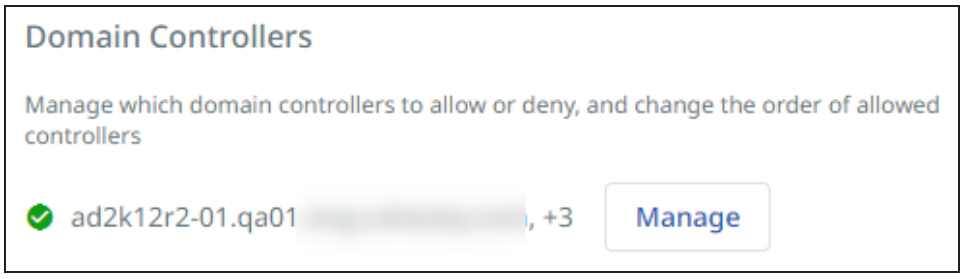

   The Cohesity cluster uses these domain controllers to join the AD domain and for all future communication with the AD domain.

7. **Machine Accounts:** Click **Add** and specify the following details:
   - **Machine Account**—Provide at least one machine account to identify the Cohesity cluster on the domain. For example, you can use the Cohesity cluster name.
   - **DNS Hostname**—Optionally, specify the DNS hostname of the machine account. For AD to work with the internal DNS server, apart from machine account names, you should add DNS hostnames that should be configured for each machine account. The DNS hostnames should have one of the In Bound DNS zone suffixes configured for the internal DNS server. For more information, see Internal DNS for Load Balancing.

> **Note:**
> - When you specify the DNS hostname, it is treated as a customized Service Principal Name (SPN), resulting in successful Kerberos authentication with this customized SPN.
> - If the DNS Hostname is modified when editing the machine accounts, all existing configuration of the machine account is overwritten on the Active Directory.

- **Encryption Types**—Optionally, select the encryption type(s) that must be used by the machine account. You can select multiple encryption types. The following encryption types are supported:
  - AES128-CTS-HMAC-SHA1-96
  - AES256-CTS-HMAC-SHA1-96
- **SPNs**—Service Principal Names (SPNs) are identifiers used to associate a service with a specific AD account. SPNs enable authentication and authorization for services. Specify the SPN in the following format:

  <service class>/<host>:<port>/<service name>

  The SPN syntax has four elements.
  - The following two elements are required:
    - **Service Class**: Specify the type of service.
    - **Host**: Specify the hostname or Fully Qualified Domain Name (FQDN) of the server on which the service is hosted.
  - If necessary, you can specify the following optional elements:
    - **Port**: Specify the port number on which the service is running.
    - **Service Name**: Specify the specific service instance on the host.

Note:

- All (or a subset) of the VIPs will resolve to the machine account. You can optionally add multiple machine accounts, which creates aliases for the same cluster. You can map these aliases to a subset (or all) of the VIPs to direct I/O to a subset (or all) of the nodes. Separate names with commas. For example: cluster1, cluster1-vip1, cluster1-vip2, cluster1-vip3, cluster1-vip4.
- If using VLANs—One of the machine accounts must be a DNS hostname for a virtual IP (VIP) of a VLAN configured on the cluster in order to recover to servers within that VLAN.
- After adding a Cohesity cluster to an AD, ensure the machine account used to register the Cohesity cluster in AD has the **Read membership** permissions across all users and groups.
- The same configuration must be done on every additional domain to which the Cohesity cluster is joined.
- The machine account must follow NetBIOS naming conventions of 15 maximum character length, those described in Naming conventions in Active Directory for computers, domains, sites, and OUs and the additional restriction that the ! @ ( ) characters are not supported.

8. **Use the above Machine Accounts even if they already exist in AD Domain:** Optionally toggle on to use the machine accounts that you specified above even if they already exist in the Active Directory domain.
9. **Mapped Provider:** Select the type of mapped provider to join the domain. By default, no mapped provider is selected. Optionally, select **LDAP** or **NIS**. If you select the **LDAP** option, you can select the provider from the **LDAP Provider** drop-down list. If you select the **NIS** option, you can select the NIS provider from the **NIS Provider** drop-down list.
10. **Organizational Unit:** Provide the organizational unit if applicable. Use the format: *OUName* or *OUName*/*SubOUName*. For example: Engineering or Sales/West or West Coast Sales/ North California.
11. **AD Workgroup / NetBIOS Name:** Provide the workgroup / NetBIOS name if applicable.
12. **Discover Trusted Domains:** Optionally, turn on this toggle to discover the trusted domains associated with an AD. You can exclude certain trusted domains after the

Cohesity cluster is joined to the AD domain. For more information, see Trusted Domains.

13. Click **Join**.
14. This step is required only for SMB access to views. On the DNS server that is integrated with the AD domain, add the A record for each machine account used to identify the Cohesity cluster to the AD domain. This is necessary to allow clients to access the Cohesity cluster as a resource. (Adding the A record is not generally necessary for Windows resources because Active Directory automatically adds it to the DNS server. However, the Cohesity cluster does not include the components that enable automatic addition to the DNS server.)
15. Optionally, repeat these steps to join the Cohesity cluster to another AD domain. For more information, see Guidelines for Joining Multiple AD Domains.
16. Optionally, you can add user ID mapping. For more information, see User ID Mapping for Multi Protocol Network Access.

Guidelines for Joining Multiple AD Domains

A cluster can join multiple Active Directory domains even if there are trust relationships between the AD domains.

After the Cohesity cluster joins the AD domain, you can add trusted domains to the allowlist or denylist. By default, all trusted domains are used for authentication unless added to the denylist. For more information, see Trusted Domains.

Only Kerberos authentication is supported for SMB shares when multiple AD domains are joined to the Cohesity cluster. NTLM authentication is disabled.

Leave Active Directory Domain

1. Go to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name and then click the leave domain icon ( ).
3. Provide the **AD Admin** and **Password** for the domain.
4. Click **Leave Domain**.

   When this succeeds, the machine accounts associated with the Cohesity cluster are removed from the Active Directory domain.

   > **Note:** If the Cohesity cluster is removed from a domain, ensure the corresponding DNS entries have been removed from the DNS server.

Enable Load Balancing

You can enable load balancing over multiple Cohesity cluster IPs by following these steps:

> **Note:** These example steps refer to the Windows DNS Manager. If you're using a different DNS Manager, ensure you perform step 3.

1. Open DNS Manager in Windows using the DNS Admin client: **Control Panel > Administrative Tools > DNS**.
2. Enable the Advanced view: **View > Advanced**.
3. Add multiple entries with the same Cohesity cluster name and do the following:
   - Give each entry a different **IP address** that corresponds to a VIP allocated to the Cohesity cluster in the DNS Zone where it will be added.
   - Generally, you can leave **Time to live (TTL)** at its default setting. You should only change this setting for specific use cases. Set TTL to a low value (such as 5 seconds or less) for backup use cases where a low number of clients are accessing the Cohesity cluster. Set TTL to a high value (such as 12 hours) for file services use cases with a high number of clients. This ensures that clients access the same Cohesity cluster nodes and benefit from caching.

Manage Domain Controllers From the Active Directory Details Page

You can also manage domain controllers from the Active Directory details page.

To manage domain controllers from the Active Directory details page:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the active directory for which you want to manage domain controllers.
   - In the **Details for AD Domain** section, click the edit icon corresponding to **Domain Controllers**:

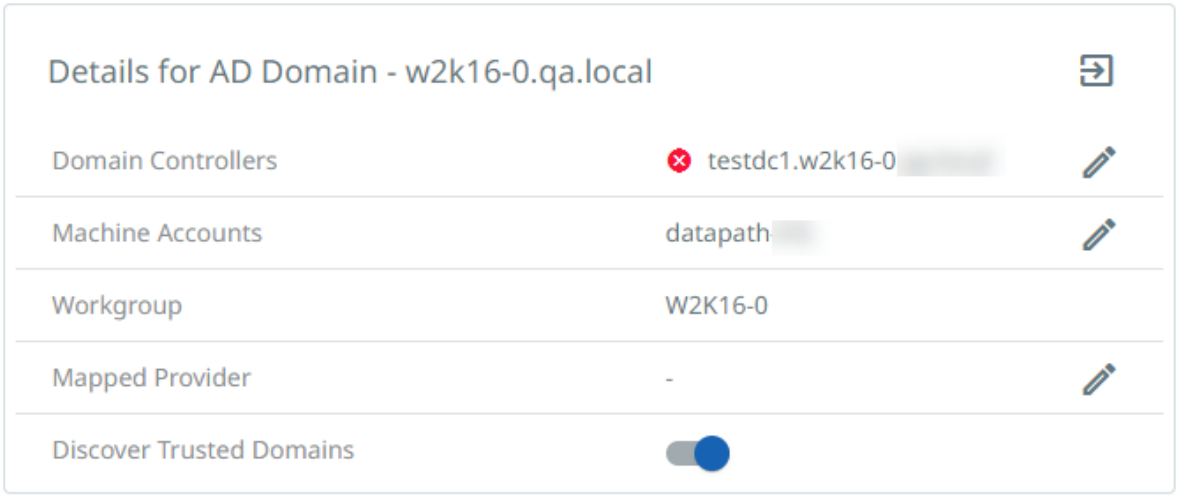

The **Manage Domain Controllers for AD Domain** dialog is displayed. You can manage the domain controllers for this AD domain.

- In the **Discovered Trusted Domains** section, the trusted domain controllers associated with the AD are discovered and listed. Click **Manage** corresponding to the domain name:

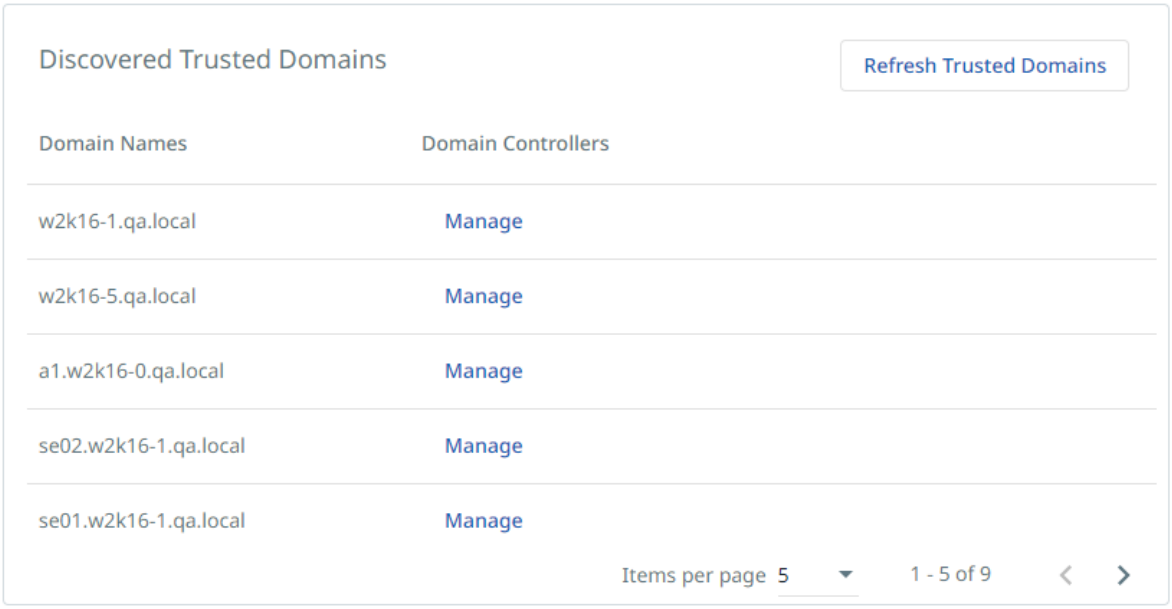

The **Manage Domain Controllers for AD Domain** dialog is displayed. You can manage the trusted domain controllers for the AD domain:

- **Allowed Domain Controllers**—All the domain controllers for the AD domain are listed in the **Allowed Domain Controllers** section. Drag a domain controller up or down to reorder the preference. To block a domain controller, hover over it and click **Deny**.
- **Denied Domain Controllers**—Domain controllers that are disallowed to connect the AD domain are listed in the **Denied Domain Controllers** section. To unblock a domain controller, hover over it and click **Allow**.

User ID Mapping for Multi Protocol Network Access

After the Cohesity cluster is joined to an Active Directory domain, you can add user ID mapping. The ID map database is stored in AD using industry-standard protocols. For more information, see ID Mapping for Multi-Protocol Views.

1. Go to **Settings > Access Management** and click the **Active Directory** tab.
2. Click **Add ID Mapping** or the edit icon ( ).
3. Select an **ID Mapping Type** and enter the required information:

- **Centrify**
  - **AD Domain:** Name of the Active Directory domain where the corresponding Centrify zone is located.
  - **Centrify Zone:** Selected zone that has the ID mappings.

  For Native security style Views, SMB will get Unix IDs from the Active Directory provider and NFS will get Unix IDs from the LDAP provider.

  For Unified or NTFS security style Views, SMB and NFS will get Unix IDs from the corresponding Centrify zone for the Active Directory provider.
- **Fixed**
  - **Unix UID to be applied to all files created:** Enter a user ID to apply to all created files.
  - **Unix GID to be applied to all files created:** Enter a group ID to apply to all created files.
- **LDAP Provider**

  The UID and GID corresponding to the Active Directory users are fetched from the configured LDAP provider. In LDAP, the most commonly used attributes to store the UID and GID are 'uidNumber' and 'gidNumber', respectively.

  All Views, regardless of their security style setting (Native, Unified or NTFS), will get Unix IDs from the LDAP provider and SIDs from the Active Directory provider.
- **RID (Relative ID of SID)**

  RID is a variable length number that is assigned to objects at creation and becomes part of the object's Security Identifier (SID) that uniquely identifies an account or group within a domain. For every Windows user or group, the RID of the corresponding SID is used as the UID or GID.

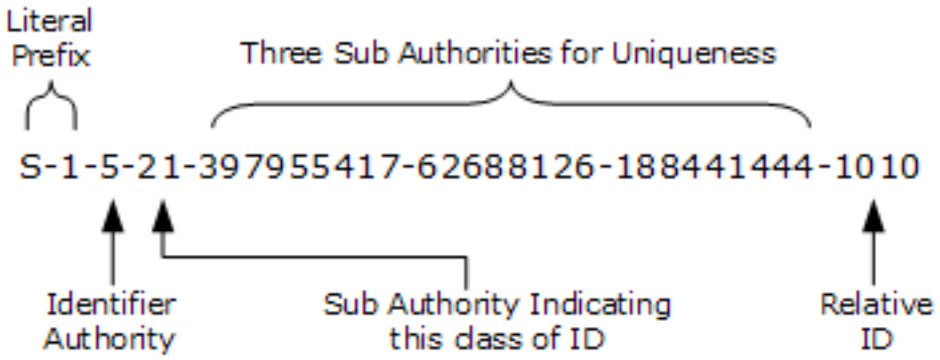

- **RFC2307**

  In RFC2307, the UID and GID are stored along with Active Directory objects in the following attributes: the UID for a user is stored in the attribute ‘uidNumber’ and GID for a group is stored in the attribute 'gidNumber'.

  For Native security style Views, SMB will get Unix IDs from the Active Directory provider and NFS will get Unix IDs from the LDAP provider.

  For Unified or NTFS security style Views, SMB and NFS will get Unix IDs from the Active Directory provider.

- **SFU (Services for Unix 3.0)**

  In SFU, the UID and GID are stored along with Active Directory objects in the following attributes: the UID for a user is stored in the attribute ‘uidNumber’ and GID for a group is stored in the attribute 'gidNumber'.

  For Native security style Views, SMB will get Unix IDs from the Active Directory provider and NFS will get Unix IDs from the LDAP provider.

  For Unified or NTFS security style Views, SMB and NFS will get Unix IDs from the Active Directory provider.

4. Select a **Fallback Option** to use if the selected ID Mapping Type is, for instance, not populated. This is available if the ID Mapping Type is not Fixed or RID.

   If you select Fixed for the Fallback Option, provide the following information:

   - **Unix UID to be applied to all files created:** Enter a user ID to apply to all created files.
   - **Unix GID to be applied to all files created:** Enter a group ID to apply to all created files.
5. Select the **AD User mapped to the Linux Root User** by selecting the AD domain and user.
6. Click **Save**.

Trusted Domains

After the Cohesity cluster joins the Active Directory domain, you can add trusted domains to the allowlist or denylist. By default, all trusted domains are used for authentication unless added to the denylist. You can add trusted domains to the allowlist so that only trusted domains added to the allowlist are used for authentication.

> **Note:** If the domain controller is not functioning correctly, the discovery of trusted domains may take a relatively long period of time or may provide inconsistent results. Therefore, a trusted domain might not be discovered at all. You must manually enter such domains in the allowlist or denylist.

A trusted domain forest can have multiple trusted domains. To provide optimized access, you can use a subset of trusted domains so that the Cohesity cluster does not send group membership queries to all trusted domains. You can choose a subset of trusted domains by using one of the following modes:

- **Denylist**—Ignores specified trusted domains. For example, if there are three trusted domains T1, T2, and T3 and if T2 is added to the denylist, T1 and T3 are the trusted domains used for authentication.
- **Allowlist**—Allows only specified trusted domains. For example, if there are three trusted domains T1, T2, and T3 and if T1 and T2 are added to the allowlist, T1 and T2 are the trusted domains used for authentication.

Add Trusted Domains to the Allowlist

To add a trusted domain to the allowlist:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page.
3. Click the edit icon in the **Denied Trusted Domains** section.
4. Click **Allowlist** and add the domain name in the **Allowed Trusted Domains** field. Click + to add another domain.

   Cohesity recommends adding at least one domain or switching to a denylist.
5. Click **Save**.

   The trusted domain is added to the allowlist.

Add Trusted Domains to the Denylist

If you want Cohesity to ignore specific trusted domains so that their principals are not assigned roles and denied access to Cohesity SMB shares, you can add the desired domains to the denylist.

To add a trusted domain to the denylist:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page.
3. Click the edit icon in the **Denied Trusted Domains** section.
4. Click **Denylist** and add the domain name in the **Denied Trusted Domains** field. Click + to add another domain.
5. Click **Save**.

   The trusted domain is added to the denylist.

Refresh Trusted Domains

If the **Discover Trusted Domains** option is enabled, the trusted domains associated with the AD are discovered and listed in the **Discovered Trusted Domains** section. After any

configuration change, trusted domains are not automatically refreshed in the **Discovered Trusted Domains** section. You must manually refresh the trusted domains.

To refresh the trusted domains list:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page:
3. In the **Discovered Trusted Domains** section, click **Refresh Trusted Domains**.

Add or Edit Preferred Domain Controllers

To add or edit preferred domain controllers of the discovered trusted domains:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page:
3. In the **Discovered Trusted Domains** section, click **Select** corresponding to the domain name.
4. On the **Edit Preferred Domain Controllers** dialog, modify the domain controllers and click **Save**.

View Trust Relationships for Domains

To view the trust relationships between the domains:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the **View Trust Relationships** button. The topology of the trust relationships between domains is displayed.This page helps you analyze complex AD topologies with multiple domains and trusts.

   You can view the following details for the trust relationships:

   - **Domain Category**: Primary and Trusted domains. The domains are color-coded for easy differentiation.
   - **Relationship Direction**: One-way trust and Two-way trust.
   - **Filters**: Primary Domains, Trusted Domains, and Trusted Relationships.

     You can also drag the AD nodes to reorganize the topology.

One-Way Trust

A one-way trust is a unidirectional authentication path created between two AD domains. In a one-way trust, the trust relationship is only one way.

The following is an illustration of one-way trust and the direction of access:

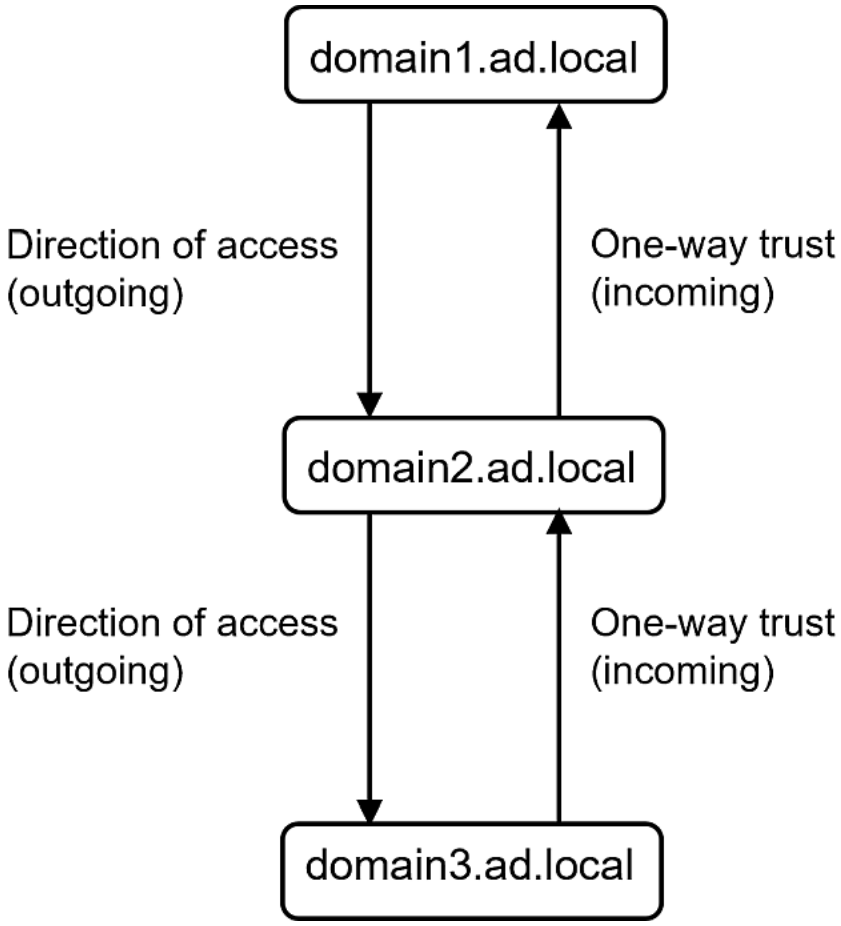

As shown in the above figure:

- domain2.ad.local has a one-way trust with domain1.ad.local. domain1.ad.local does not have a trust relationship with domain2.ad.local. Users in domain2.ad.local can access resources in domain1.ad.local. However, users in domain1.ad.local cannot access resources in domain2.ad.local.
- domain3.ad.local has a one-way trust with domain2.ad.local. domain2.ad.local does not have a trust relationship with domain3.ad.local. Users in domain3.ad.local can access resources in domain2.ad.local. However, users in domain2.ad.local cannot access resources in domain3.ad.local.

In a one-way trust, the direction of trust is the opposite of the direction of access. An outgoing trust allows incoming access and an incoming trust allows outgoing access.

If one-way trust is enabled on the AD domain, trusted domains are listed in the following way in the Cohesity cluster:

- If the Cohesity cluster joins domain2.ad.local and if trusted domains are enabled:
    - domain1.ad.local is displayed as a trusted domain because domain1.ad.local trusts domain2.ad.local

    - domain3.ad.local is not displayed as a trusted domain because domain3.ad.local does not trust domain2.ad.local
- If the Cohesity cluster joins domain3.ad.local, both domain2.ad.local and domain1.ad.local are displayed as trusted domains because domain2.ad.local trusts domain3.ad.local and domain1.ad.local trusts domain2.ad.local.
- If the Cohesity cluster joins domain1.ad.local, no domain is displayed as a trusted domain because domain1.ad.local does not have a trust relationship with domain2.ad.local or domain3.ad.local.

Two-Way Trust

Two-way trust refers to the trust relationship between two domains where each domain trusts the other. Users and resources in both domains can access resources in the other domain. In a two-way trust relationship, authentication requests can flow in both directions, allowing users from one domain to access resources in another and vice versa without encountering authentication issues. For example, if domain A establishes a two-way trust with domain B, it implies that domain B reciprocally trusts domain A. As a result, both domains can seamlessly share resources.

Cohesity supports the following types of two-way trust:

- Parent-Child Trust
- Tree-Root Trust

Parent-Child Trust

A parent-child trust is a two-way, transitive trust relationship between a parent domain and its child domain within the same forest. The parent domain sits atop the forest and governs its child domains. A child domain is a subdomain established beneath its parent, inheriting policies and settings while also capable of possessing its distinct policies and objects.

This trust relationship facilitates seamless access for users, groups, and computers within the child domain to resources located in the parent domain, and vice versa. Tree-root trust is bi-directional as it allows authentication and authorization traffic to flow in both directions between the parent and child domains.

Tree-Root Trust

A tree-root trust is a two-way, transitive trust relationship automatically established between a new domain tree and the forest root domain when the tree is added to a forest. Users from any domain in the tree can access resources in the forest root domain and vice versa, without needing additional trust relationships. Tree-root trust is bi-directional as it allows authentication and authorization traffic to flow in both directions between the two forests.

View Active Directory Join and Leave Status

You can add one or more Active Directory (AD) domains to a Cohesity cluster. Joining the AD domain to a Cohesity cluster is a multi-step process and involves the following steps:

- Connecting to the DNS server
- Getting the list of domain controllers from the DNS server
- Connecting to the DC server
- Adding machine account(s)
- Setting the encryption type

You can view the progress of each subtask. For example, while joining an AD domain, you can check whether the machine account was added successfully. This log helps in troubleshooting and identifying the exact step at which the task failed.

View Active Directory Join Status

To view the AD join status:

1. Follow the steps described on the Join Active Directory page to join the Cohesity cluster to the AD domain.
2. After you click **Join**, the following error message is displayed if the process fails:

   ```
   Joining Active Directory <ad-domain-name> failed. Show Subtasks
   ```

3. Click **Show Subtasks**.

   A dialog is displayed and the subtasks are listed.

   You can view the subtask at which the process failed.

You can also view the status when you edit an AD domain. For example, if you add or modify a machine account after the Cohesity cluster joins the AD domain and if the process fails, you can view the status by following the steps outlined below:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the AD domain that you want to edit.

   The AD Domain details page is displayed.
3. Click the edit icon corresponding to **Machine Accounts**.
4. Modify the machine account details as necessary.
5. Click **Save**.

   The following error message is displayed if the process fails:

   ```
   Failed to join Active Directory. Show Subtasks
   ```

6. Click **Show Subtasks** to view the subtask at which the process failed.

View Active Directory Leave Status

Leaving an AD domain is a multi-step process. You can view the progress of each subtask. For example, while leaving an AD domain, you can check whether the machine account was removed successfully. This log helps in troubleshooting and identifying the exact step at which the task failed.

To view the AD leave status:

1. Follow the steps described in the Leave Active Directory Domain section to remove the Cohesity cluster from an AD domain.
2. Click **Leave Domain**.

   The following error message is displayed if the process fails:

   ```
   Failed to leave Active Directory. Show Subtasks
   ```
3. Click **Show Subtasks** to view the subtask at which the process failed.

Certificate-Based Authentication

If the Cohesity cluster is joined to an Active Directory domain, certificate-based authentication can be enabled for the Cluster, requiring users to sign in using a digital certificate instead of a user name and password.

For added security, if you want to determine if the incoming certificate is from a trusted source, you must set up a trust store in the Cohesity cluster and add trusted root certificates to the trust store. For more information, see Trust Store in Cohesity Cluster.

Prerequisites and Considerations

- The Cohesity cluster must be joined to an Active Directory domain. For more information, see Join Active Directory.
- Active Directory should have Certificate Services installed if you plan to use self-signed certificates. For more information, see Install the Certification Authority.
- Each user must have a valid self-signed certificate or Authority-issued certificate in `.pem` format. The key resides on the user's computer.
- An AD user or group must be mapped to one or more Cohesity roles. For example, you might map the Administrators group to the Cohesity Admin role to grant users in the group Admin privileges in the Cohesity cluster. See the following procedure for details.
- If an AD user is defined in different AD groups and those groups are mapped to different Cohesity roles, the user will have the privileges for the combined roles on the Cluster.

To enable certificate-based authentication:

> **Note:** Enabling or disabling certificate-based authentication automatically restarts the Cluster and signs out all users. The Cluster restart can take a several minutes to complete.

1. In Active Directory Certificate Services, map the certificates to users.
2. Make sure Cohesity users have access to their certificates so they can install them when prompted to do so.

3. Map each AD user or group to one or more Cohesity cluster roles:
    a. In the Cohesity Dashboard, select **Settings > Access Management**.
    b. Click **Add AD Users & Groups**.
    c. Select **Active Directory Users & Groups**.
    d. Select a domain from the **Active Directory Domain** dropdown.
    e. Select a user or group from the **User(s) or Group(s)** dropdown.
    f. Select one or more roles for the group or user in the **Roles** dropdown.
    g. Click **Add**.
4. Enable certificate-based authentication for the Cohesity cluster:
    a. In the Cohesity Dashboard, select **Settings > Summary** and click **Configure**.
    b. Toggle on **Enable Certificate-based Authentication**.
    c. Respond to the confirmation prompt.

You and other users will be automatically signed out of the Cluster when it restarts. After the Cluster restarts:

- Users defined in AD are prompted to install the certificate the first time they access the Cohesity Dashboard. Once the certificate is installed and the user is authenticated against AD, the Cohesity Dashboard is displayed (the user bypasses the sign-in prompt).
- Users who attempt to sign in with an invalid certificate will see a `403 forbidden` error and message indicating that certificate authentication is enabled.
- Users who attempt to sign in with no certificate are presented with the sign-in prompt, but no AD domains are shown in the prompt.
- Local users (defined locally in the Cluster) can continue to sign in as usual with their user name and password.

To authenticate using the Cohesity CLI:

If certificate-based authentication is enabled on the Cohesity cluster, users who want to use the Cohesity CLI must first authenticate to the Cluster by using the following command:

```
iris_cli --server=<ip address> -cert <absolute path to public key file> -key
<absolute path to private key>
```

A sample command is shown below:

```
./iris_cli --server=192.0.2.3 -cert /a/b/c -key /a/b/c
```

The command options are:

- `-server` The IP address or hostname of the Cohesity cluster.
- `-cert` The absolute file path for the public key file (`.pem`).
- `-key` The absolute file path to the private key file (`.key`) that corresponds to the certificate file.

To change the authentication type via the Cohesity CLI:

In addition to enabling and disabling certificate-based authentication in the Cohesity Dashboard, you can use the following Cohesity CLI command. Note that the Cluster must be restarted after changing authentication with the CLI command.

```
iris_cli cluster update authentication-type=<string>
```

Specify one of the following options for `<string>`:

- `password-only` Disables certificate-based authentication. All users must provide passwords to sign in.
- `certificate-only` Enables certificate-based authentication. Users defined in the joined AD domain must have certificates to sign in to the Cluster. Local users must provide passwords.

Certificate Mapping-Based Authentication

If the Cohesity cluster is joined to an Active Directory (AD) domain, Active Directory (AD) users can authenticate to the Cohesity cluster using certificate mapping-based authentication.

In certificate mapping-based authentication, the AD user is authenticated based on the attribute mappings you have configured for the certificate and AD. In the Cohesity cluster, you can define the following attributes of the certificate and AD that must map for the user authentication:

- Certificate Attributes:
  - CommonName
  - EmailAddress
  - UserPrincipalName
- Active Directory Attributes:
  - SamAccountName
  - UserPrincipalName

Configure the Authentication Attributes

In the Cohesity cluster, you can define the certificate and AD attributes that must map for user authentication.

To configure the authentication attributes:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

```
./iris_cli -server 172.16.3.101 -username=admin
```

2. To define the certificate and AD attributes that must map for user authentication, run the following command:

   ```
   iris_cli security-config update ad-mapping=<AD_attributes> certificate-mapping=<cert_attributes>
   ```

   Where:

   - `<AD_attributes>` - Any one of the following AD attributes:
     - SamAccountName
     - UserPrincipalName
   - `<cert_attributes>` - Any one of the following certificate attributes:
     - CommonName
     - EmailAddress
     - UserPrincipalName

   Sample:

   ```
   iris_cli security-config update ad-mapping=UserPrincipalName certificate-mapping=EmailAddress
   ```

Enable Certificate Mapping-Based Authentication

In the Cohesity cluster, you can enable certificate mapping-based authentication using the Cohesity CLI.

To enable certificate mapping-based authentication:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. To enable certificate mapping-based authentication, run the following command:

   ```
   iris_cli security-config update enable-mapping-based-auth=true
   ```

The AD users are now authenticated to the Cohesity cluster using certificate mapping-based authentication. For added security, if you want to determine if the incoming certificate is from a trusted source, you must set up a trust store in the Cohesity cluster and add trusted root certificates to the trust store. For more information, see Trust Store in Cohesity Cluster.

Trust Store in Cohesity Cluster

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

For security purposes, you can set up a trust store in the Cohesity cluster and add trusted root certificates or certificate authority (CA) to the trust store. If the incoming certificate forms a valid certificate chain with at least one of the trust store certificates, then the incoming certificate is determined to be from a trusted source.

Once the AD users are authenticated using certificate mapping-based authentication and the digital certificate is determined to be from a trusted source, the AD user can log into the Cohesity cluster.

Set Up Trust Store

You can set up a trust store in the Cohesity cluster using the Cohesity CLI.

Prerequisites

Ensure that:

- The certificate is in PEM format.
- You copy the certificate to the Cohesity cluster.

To set up a trust store:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Validate the certificate. Run the following command:

   ```
   iris_cli trusted-cas register name=rootCert desc=testing file-
   path=rootCert.pem only-validate=true
   ```

3. If the certificate is valid, add the certificate to the trust store. Run the following command:

   ```
   iris_cli trusted-cas register name=rootCert desc=testing file-
   path=rootCert.pem
   ```

4. You can view the added certificates to the trust store using the following command:

   ```
   iris_cli trusted-cas ls
   ```

5. Enable the feature that validates if the incoming certificate forms a valid certificate chain. Cohesity Support to enable the feature.
6. If you have not enabled certificate-based authentication or certificate mapping-based authentication on the Cohesity cluster, enable it. For more information, see Certificate-Based Authentication or Certificate Mapping-Based Authentication.

The AD users are now authenticated to the Cohesity cluster based on the type of authentication you have configured, such as certificate-based authentication or certificate mapping-based authentication. And, the user can log in to the Cohesity cluster using only the trusted certificates.

Kerberos

Cohesity supports Kerberos authentication using both NFSv3.0 and NFSv4.1. Kerberos prevents unauthorized users from accessing Cohesity NFS exports.

Kerberos is a network authentication protocol that allows users to securely access services over a physically insecure network.

Kerberos has three parts:

- Client
- Server
- Kerberos Key Distribution Center (KDC) - KDC mediates between client and server.

Clients obtain tickets from the Kerberos Key Distribution Center(KDC), and they present these tickets to servers when connections are established. For more information, see Kerberos.

Supported Kerberos Flavors

Cohesity supports the following two flavors of Kerberos:

- Microsoft AD Kerberos
- MIT Kerberos

Kerberos Terminology

Realm

A network that uses Kerberos, composed of one or more servers called KDCs and a potentially large number of clients.

Key Distribution Center (KDC)

A service that issues Kerberos tickets, usually run on the same host as the Ticket Granting Server.

Ticket Granting Service (TGS)

A server that issues tickets for the desired service which are in turn given to users for access to the service. The TGS usually runs on the same host as the KDC.

Ticket Granting Ticket (TGT)

A special ticket that allows the client to obtain additional tickets without applying for them from the KDC.

Principal

A user or service that can authenticate using Kerberos. A principal's name is in the form root [/instance]@REALM. For a typical user, the root is the same as their login ID. The instance is optional. If the principal has an instance, it is separated from the root with a forward slash ("/"). An empty string ("") is considered a valid instance (which differs from the default NULL instance), but using it can be confusing. All principals in a realm have their own key, which is derived from their password or randomly set for services.

For more information on Service Principal Name and Host Principal Name, see Service Principal Names.

Client

An entity on the network (a user, a host, or an application) that can get a ticket from Kerberos.

Microsoft AD Kerberos

Kerberos is the authentication service in Microsoft Active Directory.

Configure AD Kerberos on Cohesity Cluster

Components in AD Kerberos configuration on the cluster:

- AD Server
- LDAP Server
- Unix Clients
- Cohesity cluster

NFS User and Group Configuration Parameters in AD

When you use AD as an LDAP provider for NFS with Kerberos, you must set the following parameters for a user or group.

You must set the following parameters for a user:

- uid—username (same as sAMAccountName)
- uidNumber—User identifier
- gidNumber—Group identifier of the primary group

- objectClass—posixAccount

  Instead of setting the objectClass in AD, you can contact Cohesity Support to help use one of the objectClasses for a user in AD.

You must set the following parameters for a group:

- gidNumber—Group identifier of the group
- objectClass—posixGroup

  Instead of setting the objectClass in AD, you can contact Cohesity Support to help use one of the objectClasses for a group in AD.
- If a group is a secondary group for some users, then set:

  memberUid—List of uidNumbers

> **Note:** Make sure the AD group memberships reflect the same as these UNIX memberships.

Configure AD Kerberos on a Cohesity Cluster

To configure AD Kerberos on a Cohesity cluster, complete the following high-level steps:

1. Create a DNS entry (in the DNS Manager option) on the AD server with the machine account name to provide IP mapping.
2. Add the AD domain to the Cohesity cluster. For more information, see Join the Cohesity cluster to an AD Domain.
3. Add an LDAP server to the Cohesity cluster and map it to the AD server.
4. Log in to the AD server and complete the following steps:
    1. Right-click the corresponding machine account and select **Properties**.
    2. Go to the **Attribute Editor** tab and double-click the **servicePrincipalName** parameter.
    3. In the **Value to add** text box, add **nfs** followed by the machine account name and domain name and click **Add**. For example, considering **datapath-stress** is the machine account name and **w2k16-1.test.local** is the domain name, the

value is: **nfs/datapath-stress.w2k16-1.test.local**:

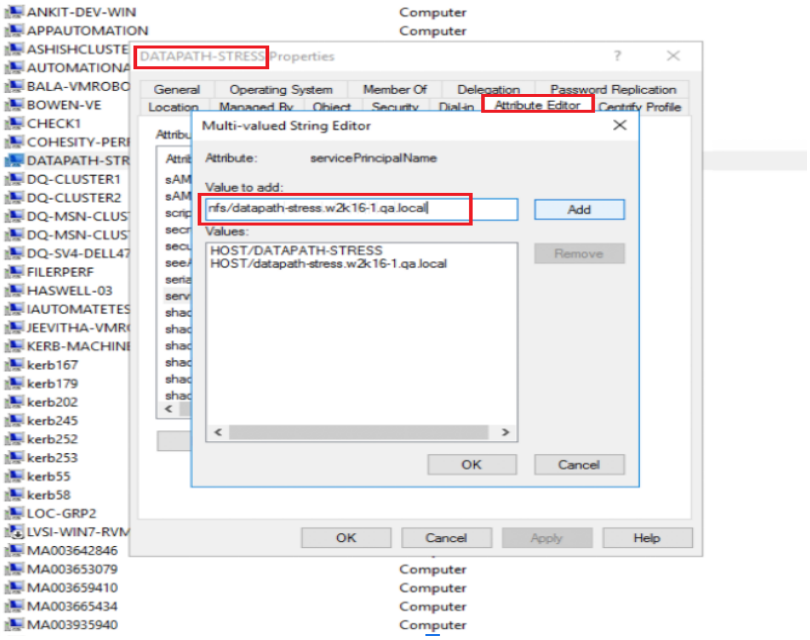

5. If your environment has multiple domain controllers for the same AD domain, you must sync all of them.
    1. To sync, run the following command on all domain controller:

    ```
    repadmin /syncall /adep
    ```

    > **Note:** If an IP address for the existing entry is changed, make sure that DNS cache is flushed at AD DCs.

6. Add the Mapped LDAP+ AD as Authorization Provider for a Storage Domain. For more information, see Map LDAP+AD as an Authorization Provider to a Storage Domain.
7. If your environment has multiple domain controllers for the same AD domain, you must sync all of them.
    1. To sync, run the following command on all domain controller:

    ```
    repadmin /syncall /adep
    ```

    > **Note:** If an IP address for the existing entry is changed, make sure that DNS cache is flushed at AD DCs.

8. Configure Unix client for mounting Kerberos based View: For more information, see Configure Unix client for Mounting Kerberos Based View
9. Mount Kerberos-based NFS View. For more information, see Mount a View Using NFS.

Map LDAP+AD as an Authorization Provider to a Storage Domain

To map LDAP+AD as an authorization provider to a storage domain, complete the following steps:

1. In the Cohesity UI, select **Settings > Summary** and select the **Storage Domains** tab.
2. On the **Storage Domains** page, complete one of the following actions:
    1. Edit an existing Storage Domain by clicking the Actions Menu (⋮) and selecting **Edit**.
3. In the **Authentication Provider** drop-down, select required LDAP +AD.
4. Click **Save**.

Configure Unix client for Mounting Kerberos Based View

> **Note:** The commands in the steps are CentOS commands.

To configure a Unix client for mounting Kerberos-based View, complete the following steps:

1. Stop SELinux and disable the service.
2. Open the `/etc/selinux/config` file and set the SELINUX mode to be disabled.
3. Run the following commands and stop the firewall and disable the service.
    - `systemctl stop firewall`
4. Reboot the Unix client node.
5. Open `/etc/hosts` and change the hostname of the node.

    For example, consider **Kerb245** is the hostname of the client and **w2k16-1.test.local** is the domain name

    "198.51.100.12 kerb245.w2k16-1.test.local kerb245"
6. Change the hostname using the following command.
    - `hostnamectl set-hostname`
7. Install all the required packages for `realm/krb5`.
    1. Run the following command
        - `yum update;`
        - `yum install realmd sssd oddjob oddjob-mkhomedir adcli samba-common samba-common-tools ntpdate ntp`
        - systemctl enable ntpd.service
    2. realm discover w2k16-1.qa.local

3. realm join w2k16-1.qa.local --user=Administrator@W2K16-1.QA.LOCAL --verbose
4. realm list

8. Starting from Cohesity Data Cloud (Self-managed) 6.6.0a onwards, you need not create a user for each client machine that you intend to mount from. Instead, you can map all clients to the default (nobody) hostname of the client. You must create a **nobody** user in the AD server and configure the attributes as described in NFS User and Group Configuration Parameters in AD.

   > **Note:** If the Cohesity Data Cloud (Self-managed) version is 6.6, you must create a user and ensure that the username is the same as the hostname of the client.

   If you create a user, ensure that the username is the same as the hostname of the client. For example, if the hostname of the client is **krb245**, you should configure the attributes as follows:

   - Username—*krb245*
   - objectClass—posixAccount
   - Uid—*krb245$*
   - UidNumber—User identifier. The value can be any integer.
   - GidNumber—Group identifier of the group this user is a part of. This group must have the value **posixGroup** in the objectClass attribute.

   > **Note:** For the AD server user, UID number and UID value are different fields.

9. Create a group on the AD server and add the user to the group.

   Once you are done with configuring the Unix client, you can mount a NFS View using Kerberos.

> **Note:** If you are using an AD user from a Linux client, ensure to set the following parameters on the client:

For AD users to work, `/etc/sssd/sssd.conf` file should have the following values:

```
ldap_id_mapping = False
use_fully_qualified_names = False
fallback_homedir = /home/%u
```

To mount a Kerberos-based NFS View, see Mount a View Using NFS.

MIT Kerberos

MIT Kerberos is another flavor of Kerberos similar to AD Kerberos. You need to add the MIT Kerberos provider to the Cohesity cluster from the Access Management page and configure. For more information, see MIT Kerberos.

> **Note:** MIT Kerberos does not support NTFS View type.

Prerequisites for Configuring MIT Kerberos on NFS Client

The following are the prerequisites for NFS client to configure Kerberos:

- Have DNS entries for all servers involved in the Kerberos environment.
- Ensure that the NFS users of client machines should have corresponding accounts in both LDAP and KDC.
- Host principal name for clients and servers.
- Configure KDC.
- Configure OpenLDAP server.
- Prepare a Client-Machine as:
  - OpenLDAP client: Validate client-ops for OpenLdap.
  - KDC-client: Validate client-ops for KDC.
  - NFS-client.

Configuring MIT Kerberos on a Cohesity Cluster

> **Note:** Use iris_cli cluster add-etc-hosts if servers are not DNS routable from Cohesity cluster.

Configuring MIT Kerberos on a Cohesity Cluster involves the following high-level steps:

1. Add OpenLDAP server to Cohesity cluster. For more information, see Add an LDAP Provider to the Cohesity cluster.
2. Add MIT Kerberos Provider to Cohesity Cluster. Adding MIT Kerberos Provider procedure allows you to map it to the LDAP.
3. Map MIT Kerberos + LDAP as an Authorization Provider to a Storage Domain.
4. Optional: Create a View (if not already exists), with kerberos security modes enabled (default on).
5. Log in as one of the LDAP-user, or change the user;

1. Get user Ticket Granting Ticket (TGT). If not acquired on login: Run '`kinit`
6. Access the mounted directory.

Add MIT Kerberos Provider to Cohesity Cluster

You need to add the MIT Kerberos provider to the Cohesity cluster to authenticate Unix users to access the Cohesity view in a Kerberos environment.

Prerequisites

- Realm name
- KDC server name and credentials

To add MIT Kerberos Provider to the Cohesity cluster, complete the following steps:

1. Click **Settings > Access Management > Kerberos**
2. In the Kerberos tab, click **Add Kerberos Provider**.
3. In the Add Kerberos Provider screen, complete the following steps:
    1. Type **Realm Name**.
    2. Type **Host Alias**.

        > **Note:** Host alias is Cohesity cluster names. Default host alias is your Cohesity cluster name. You can add more than one host alias names.

    3. (Optional) Toggle the **Use existing Host Aliases**.
    4. Type **KDC Servers** names.

        > **Note:** You can add multiple KDC Server host names. Use a comma to separate the names.

    5. Add **Admin Server** (KDC).

        > **Note:** You can select one of the added KDC servers as an admin server.

    6. Type **Admin password**.
    7. Select Mapped LDAP Provider.

        You can add the LDAP provider if it is not added already.
    8. Click **Add**.

Map MIT Kerberos + LDAP as an Authorization Provider to a Storage Domain

After adding the MIT Kerberos to the Cohesity cluster that is mapped to an LDAP server, you must map MIT Kerberos + LDAP server to a storage domain. The Views created under the storage domain mapped with MIT Kerberos + LDAP uses the MIT Kerberos for authentication.

To map MIT Kerberos Provider with a Storage Domain, complete the following:

1. In the Cohesity UI, select **Settings > Summary** and select the **Storage Domains** tab.
2. On the Storage Domains page, complete one of the following actions:
    1. Edit an existing Storage Domain by clicking ⋮ and select **Edit**.
3. In the **Authentication Provider** drop-down, select the required **MIT Kerberos +LDAP** in the **Kerberos + LDAP** section.
4. Click **Update Storage Domain**.

Edit MIT Kerberos Provider

You can edit the MIT Kerberos provider to update any modifications made in the Kerberos server.

To edit the MIT Kerberos provider, complete the following steps:

1. Click **Settings > Access Management > Kerberos**.
2. In the **Kerberos** tab, select the required **MIT Kerberos Provider** and select the edit icon for the required **MIT Kerberos**.
3. In the **Edit Kerberos Provider** screen, you can update the following:
    - Host Alias. Host alias is Cohesity cluster names.
    - (Optional) Toggle the Use existing Host Aliases
    - KDC servers names
    - Admin Server (KDC). You can select one of the KDC servers as an admin server.
    - Mapped LDAP provider.
4. Click **Save**.

Unmap MIT Kerberos Provider from Storage Domain

You first need to unmap the MIT Kerberos provider from the storage domain when you want to stop using it for authentication and delete it from the Cohesity cluster. If you do not unmap Views from the storage domain, the existing clients observe errors.

To unmap MIT Kerberos provider from the Storage Domain, complete the following:

1. In the Cohesity UI, select **Settings > Summary** and select the **Storage Domains** tab.
2. On the Storage Domains page, complete one of the following actions:
    1. Edit an existing Storage Domain by clicking ⋮ and select **Edit**.
3. In the **Authentication Provider** dropdown, select **None** in the Kerberos + LDAP section.
4. Click **Save**.

Delete MIT Kerberos Provider

When you no longer want to use the MIT Kerberos for authentication, you may delete. You can delete the MIT Kerberos Provider once you unmap from all storage domains.

To delete the MIT Kerberos from the Cohesity cluster, you must have MIT Kerberos password.

Prerequisite

You must unmap MIT Kerberos from the storage domain before deleting it.

To delete the MIT Kerberos Provider, complete the following:

1. Click **Settings > Access Management > Kerberos**.
2. In the **Kerberos** tab, select the required **MIT Kerberos Provider** and click the delete icon for the required **MIT Kerberos**.
3. In the **Delete MIT Kerberos** screen, type the **MIT Kerberos** admin password.
4. Click **Delete**.

Secure Login

Secure Login provides additional layers of security that enable the verification of the identity of a user, restricting unauthorized access and malicious activity in the cluster environment. Secure Login includes,

- Multi-factor Authentication (MFA)
- Split Key

For enhanced security, Cohesity strongly recommends that you enable MFA for all users and groups. After MFA is enabled by the administrator, users are prompted for an additional authentication factor along with the Cohesity cluster user credentials during the sign-in process.

> **Note:** A user must have a user account in the Cohesity Data Cloud (Self-managed) to use MFA provided by Cohesity.

Cohesity supports multi-factor authentication for the following types of users and groups:

- Local user
- Active Directory users with user accounts created in Cohesity
- SSO users and groups
- Support user

Cohesity supports split key for the following types of users:

- Cohesity Console user
- Support user

> **Important:**
> If you are using SSO for authentication, then Cohesity recommends that you enable MFA in your SSO system.

By default, MFA is enabled for Active Directory users with user accounts created in the Cohesity. You can enable MFA for local users and the **support** user account in the Cohesity. For more information, see:

- Multi-factor Authentication for Local and Active Directory Users
- Multi-factor Authentication for the Support User Account
- Enable Split Key

Multi-factor Authentication for Local and Active Directory Users

Multi-factor Authentication (MFA) is an additional layer of security used to verify the identity of a user. Cohesity supports MFA for local users. MFA enforces users to provide an additional authentication factor to access the Cohesity GUI or Cohesity CLI.

By default, MFA is enabled for Active Directory users with user accounts created in Cohesity. Administrators can enable MFA so that all or specific local users are prompted during the sign-in process for an additional authentication factor along with the local user credentials that they use to log in to the Cohesity cluster.

Cohesity supports the following authentication methods:

- **Authenticator App**—Users must enter the verification code generated by a TOTP authenticator app such as Okta Verify.

> **Note:** For Active Directory users, the Authenticator App is the only supported MFA method.

- **Email**—Users must enter the verification code sent to their email address.

After MFA is enabled by the administrator, an additional authentication factor is enforced for all users on their next login.

Points to Note

- Only administrators or users with the following privileges can enable or disable MFA for local users:
  - Users with the **Manage MFA** privilege can enable or disable MFA at the cluster level:

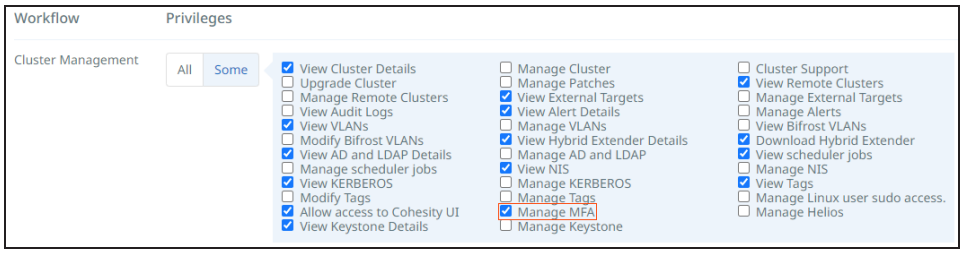

  - Users with the **Manage Users** privilege can enable or disable MFA at the user level:

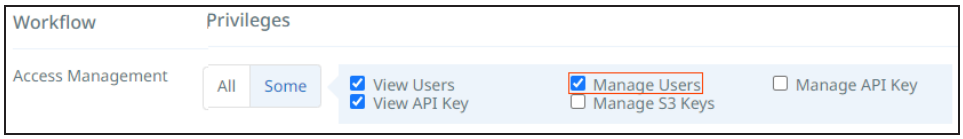

    For more information about users and roles, see Access Management.
- Users must download, install, and enroll in a TOTP authenticator app on their devices. See Supported TOTP Apps.
- Administrators must configure the email server on the Cohesity cluster. For more information about configuring an email server, see Email Configuration.
- Administrators must configure an email address for each user. Verification code is sent to this email address.
- Administrators can select one or both of the following authentication methods:
  - Authenticator app
  - Email
- Although administrators can select any authentication method as an additional factor, Cohesity strongly recommends that you:

    - Use the authenticator app as the primary MFA method
    - Do not use email as the primary MFA method because email messages can be intercepted and their contents potentially compromised by intermediary parties, especially when messages are routed over the internet.
- The lifetime of an email OTP code is two minutes.
- Disabling MFA for AD (Active Directory) users is not supported due to security reasons.

Supported TOTP Apps

Cohesity supports the following TOTP authenticator apps:

- Okta Verify
- Google Authenticator
- Duo Mobile
- Microsoft Authenticator

Enable MFA for Local Users

To enable MFA:

1. Navigate to the **MFA** tab using one of the following methods:
    - If MFA is not enabled on the Cohesity cluster, a banner appears at the top of the page; click **Manage Settings** to navigate to the **Secure Login** tab:

    - Go to **Settings > Access Management** and click the **Secure Login** tab.
    - Turn on the **MFA** toggle, select one or both authentication methods:

      > **Note:** Cohesity recommends that you select both authentication methods. If you select only one authentication method, users may get locked out if they do not have an alternative method to access the system.

        - **Authenticator App**—All users must configure an authenticator app and use a verification code generated by the app for verification.

          > **Note:** If you are using a TOTP authenticator app, the Cohesity cluster need not be connected to the internet.

- **Email**—The administrator must configure an email address for each user. The verification code is sent to this email address.

  > **Note:** If you are using email for authentication, the SMTP should be configured in the Cohesity cluster and should be active.

2. Click **Save**.

   MFA is enabled for all users. Navigate to **Settings > Access Management > Users & Groups** to view the **MFA** column.

   > **Note:** Administrators or users with the **Manage Users** privilege can turn on or turn off the toggle to enable or disable MFA for the corresponding user.

Set Up MFA for Local and Active Directory Users

After MFA is enabled, users are prompted to set up MFA when they log in for the first time.

> **Note:** For Active Directory users, the Authenticator App is the only supported MFA method.

To set up MFA through the Cohesity GUI, users must do the following:

1. Open a web browser and enter the hostname or IP address of the Cohesity cluster.
2. Enter the credentials and click **Sign In**.
3. Scan the QR code and click **Continue**:

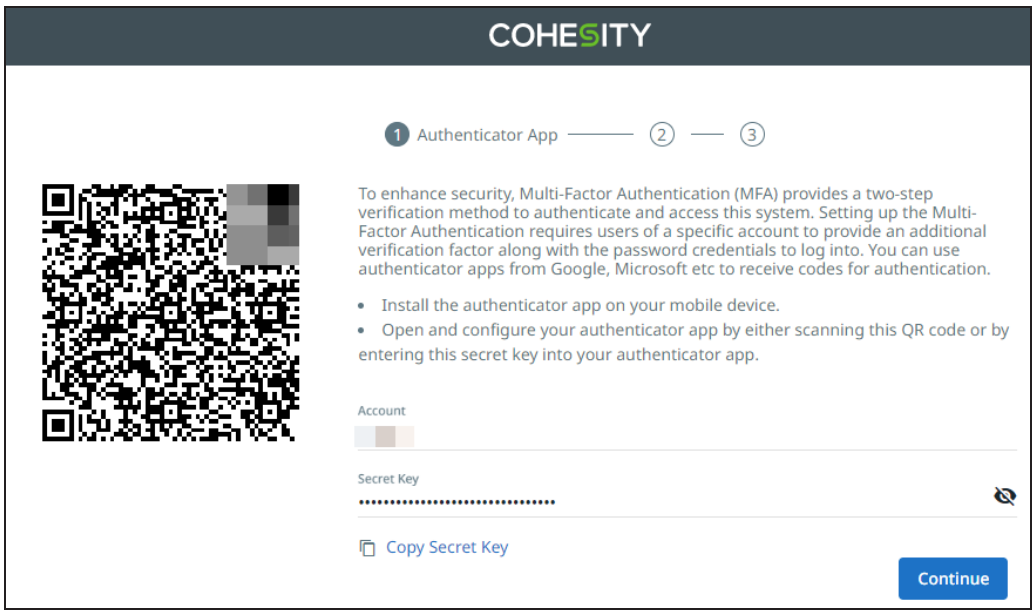

4. Enter the verification code generated by the authenticator app and click **Verify**.
5. A message is displayed after the code is successfully verified:
    - If the administrator has enabled only one authentication method, click **Complete** to finish the setup process and log in to the Cohesity cluster:

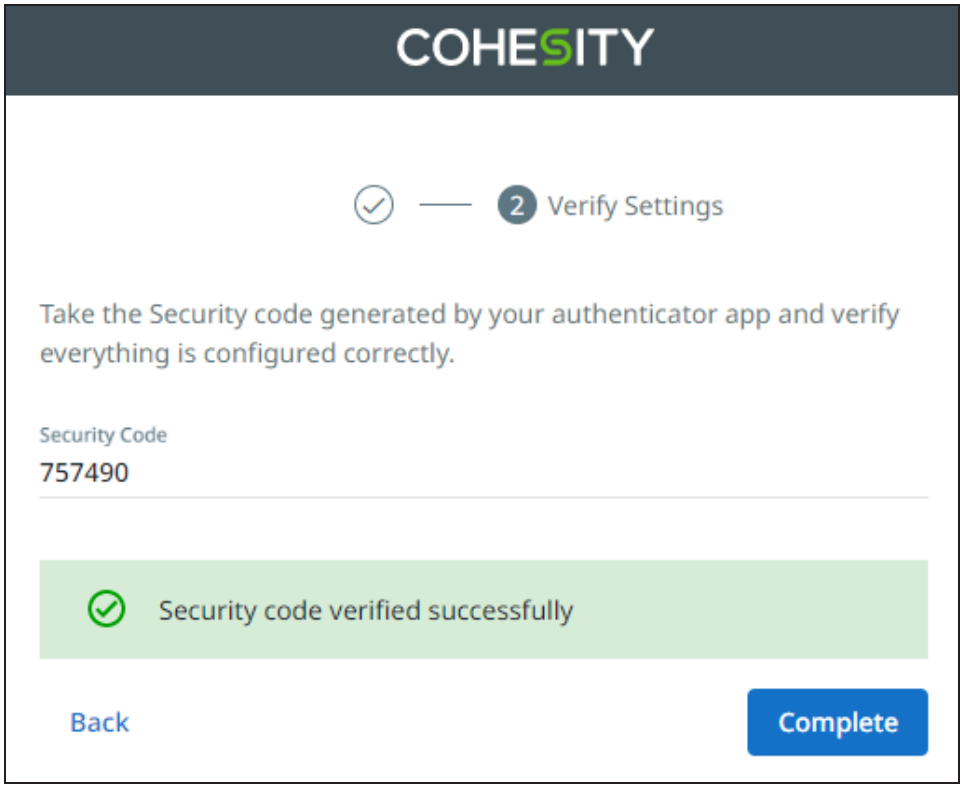

    - If the administrator has enabled both authentication methods:
        i. Click **Next** to verify the email authentication method.
        ii. Enter the verification code received on your email address and click **Verify**.
        iii. Click **Complete** to finish the setup process and log in to the Cohesity cluster:

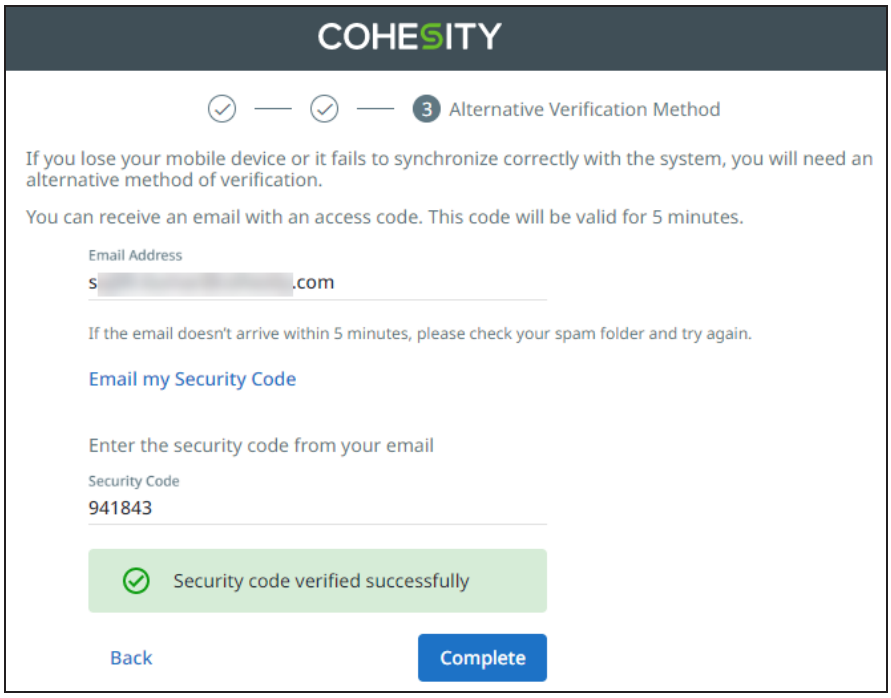

> **Note:** Users cannot update their email addresses. Users must contact the administrator to update the email address.

Log In to the Cohesity GUI After MFA is Enabled

After MFA is enabled, users are prompted to enter the verification code after entering their credentials.

> **Note:** If the administrator has enabled both authentication methods, the default authentication method used is the authenticator app.

To log in to the Cohesity GUI:

1. Open a web browser and enter the hostname or IP address of the Cohesity cluster.
2. Enter the username and password and click **Sign In**.
3. Enter the verification code from the authenticator app and click **Sign In** to log in to the Cohesity cluster.

If you have lost access to the device on which the authenticator app is installed, complete the following steps to log in to the Cohesity cluster:

1. Open a web browser and enter the hostname or IP address of the Cohesity cluster.
2. Enter the username and password and click **Sign In**.

3. Click **Use Alternative Verification Method**:

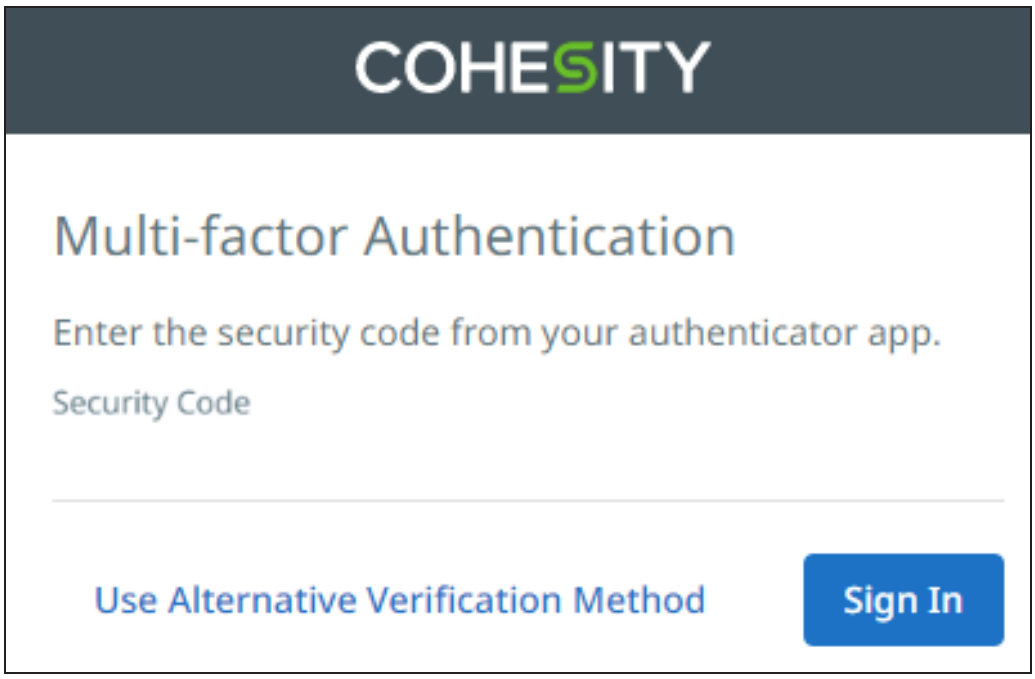

4. Click **Email my Security Code**:

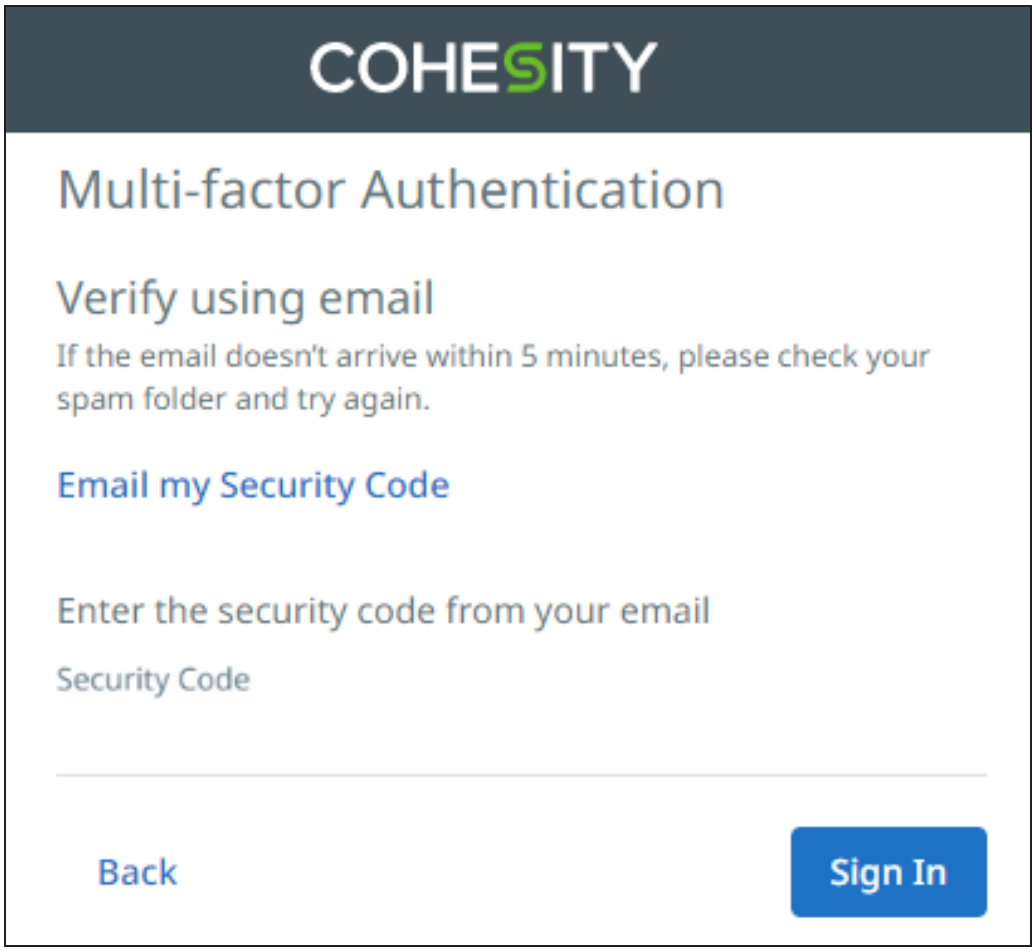

5. Enter the verification code received on your email address and click **Sign In**.

Log In to the Cohesity CLI After MFA is Enabled

After MFA is enabled through the Cohesity GUI, MFA is enabled for users logging in to the Cohesity CLI as well. Note the following important points:

- Enable the **support** user account on the Cohesity cluster. For more information, see Multi-factor Authentication for the Support User Account.
- Currently, Cohesity does not support setting up MFA through the CLI.
- If the administrator has enabled only the authenticator app, users must set up MFA through the Cohesity GUI. For more information, see Set Up MFA for Local and Active Directory Users.
- If the administrator has enabled only email-based TOTP, users need not set up MFA as the verification code is sent to their email address.

To log in to the Cohesity CLI:

1. Open the SSH terminal and execute the following command to log in to the Cohesity cluster Linux shell as a **support** user:

   ```
   ssh support@<ip address or hostname of the Cohesity cluster>
   ```

   Examples

   - IP address example: `ssh support@192.0.2.1`
   - Hostname example: `ssh support@cohesityhost`

2. Enter the following command:

   ```
   iris_cli
   ```

3. Enter the cluster username and password.
4. If the administrator has enabled both authentication methods, enter `Email` or `Totp` to select the authentication method that you want to use:

   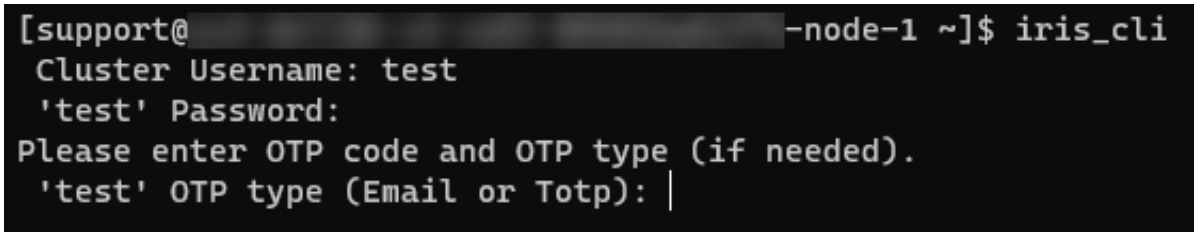

   - If you selected Email authentication, enter `Y` to trigger the email and enter the verification code received in your email:

     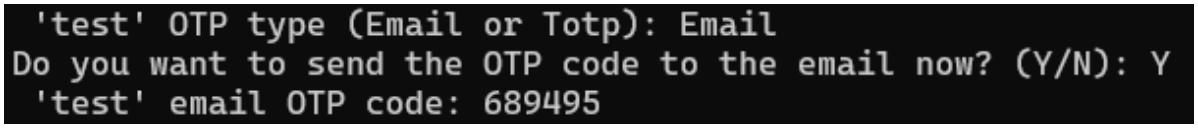

     > **Note:** You can generate the email verification code only once every two minutes.

   - If you selected TOTP authentication, enter the verification code generated by the authenticator app.

     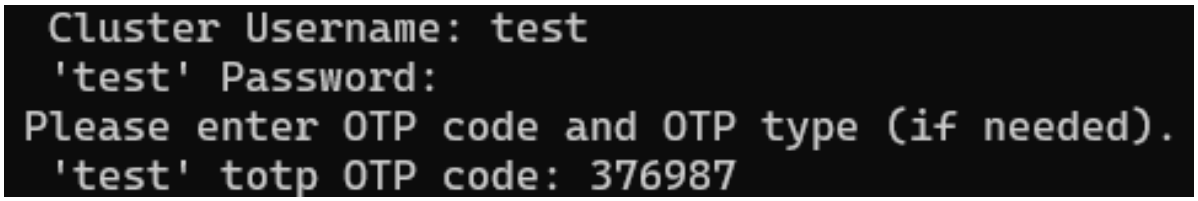

> **Note:** If you have not set up MFA, the following message is displayed:
>
> ```
> 'test' OTP type (Email or Totp): Totp
> The user has not setup Time-base One-time Password (TOTP). Please finish the TOTP setup via UI.
> ```
>
> You must set up MFA through the Cohesity GUI to continue. For more information, see Set Up MFA for Local and Active Directory Users.

5. After you enter the verification code, press **Enter** to log in to the Cohesity CLI.

Disable MFA for Local Users

Administrators or users with the requisite privileges can disable MFA for a user or all users. Cohesity recommends that you keep MFA enabled for enhanced security.

> **Note:**
> - If the administrator re-enables MFA, users must set up MFA again on their device. For more information on setting up MFA, see Set Up MFA for Local and Active Directory Users.
> - Disabling MFA for AD (Active Directory) users is not supported due to security reasons.

To disable MFA for a local user:

1. Navigate to **Settings > Access Management > Users & Groups**.
2. In the **MFA** column, turn off the toggle corresponding to the user.
3. Click **Disable**.

To disable MFA for all users:

1. Navigate to **Settings > Access Management** and click the **Secure Login** tab.
2. Turn off the **MFA** toggle and click **Disable**.
3. Click **Save**.

Frequently Asked Questions

1. How can I log in to the Cohesity cluster if I lose my device?

   If the authenticator app is the primary MFA method and if you lose your device, you can use the alternative verification method. For more information, see Click Use Alternative Verification Method:

   Cohesity recommends that administrators enable both authentication methods for local users:

- Authenticator app
- Email

If email is not enabled as an authentication method, users can contact the administrator to reset MFA for their account.

2. How do I connect the authenticator app to my new device?

   If you are changing your device, check whether the authenticator app allows transferring accounts to your new device. For example, Google Authenticator supports transferring accounts to your new device. For more information, see the **Transfer Google Authenticator codes to a new phone** section on this page.

   As a best practice, Cohesity recommends that you check the authenticator app that you are using and confirm whether you can transfer accounts to your new device as different authenticator apps handle this in different ways. Users can also contact the administrator to reset MFA for their account.

3. How can the administrator reset the MFA for a user?

   If the authenticator app does not support transferring accounts or if a user does not have access to their old device, the administrator or users with the **Manage Users** privilege can reset MFA for users.

   > **Note:** MFA can be disabled only for local users and not Active Directory user with user accounts created in the Cohesity.

   To reset MFA:

   a. Navigate to **Settings > Access Management > Users & Groups** tab.
   b. Turn off the toggle and disable MFA for the user.
   c. Turn on the toggle and reenable MFA for the user.
   d. The user must install and set up the authenticator app on their new device.

4. How can the administrator reset MFA for all local users?

   Administrators or users with the **Manage MFA** privilege can reset MFA for all local users. To reset MFA:

   1. Navigate to **Settings > Access Management > Secure Login** tab.
   2. Turn off the **MFA** toggle and disable MFA at the cluster level.
   3. Turn on the **MFA** toggle and re-enable MFA at the cluster level.

      MFA is reset for all local users. Users must set up MFA again. For more information, see Set Up MFA for Local and Active Directory Users.

5. What should I do if my email-based verification code expires?

If email is the primary MFA method and if there is a delay in receiving the email, the verification code expires. Users can request a new verification code. Cohesity recommends that you use the authenticator app method instead. An authenticator app can generate verification codes even if the device has no connectivity.

6. What should I do if I am locked out of my account?

   A login attempt may fail if users enter an incorrect password or verification code. User accounts are locked out after three consecutive failed attempts.

   If the user account is locked out because of an incorrect verification code, retry after one minute. If the user account is locked out because of an incorrect password, retry after 15 minutes or contact the administrator and request to unlock the user account.

   If an administrator is locked out, contact the peer administrator to reset MFA.

7. Can the default MFA be disabled for AD accounts?

   Yes, the default MFA can be disabled for the AD accounts through Cohesity CLI. Contact the Cohesity Support to disable the default MFA.

Related Topics

- Multi-factor Authentication for Organizations
- Multi-factor Authentication for the Support User Account

Multi-factor Authentication for the Support User Account

Multi-factor Authentication (MFA) is an additional layer of security used to verify the identity of a user. Cohesity supports MFA for the **support** user account. MFA enforces users to provide an additional authentication factor to access the Cohesity or Cohesity CLI.

Administrators can enable MFA for the **support** user account through Cohesity. For more information about the **support** user, see Manage Support User.

> **Warning:** MFA is not supported for the Cohesity console user accounts.

The administrator can select only one of the following authentication methods:

- Authenticator app
- Email

> **Note:** If you select **Email**, then you should establish a connection between Cohesity and SMTP server. For more information, see Configuring MFA for the support user to work with the SMTP server.

To enable MFA for the support user account:

1. Navigate to **Settings > Access Management > Support**.
2. In the **Support User** section:
    - Enter the email address of the **support** user and click **Update Email**.
    - To update the Linux Shell password, click **Update Password** and update the password.
3. Turn on the **Multi-factor Authentication** toggle, select one authentication method, and click **Update MFA**:
    - **Authenticator App**—The **support** user has to set up MFA and enter the verification code from the authenticator app to access the Cohesity cluster Linux shell through SSH.
    - **Email**—The **support** user has to enter the verification code received on the email address to access the Cohesity cluster Linux shell through SSH.
4. In the **Enable MFA** pop-up window that appears, enter the support user password and click **Authorize**.
    - Authenticator App – To enable MFA using an authenticator app, perform the following steps:
        a. Scan the QR code displayed on the screen using your authenticator app, or manually enter the account name and secret key.
        b. Click **Continue**.
        c. Enter the security code generated by the authenticator app and click **Verify**.

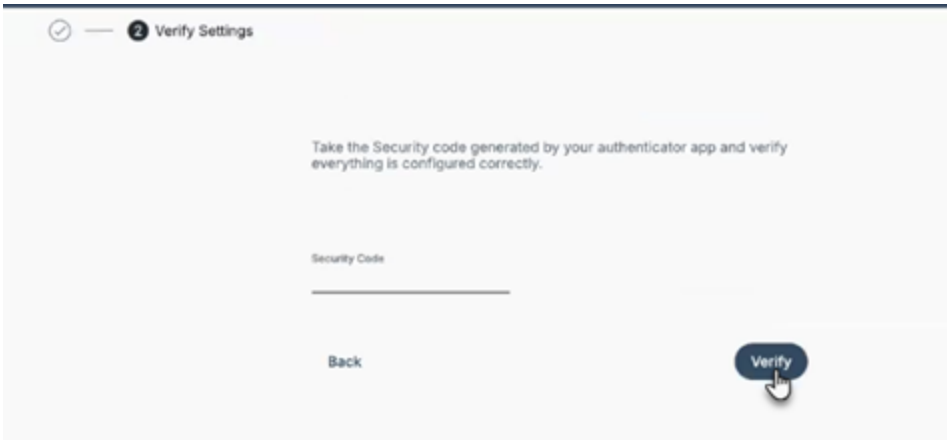

        d. Upon successful verification, a check mark appears next to the code. Click

Complete to enable MFA for the support user.

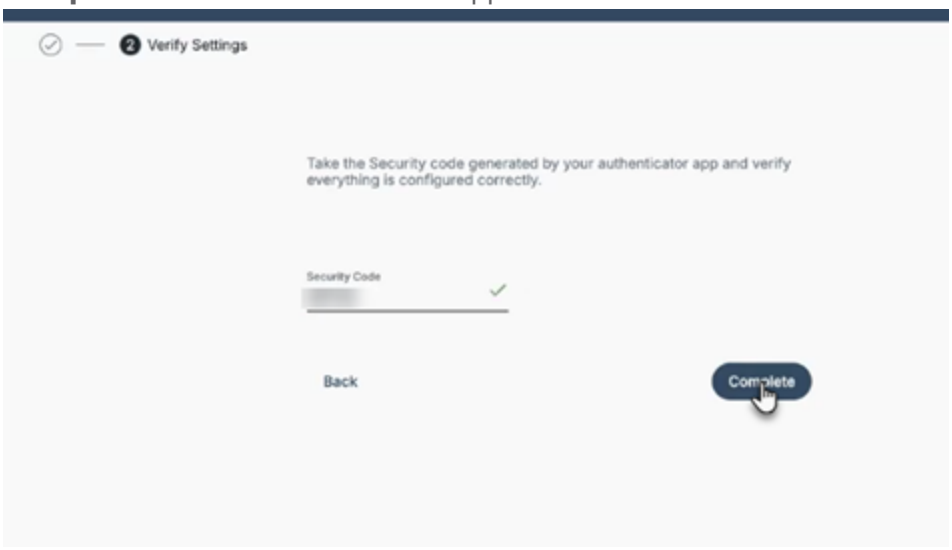

- Email – To enable MFA through email, perform the following steps:
  a. Enter the registered email address of the support user.
  b. Click **Send Email**. Cohesity sends a verification code to the email address provided.

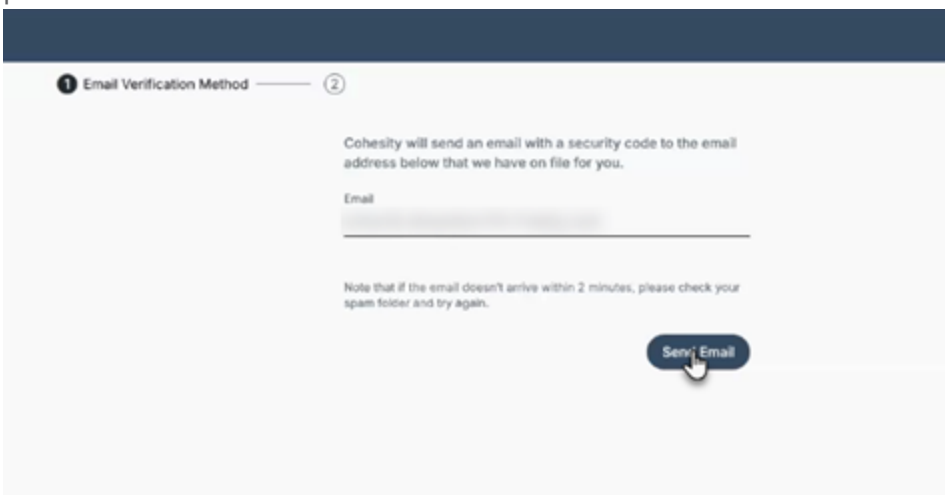

  c. Enter the security code you received and click **Verify**.
  d. Upon successful verification, a check mark appears next to the code. Click **Complete** to enable MFA for the support user.

Set Up MFA for the Support User Account

If MFA is enabled for the **support** user account, the **support** user is prompted to set up MFA during the subsequent login attempt.

To set up MFA through the Cohesity, the **support** user must complete the following steps:

1. Open a web browser and enter the hostname or IP address of the Cohesity cluster.
2. Enter the credentials and click **Sign In**.
3. Scan the QR code and click **Continue**:

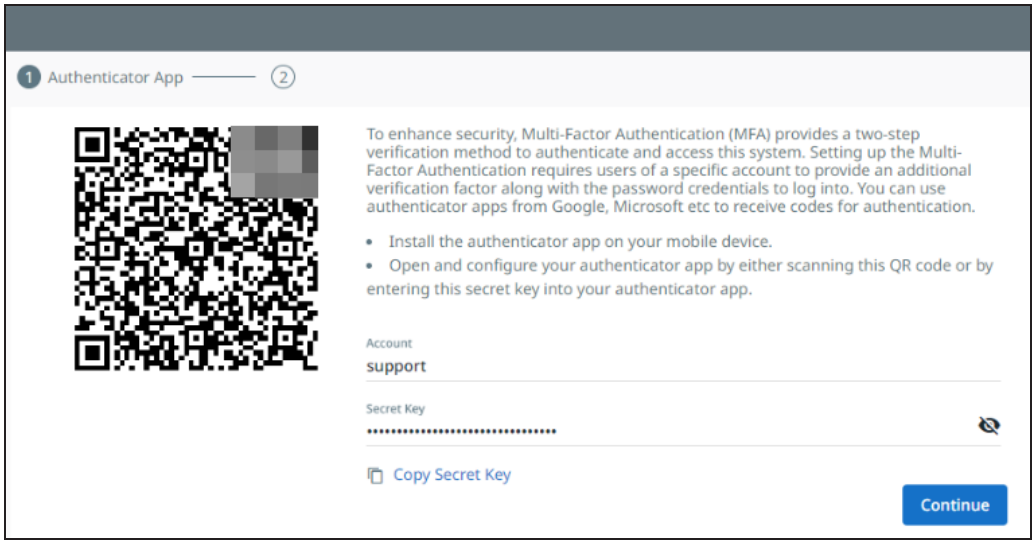

4. Enter the verification code generated by the authenticator app and click **Verify**.
5. After the code is successfully verified, click **Complete** to finish the setup process and log in to the Cohesity cluster:

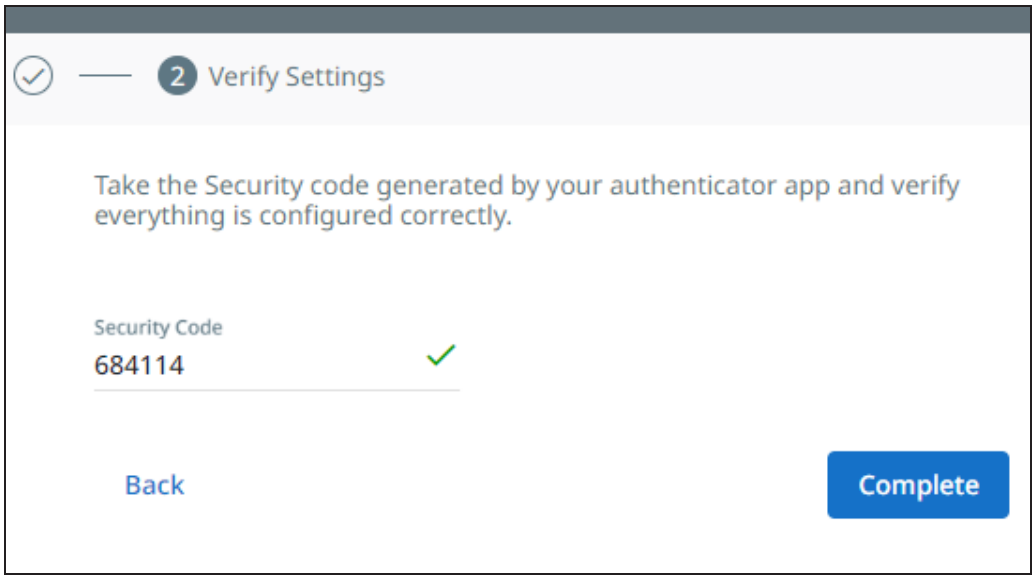

> **Note:** If the administrator has configured **Email** as the authentication method, the **support** user need not set up MFA as the verification code is sent to the email address.

Log In to the Cohesity Cluster Linux Shell

After MFA is enabled, the **support** user is prompted to enter the verification code after entering the password.

To log in to the Cohesity cluster Linux shell as a **support** user:

1. Open the SSH terminal and execute the following command:

   ```
   ssh support@<ip address or hostname of the Cohesity cluster>
   ```

   Examples

   - IP address example: `ssh support@192.0.2.1`
   - Hostname example: `ssh support@cohesityhost`
2. Enter the **support** user password.
3. Enter the verification code:
   - If the administrator has configured **Authenticator App** as the authentication method, enter the verification code generated by the authenticator app at the **Verification code** prompt:

     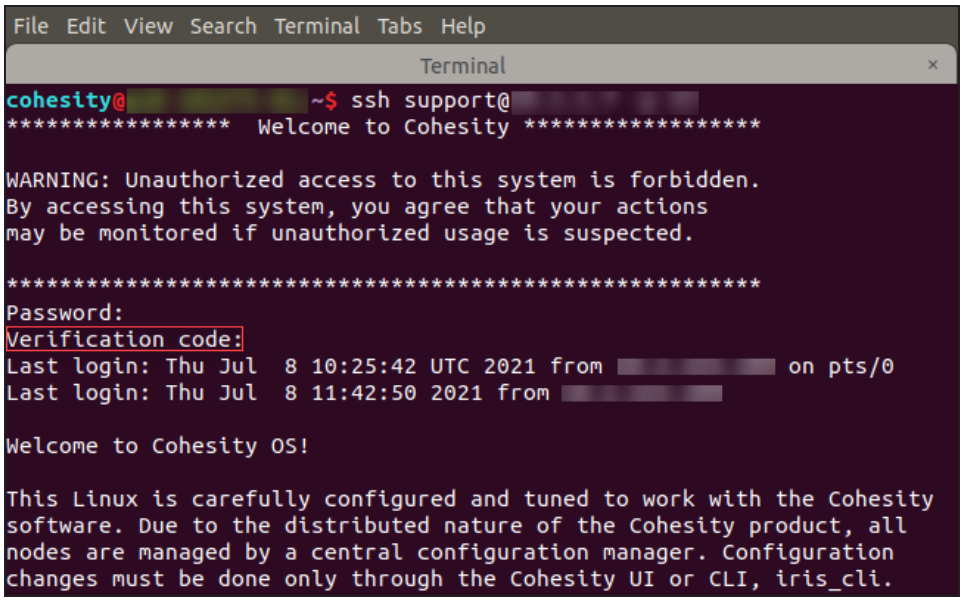

   - If the administrator has configured **Email** as the authentication method, enter the verification code received in the email at the **Enter the email based auth**

token prompt:

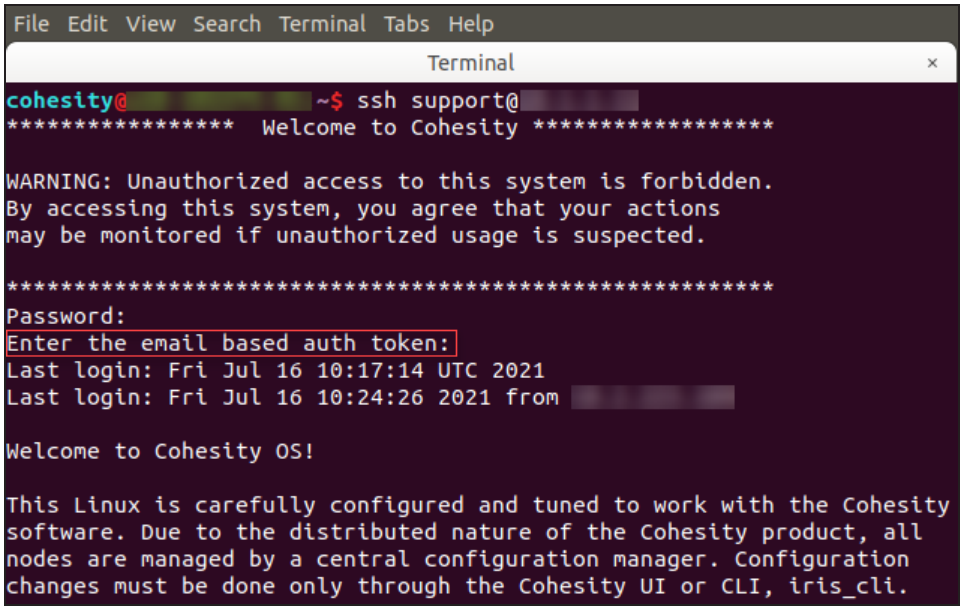

Disable MFA for the Support User Account

Administrators can disable MFA for the **support** user account. Cohesity recommends that you keep MFA enabled for enhanced security.

> **Note:** If the administrator reenables MFA, the **support** user must set up MFA again on the device. For more information on setting up MFA, see Set Up MFA for the Support User Account.

To disable MFA for the support user account:

1. Navigate to **Settings > Access Management > Support**.
2. In the **Support** tab,turn off the **Multi-factor Authentication** toggle.
3. In the **Disable MFA** pop-up window that appears, enter the support user password and the security code received via the authenticator app or email, depending on your MFA settings.

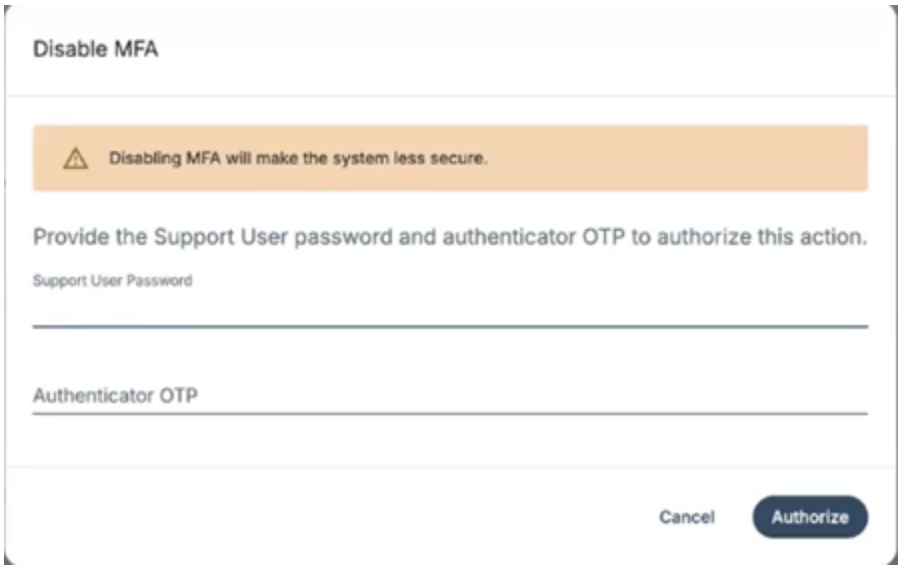

4. Click Authorize to disable MFA for the support user.

Related Topics

- Points to Note
- Supported TOTP Apps
- Frequently Asked Questions
- Multi-factor Authentication for Organizations

Enable Split Key

Split Key is an additional layer of security that protects the Cohesity host OS against any malicious activity. It restricts access to the host OS through multi-layer security that requires login credentials, and a challenge and response workflow. For more information, see Access Cohesity Host OS.

Follow these steps to enable Split Key on your cluster,

1. Navigate to **Settings > Access Management > Secure Login**.
2. In the **Secure Login** tab, turn on the **Split Key** toggle.
3. Make sure the **Split Key** enable confirmation screen appears.

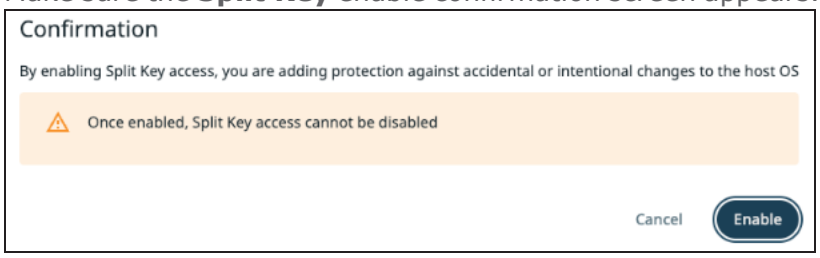

4. Click Enable to enable the Split Key.

> **Important:** Once you enable the Split Key, it cannot be disabled.

For clusters upgraded from 6.8.1 to 7.1, the UI offers the option to enable the Split Key.

By default, the Split Key is enabled for 7.1 fresh installations or upgrades from 7.0.1.

Certificates

The **Certificates** page lists all the certificates available in the certificate store. Certificates in this store are used to authenticate Cohesity to registered sources, enabling Cohesity to protect your workloads. You can generate a new certificate or import a certificate in the Personal Information Exchange (PFX) format. You can also add, edit, and delete the certificates.

> **Note:**
> - This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.
> - The Certificates page is used to manage the certificate store for workload authentication. To view or update the SSL certificate that secures HTTPS access to the cluster UI, see Web Server Certificates. For CLI-based certificate updates, see Manage SSL Certificates.

Generate Cohesity Managed Certificate

To generate a Cohesity Managed Certificate:

1. Navigate to **Settings** > **Access Management** > **Certificates**.
2. Click on the **Certificates** drop-down and select **Generate**.
3. In the **Generate Cohesity Managed Certificate** window, enter the following details:
   - **Given Name** – A name for the certificate that is easy to recall.

     This can be edited later.

- **Common Name (CN)** - Domain name or the entity for which the certificate is issued.
- **Cluster**/**Region** – Cluster/Region for which the certificate is generated. Only required when accessing from global mode or Data Protect mode.
- **Valid Till** – Validity of the certificate. The default validity is one year.
- **Organization Issued To** - *[Optional]* Organization to which the certificate is issued. The default value is **Cohesity Inc**.
- **Department Issued To** - *[Optional]* Department to which the certificate is issued. The default value is **Cohesity Inc**.

4. Click **Generate**.
5. The new certificate will be generated and displayed on the **Certificates** page.

Guidelines to Generate Custom Certificate using OpenSSL

To generate a custom certificate using OpenSSL:

1. To generate a private key:

   ```
   #openssl genrsa -out tls.key 2048
   ```

2. To generate a CSR:

   ```
   #openssl req -new -key tls.key -out tls.csr
   ```

3. To self sign the certificate:

   ```
   #openssl x509 -req -in tls.csr -signkey tls.key -out tls.crt -days 365
   ```

If you require a password protected certificate, then instead of command 4 run command 5.

4. To export the certificate and the private key to PFX:

   ```
   #openssl pkcs12 -export -out tls.pfx -inkey tls.key -in tls.crt -
   certfile ca-chain.crt
   ```

5. To generate a password protected certificate.

   ```
   #openssl pkcs12 -export -out tls.pfx -inkey tls.key -in tls.crt -
   certfile ca-chain.crt -passout pass:MyStrongPassword
   ```

If you already have the certificate and private key in PEM format and only require the PFX file, then only run the commands in steps 4 or 5.

Import Custom Certificate

> **Note:** Certificates only in the Personal Information Exchange (PFX) format are supported.

To import custom certificates:

1. Navigate to **Settings > Access Management > Certificates**.
2. Click on the **Certificates** drop-down and select Import.
3. In the **Import a Certificate** window, enter the following details:
    - **Given Name** – A name for the certificate. This can be edited later.
    - Drag and drop the PFX certificate or click and upload the certificate file.
    - **Password** – Password of the protected file.
4. Click **Import**.

The certificate will be imported and displayed on the Certificates page.

Download a Certificate

You can download a certificate from the certificate store.

1. In the Cohesity Dashboard, select **Settings > Access Management** and select the **Certificates** tab.
2. Click the actions icon (⋮) for the required certificate and click **Download**.

The certificate will be downloaded in the **.cer** format.

Edit a Certificate

You can edit the name of the certificate.

1. In the Cohesity Dashboard, select **Settings > Access Management** and select the **Certificates** tab.
2. Click the actions icon (⋮) for the required certificate and click **Edit**.
3. Update the name of the certificate as needed.
4. Click **Update**.

Delete a Certificate

You can delete a certificate if you no longer need it.

1. In the Cohesity Dashboard, select **Settings > Access Management** and select the **Certificates** tab.
2. Click the actions icon (⋮) for the required certificate and click **Delete**.
3. Respond to the confirmation prompt.
4. Click **Delete**.

Certificate Alert Threshold Configuration

Certificates generated by Cohesity are valid for one year by default. Default severity thresholds for non-auto-renewed certificates are 15% / 10% / 5% remaining validity (Info / Warning / Critical).

You can perform a retrieval of certificate inventory and configure alert thresholds using the API and CLI. A retrieval is faster than relying on scheduled inventory scans. It allows you to fetch certificate details on demand and proactively manage certificate expiry alerts.

The CLI commands internally invoke the same REST APIs for threshold configuration.

Run the following command from a node to enable the certificate_inventory_api_and_alerts feature:

```
iris_cli cluster update-feature-flag name=certificate_inventory_api_and_
alerts is-approved=true reason="enable cert inventory"
```

Ensure the threshold values meet the following requirements:

- Values must be in the range 1–50.
- criticalPercent < warningPercent < infoPercent.

View Certificate Inventory using API

Endpoints

```
GET /v2/cert-manager/cert-inventory
```

Perform the following steps to retrieve certificate inventory using API:

1. Authenticate using `POST /irisservices/api/v1/public/accessTokens`
2. Call `GET /v2/cert-manager/cert-inventory`

Parameters

Parameter	Type	Required	Description
maxStaleness	string	No	Controls freshness of certificate inventory data
nodeIp	string	No	Retrieves inventory for a specific node
pageSize	integer	No	Reserved for future batching/pagination

Configure Certificate Alert Threshold using CLI

To view the stored override configuration, run:

```
iris_cli cert-manager inventory-alert-threshold-get | jq .
```

You can update threshold configurations using either a JSON file or CLI flags.

To update thresholds using a JSON file, run:

```
iris_cli cert-manager inventory-alert-threshold-set threshold-json-file-
e=/tmp/threshold-by-cert.json | jq .
```

The JSON file supports the following structure:

```
{ "byCertId": { "<cert_id>": { "infoPercent": , "warningPercent": , "crit-
icalPercent": } }, "byCategory": { "": { "infoPercent": , "warningPercent":
, "criticalPercent": } }, "global": { "infoPercent": , "warningPercent": ,
"criticalPercent": } }
```

To update thresholds using flags, run:

```
iris_cli cert-manager inventory-alert-threshold-set
 global-info-percent=
 global-warning-percent=
 global-critical-percent= | jq .
```

Example:

```
iris_cli cert-manager inventory-alert-threshold-set
 global-info-percent=28
 global-warning-percent=17
 global-critical-percent=9 | jq .
```

Configure Certificate Alert Threshold using API

Endpoints

- `GET /v2/cert-manager/cert-inventory/alert-threshold`
- `PUT /v2/cert-manager/cert-inventory/alert-threshold`

Parameters

Parameter	Type	Required	Description
maxStaleness	string	No	Controls freshness of certificate inventory data
nodeIp	string	No	Retrieves inventory for a specific node
pageSize	integer	No	Reserved for future batching/pagination

Perform the following steps to retrieve stored threshold configuration using API:

1. Authenticate using `POST /irisservices/api/v1/public/accessTokens`
2. Call `GET /v2/cert-manager/cert-inventory/alert-threshold`

Perform the following steps to configure thresholds using API:

1. Authenticate using `POST /irisservices/api/v1/public/accessTokens`
2. Call `PUT /v2/cert-manager/cert-inventory/alert-threshold`

Request Body

The request body for `PUT /v2/cert-manager/cert-inventory/alert-threshold` should include a JSON object with the following structure:

```
{
  "byCertId": {
    "<cert_id>": {
      "infoPercent": <int>,
      "warningPercent": <int>,
      "criticalPercent": <int>
    }
  },
  "byCategory": {
    "<category>": {
      "infoPercent": <int>,
      "warningPercent": <int>,
      "criticalPercent": <int>
    }
  },
  "global": {
    "infoPercent": <int>,
    "warningPercent": <int>,
    "criticalPercent": <int>
  }
}
```

Request Body Schema

Field	Type	Required	Description
byCertId	Object	No	Map of certificate IDs to threshold configurations. If omitted, the existing byCertId configuration is preserved.
byCategory	Object	No	Map of certificate categories to threshold configurations. If omitted, the existing byCategory configuration is preserved.
global	Object	No	Global threshold configuration. If omitted, the existing global configuration is preserved.

Configure Account Security Settings

The Account Security page on a Cohesity cluster allows you to secure your Cohesity cluster through policy-driven user settings for managing passwords, user sessions, user inactivity, and locking the user account after multiple wrong password attempts. You can refer to the sections below to learn more about using and configuring these policies.

- View the Default Password Policy
- Define the Account Lockout Policy
- Define the Account Session Policy

View the Default Password Policy

On the Cohesity cluster, by default, a password policy is assigned to the local users. For added security, you can edit and customize the password policy based on your organization's requirements. It's important to change the passwords immediately as a response to ransomware attacks. This helps in minimizing the impact of the attack while keeping the data secured.

In the default password policy, you can update:

- Password strength
- Restrict the reuse of recently used passwords
- Set the password lifetime

From the Cohesity Dashboard, you can view the password policy that is applied by default to the local user account in the Cohesity cluster.

To view the default password policy:

1. In the Cohesity Dashboard, select **Settings** > **Account Security**.
2. From the **Password Policy** tab, you can view the configuration of the default password policy.

Update the Password Strength

The strength of a password depends on its length, complexity, and unpredictability. You can make the local user access to your Cohesity cluster more secure by changing the default password strength defined in the password policy.

To update the password strength:

1. In the Cohesity Dashboard, select **Settings** > **Account Security**.
2. On the **Password Policy** tab, click the **Edit** icon.
3. In the **Strength** section, based on your requirement, set the following values:
    1. In the **Minimum Length** field, enter the length of the password. By default, the length of the password is set to eight characters.

        > **Note:** Ensure that you select any one of the following options.

    2. Select **At least one uppercase letter** to mandate, including one uppercase letter in the password.

3. Select **At least one lowercase letter** to mandate, including one lowercase letter in the password.
4. Select **At least one number** to mandate, including one numerical digit in the password
5. Select **At least one special character** to mandate, including one special character in the password. For more details on what special characters are allowed in the local user password, see the Users and Groups section.

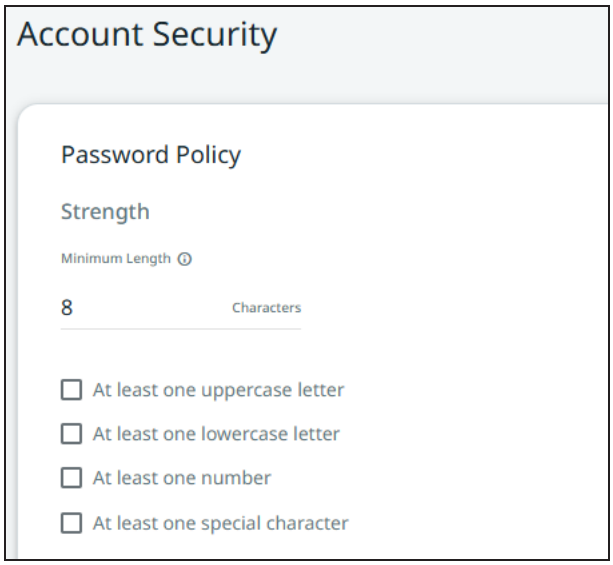

4. Click **Save**.

Restrict Reuse of Passwords

You can prevent local users from reusing a specified number of previous passwords. By default, the password policy is set to restrict five previously used passwords.

To update the password reuse limit:

1. In the Cohesity Dashboard, select **Settings** > **Account Security**.
2. On the **Password Policy** tab, click the **Edit** icon.
3. Under the **Reuse** section, do the following:
   - In the **Minimum Characters Different** field, set the minimum character that must be different than the characters used in the previous password.
   - In the **Minimum Reuse Difference** field, set the password reuse limit.
4. Click **Save**.

Set Password Lifetime

In the password policy, you can restrict the local user from resetting the password for a certain number of days. Also, you can set the local user password to expire after a specific number of days. By default, passwords are set to expire in 60 days. You can set the minimum and maximum number of days for the password lifetime.

The minimum number of days defines the number of days the local user is restricted to reset the password. For example, if the minimum number of days is set to five, the local user cannot reset his password for five days after the password is set for the local user.

The maximum number of days defines the number of days the local user password is valid. After the defined number of days, the password expires, and the local user is automatically prompted to change the password.

To set the password lifetime:

1. In the Cohesity Dashboard, select **Settings** > **Account Security**.
2. On the **Password Policy** tab, click the **Edit** icon.
3. In the **Lifetime** section, based on your requirement, set the following values:
    1. In the **Minimum Password Reset Duration** field, enter the minimum number of days to restrict the local user from resetting the password.
    2. In the **Password Expiry** field, enter the maximum number of days after which the local user password must expire.
4. Click **Save**.

> **Note:** The password lifetime can be set to a maximum value of 365.

Define the Account Lockout Policy

The Account Lockout policy in your password policy enables you to automatically lock the local user account if an attempt has been made to brute-force a local user password.

> **Note:** Users with the admin role are exempt from lockouts even after the failed login attempts.

In the Account Lockout policy, you can define:

- The number of consecutive failed login attempts causes a local user account to be locked.
- The time duration within which the consecutive failed login attempts cause a local user account to be locked.
- The inactivity timeout period for the user account. If the local user account is inactive for the configured inactivity timeout period, the local user account is automatically locked.

Once a user account is locked, the System Admin can unlock the account from the **Settings** > **Account Security** page in the Cohesity Dashboard.

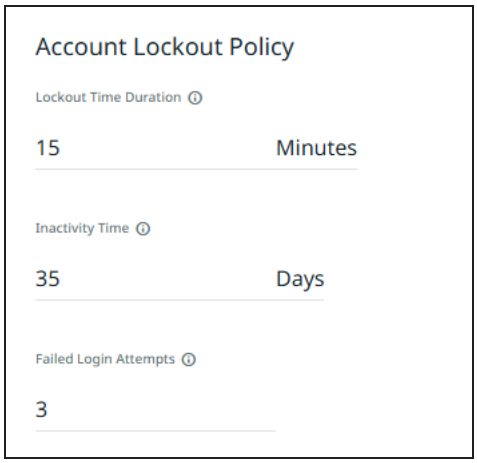

To define the account lockout policy:

1. In the Cohesity Dashboard, select **Settings** > **Account Security**.
2. On the **Password Policy** tab, click the **Edit** icon.
3. In the **Account Lockout Policy** section, based on your requirement, set the following values:
    1. In the **Lockout Time Duration** field, enter the time duration within which the consecutive failed login attempts must cause a local user account to be locked. By default, the time duration is set to 15 minutes.
    2. In the **Inactivity Time** field, enter the inactivity timeout period for the user account after which the account is automatically locked. By default, the inactivity timeout period is set to 35 days.
    3. In the **Failed Login Attempts** field, enter the number of consecutive failed login attempts that must cause a local user account to be locked. By default, the number of failed attempts is set to 3.
4. Click **Save**.

Define the Account Session Policy

The Session policy enables you to define the session timeout for various interfaces such as UI, CLI, and SSH.

In the Account Session policy, you can define:

- The maximum time limit after which the user is logged out from Cohesity UI and Cohesity CLI.
- The time limit after which any user on the Cohesity cluster is logged out from Cohesity Dashboard.
- The time limit after which any user on the Cohesity cluster is logged out when the user accesses the Cohesity through an SSH tunnel.

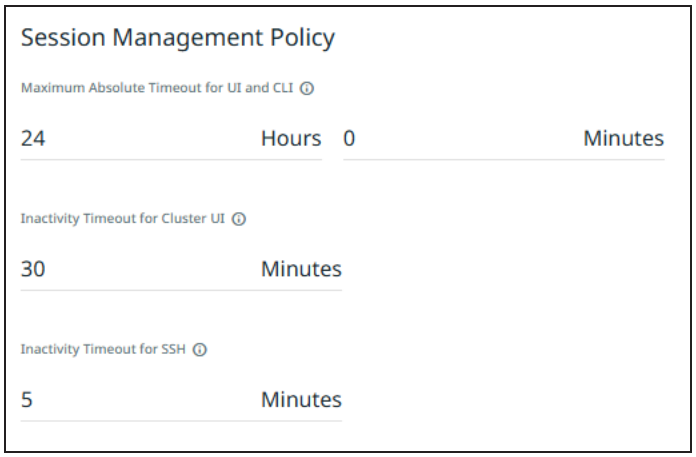

To define the account session policy:

1. In the Cohesity Dashboard, select **Settings** > **Account Security**.
2. On the **Password Policy** tab, click the **Edit** icon.
3. In the **Session Management Policy** section, based on your requirement, set the following values:
    1. In the **Maximum Absolute Timeout for UI and CLI**, enter the maximum time duration for UI and CLI after which the user account will be logged out irrespective of the status of the session. The default value is set to 1440 minutes, i.e. 24 hours. The minimum value can be 1 hour and the maximum value can be 24 hours.
    2. In the **Inactivity Timeout for Cluster UI** field, enter the inactivity time limit on Cluster UI for the user account to get logged out. The default value is set to 30 minutes. The minimum value can be 1 minute and the maximum value can be 12 hours.
    3. In the **Inactivity Timeout for SSH**, enter the inactivity timeout limit on SSH sessions for the user account to get logged out. The default value is set to 5 minutes. The minimum value can be 1 minute and the maximum value can be 24 hours.
4. Click **Save**.

Best Practices for Managing Users

This page provides best practices for managing and administering various user accounts, roles, and access to Cohesity products. You can access the system through:

- Cohesity Data Cloud (SaaS)
- Local Cohesity cluster

This page does not provide security-level recommendations. You must refer to the Cohesity Security Hardening Guide and configure your Cohesity products to meet your security requirements.

About MyCohesity

MyCohesity is a secure, easy-to-access portal, giving users fast, centralized access to the following Cohesity resources:

- **Helios**—Access all your data management services, whether self-managed or SaaS, through our one simple UI.
  - The Helios Super Admin can grant access to other MyCohesity users from the same organization. For more information, see this knowledge base article.
  - If you are unable to access Cohesity Data Cloud, see Why can't I access Helios from MyCohesity?
- **Support**—Manage support tickets and download software updates. Connect with our support resources for expert insights.
- **Product Documentation**—Access the latest product documentation to make the most of Cohesity for your business.
- **Academy**—Build your next-generation data management skill set to keep learning with our training, accreditation, and certification programs for sales and technical professionals.
- **User Group**—Learn from and connect with other local IT professionals, and enjoy exclusive access to Cohesity resources.
- **Galaxy Community**—Be a part of our vibrant global customer community to network with your peers.

For more information about MyCohesity, review this page.

Manage Helios (SaaS) Users

Cohesity Data Cloud is a next-gen data management platform that provides a comprehensive range of data services. You can log in to the SaaS-based management portal and manage multiple Cohesity clusters. This is a service managed by Cohesity and Cohesity Support grants the initial access. Go through the following subsections:

- Understanding Helios Users
- Best Practices for Managing Helios Users
- Shared User Accounts
- Enable Multi-factor Authentication for Cohesity Data Cloud Users
- SSO Usage Notes
- Summary

Understanding Helios Users

When you request access to the Cohesity Data Cloud platform, an initial user account is created. To get Super Admin access, the user has to:

1. Self-register and sign up for a MyCohesity account.
2. Log in to MyCohesity.
3. On the **Helios** tile, click **Launch**:

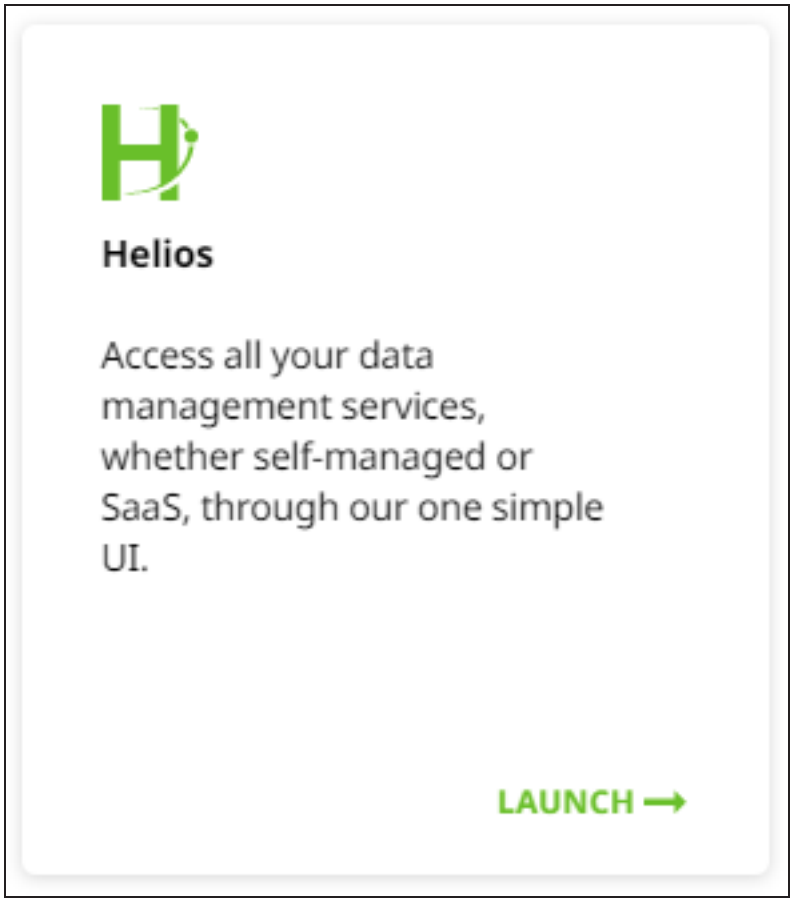

The user is granted Super Admin access.

> **Important:** If there are multiple users in the same domain, the first user who clicks the **Helios** tile gets Cohesity Data Cloud Super Admin access.

To keep your digital and cloud experience safer and compliant to your business needs, Cohesity supports Multi-Factor Authentication (MFA) for your MyCohesity account. For more information, see this knowledge base article.

> **Note:** By default, MFA is not enabled for this initial Super User account.

Best Practices for Managing Helios Users

Firstly, Cohesity recommends that you decide the primary authentication mechanism you wish to use. You have two options:

- Helios (SaaS) authentication
- Your own SSO provider

  For more information about configuring Cohesity Data Cloud to use an Identity Provider (IdP), see Add a Single Sign-on Provider.

Cohesity does not recommend one over the other. However, regardless of which authentication mechanism you choose, there are a few things you should do right away to avoid any access issues in the future:

- Create custom roles to provide access to individuals/groups in your organization. Cohesity recommends following the principle of least privilege.
- Cohesity recommends that you create a custom role with privileges equivalent to the built-in Super Admin role as you cannot modify the privileges of the built-in account. For more information about roles and privileges, see Custom Roles and Privileges.
- You must have multiple system administrator level users. Cohesity recommends that you have at least three users assigned with the built-in Super Admin role or custom role that you created with equivalent permissions.

  > **Note:** Cohesity recommends assigning the custom role.

  > **Warning:** If you have only one user with the Super Admin role or the equivalent custom role and if this user leaves your organization, you may experience issues, so ensure that you create additional users with the Super Admin role or the equivalent custom role.

- Enable MFA for all users:
  - To enable MFA for Cohesity Data Cloud local accounts, you should follow the steps mentioned in Multi-factor Authentication.
  - To enable MFA for SSO users, you must set up MFA in your SSO provider's system. For example, if you use Okta as the Identity Provider, you must set up and enable MFA in Okta.

Shared User Accounts

Cohesity does not recommend using a shared user account for any administration of our platform. If you use a shared user account, you will not be able to determine who made changes within your Cohesity Data Cloud account as audit logs are recorded for actions performed at the user account level.

> **Warning:** A shared user account is considered insecure and Cohesity strongly recommends you to not use a shared user account.

If you decide to use the shared user account method, Cohesity strongly recommends that you enable MFA by following the steps outlined below.

Enable Multi-factor Authentication for Cohesity Data Cloud Users

Multi-factor Authentication (MFA) provides a two-step verification method to authenticate and access Cohesity Data Cloud (SaaS). Setting up the MFA requires users to provide an additional verification factor along with the password credentials to log in to Cohesity Data Cloud. You can use codes from any authenticator app to log into Cohesity Data Cloud. After you enable MFA, access to your Cohesity Data Cloud is secured. MFA prevents unauthorized access and protects your data from ransomware attacks. For more information, see Multi-factor Authentication.

SSO Usage Notes

If you enable SSO to maintain your organization's users within Cohesity Data Cloud, then Cohesity recommends that you maintain two native accounts in case of an outage or impact to your SSO provider.

If you choose to use your own SSO provider, after all SSO users are created, you can remove the initial Helios/MyCohesity account from the Cohesity Data Cloud portal. Subsequently, only users from your SSO organization can access the Cohesity Data Cloud portal. However, note that SSO users cannot access MyCohesity, so you need MyCohesity user accounts to manage support tickets and download software updates.

If you need additional users strictly for MyCohesity, you can create them in the Cohesity Data Cloud portal and once the user logs in for the first time, you can remove the application access. These users would still maintain MyCohesity access. Cohesity strongly recommends that you enable MFA for MyCohesity user accounts as well. For more information, see this knowledge base article.

Summary

- If you have followed the guidelines outlined in the Manage Helios (SaaS) Users section, you should have multiple Cohesity Data Cloud and MyCohesity users with elevated permissions to perform additional functions.
- Create custom roles based on your operational requirements. You can then create additional users and assign custom role(s) based on your requirement. For more information about roles and privileges, see Custom Roles and Privileges.
- Periodically evaluate your three elevated users to ensure that as people change roles, you are maintaining a minimum of three active users with elevated permissions.

Manage Users Locally on the Cohesity Cluster

You can directly access the Cohesity cluster UI through the Virtual IP address (VIP) or the FQDN associated with the VIP. This management plane is out-of-band from the Cohesity Data Cloud application, so you can access the cluster even if the cluster is not connected to it. Go through the following subsections:

- About Cohesity Cluster Local Users
- Best Practices for Managing Local Cohesity Cluster Users
- Important Points to Note Regarding Initial Local Admin User
- Summary

About Cohesity Cluster Local Users

You can add local users to a Cohesity cluster. By default, when a new cluster is created, default built-in local users and groups are also created. The first user created is assigned the built-in Local Admin role. You have multiple options to add users in the Cohesity cluster:

- Local users and groups (GUI and iris_cli)
- SSO users (GUI only)
- Active Directory (GUI and iris_cli)
- Generic LDAP (GUI and iris_cli)

Local users have direct access to the cluster if the Cohesity Data Cloud application is unavailable for cloud-based multi-cluster management. The **support** user can be used to access the bash shell of the Cohesity cluster using SSH.

Cohesity does not recommend one over the other. However, regardless of which authentication mechanism you choose, there are a few things you should do right away to avoid any access issues in the future.

SSO users cannot execute `iris_cli` commands. If you decide to use SSO on the local cluster and if you want to avoid using the initial admin user and local user accounts, you can import users from Active Directory or any other domain for executing `iris_cli` commands.

Best Practices for Managing Local Cohesity Cluster Users

- Create custom roles to provide access to individuals/groups in your organization. Cohesity recommends following the principle of least privilege.
- Cohesity recommends that you create a custom role with privileges equivalent to the built-in Admin role as you cannot modify the privileges of the built-in account.
- You must have multiple system administrator level users. Cohesity recommends that you have at least three users assigned with the built-in Admin role or custom role that you created with equivalent permissions.

> **Note:** Cohesity recommends assigning the custom role.

> **Warning:** If you have only one user with the Admin role or the equivalent custom role and if this user leaves your organization, you may experience issues, so ensure that you create additional users with the Admin role or the equivalent custom role.

- Leverage named local users such as Active Directory users so that you can track `iris_cli` usage through audit logs.
- Enable MFA for all users:
  - To enable MFA for SSO users, you must set up MFA in your SSO provider's system. For example, if you use Okta as the Identity Provider, you must set up and enable MFA in Okta.

Important Points to Note Regarding Initial Local Admin User

Currently, you cannot remove or disable the initial local user. Therefore, Cohesity recommends that you edit the local user and add a shared email address that is a distribution list:

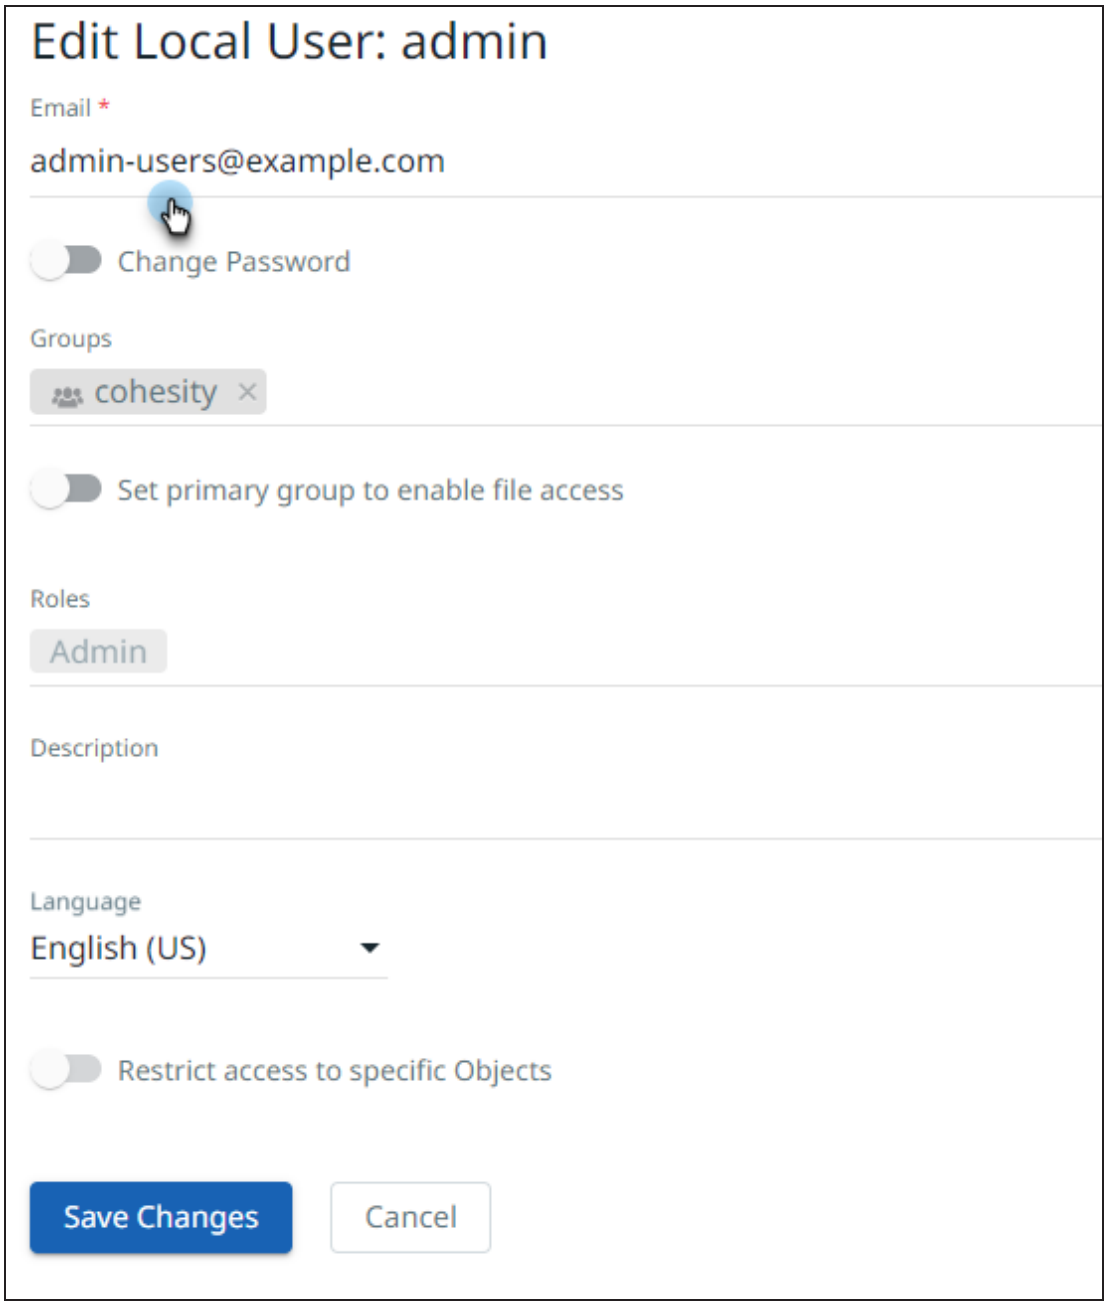

> **Note:** Even if you add a distribution list as an email address, you can still enable MFA for this local user.

To enable MFA for the local admin user, complete the following steps:

1. Log in to the Cohesity cluster GUI as the admin user.
2. Enable SMTP Server as you will be enabling email for authentication.
3. Turn on the **MFA** toggle:

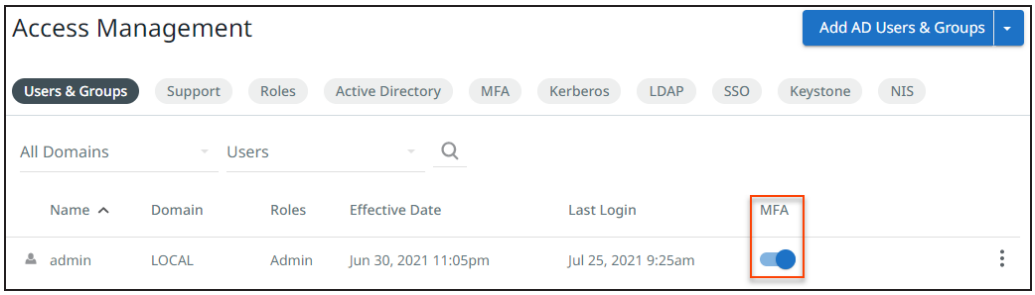

4. When you log in for the first time after MFA is enabled, save the QR code and/or verification code and securely share the code with all users in the distribution list to set up MFA on their own devices.

> **Note:** The Cohesity cluster supports the authenticator app and email-based authentication. You can use one or both authentication methods.

> **Important:** Cohesity 6.8 and above supports MFA for Active Directory users. For more information, see this knowledge base article.

Summary

- If you have followed the guidelines outlined in the Manage Users Locally on the Cohesity Cluster section, you should have multiple local users with elevated permissions to perform additional functions.
- Create custom roles based on your operational requirements. You can then create additional users and assign custom role(s) based on your requirements.
- Periodically evaluate your three elevated users to ensure that as people change roles, you are maintaining a minimum of three active users with elevated permissions.

Open ID Connect Support for Cluster APIs

You can integrate your identity provider solution with Cohesity Data Cloud (Self-managed) to access Cohesity REST APIs through single sign-on.

> **Note:** OpenID Connect is supported only for cluster APIs, and is not available for Single Sign-On (SSO) login via the UI.

Prerequisites

1. You should get the JSON Web Token from your identity server. To create JWT, enter the following parameters:
   - **alg** - The type of algorithm used to generate JWT. For example, **RS256**.
   - **typ** - The type should be JSON Web Token and the value for it is **JWT**.
   - **x5t** - This is X509 Thumbprint. It's a Base64 encoded string.
   - **exp** - This is a Timestamp.
   - **aud** - The audience ID of the client application on the identity server.
   - **upn** - The username or email address of the user.

   > **Note:** The validity of JWT is defined on the Open ID Connect server.

2. You should create an Open ID user on Cohesity Data Cloud (Self-managed) to access Cohesity REST APIs through SSO. For more information, see Create an Open ID User.
3. You should register Cohesity as an application on your identity server.
4. You should create a public key and get the URL to it. Refer to the sample key below,

```
{
"n": "<Modulus_of_Public_Key_in_Base64>",
"e": ["<Exponent_of_Public_Key_in_Base64>"]
"x5t": "<Certificate_Thumbprint_of_Public_Key_in_Base64>"
}
```

How it works

You can use Cohesity REST APIs to run cluster operations. To use Cohesity REST APIs, you can use your credentials or configure single sign-on (SSO) to authenticate without credentials. Cohesity uses OAuth 2.0 authorization framework to allow access to its REST APIs.

Configure SSO

When you configure SSO, Cohesity fetches the public key from the identity server based on the time defined for the parameter **pollingFrequencyMins**.

Access Cohesity REST APIs

The API request is sent from a client application to Cohesity. As a prerequisite, the client application should have the access token in JSON Web Token (JWT) format. This token consists of information about the user's claims, the algorithm used for encoding, and the public and private keys. The client application fetches this token from the identity server such as the Open ID Connect server.

Once the client application receives the token, it sends an API request to the Cohesity REST API endpoint with the JWT in its headers. Cohesity decodes the JWT from the API request

and matches the information against the public key. If the public key matches, then Cohesity allows the user mentioned in the JWT to access its REST APIs.

Configure Identity Provider

This API endpoint allows you to configure your identity provider on Cohesity. As a result, Cohesity requests the public key from the identity server regularly and uses it to authenticate users.

API URL

```
https://<IP_Address>/v2/identity-provider
```

> **Important:** Make sure to use the POST method.

Sample Body

```
{
"id": 1,
"domain": "<Domain_Name>",
"isEnabled": true,
"identityProviderType": "OpenIdConnect"
"openIdConnectParams": {
"audienceIds": [
"<Audience_ID1>"
"<Audience_ID2>"
]
"publicKeyURL": "<URL>",
"pollingFrequencyMins": 1
}
```

Parameters

Parameter	Description
id	Enter the value.
domain	Enter the domain name of the identity server. The domain should be the same as available in the JWT.
isEnabled	If this is set to False, then the OAuth access is disabled. If no value is entered by default, it is set to True.
identityProviderType	Enter the value “OpenIdConnect”.

Parameter	Description
audienceIds	The audience ID of the client application on the identity server.
publicKeyUrl	The URL to the public key on the identity server.
pollingFrequencyMins	The rate at which Cohesity sends the request for the public key to the identity server. The minimum value is one minute and the maximum value is one week. By default, it is set to one day.

Response

The status code 200 is returned.

Create an Open ID User

API URL

```
https://<IP_Address>/irisservices/api/v1/public/idps/principals
```

> **Important:** Make sure to use the POST method.

Sample Body

```
[
{
"principalName": "<Email_Address>",
"objectClass": "kUser",
"Roles":
[
"COHESITY_ADMIN"
],
"domain": "<Domain_Name>"
"restricted": false
}
]
```

Parameters

Parameter	Description
principalName	Enter the username or email address of the user. This should match the value of 'upn' in JWT.
objectClass	You should enter the value "kUser" to create a user.

Parameter	Description
roles	Enter the roles you want to define for the user.
domain	Enter the domain name of the identity server. The domain should be the same as available in the JWT.
restricted	Enter the value “false”.

Response

The status code 200 is returned.

Refresh Public Key

This API endpoint allows you to fetch the public key from the identity server.

API URL

```
https://<IP_Address>/v2/identity-providers/actions
```

> **Important:** Make sure to use the GET method.

Sample Body

```
[
{
“identityProviderType”: “OpenIdConnect”,
“openIdConnectParams”:
{
“action”: ”RefreshPublicKeys”
}
}
]
```

Parameters

Parameter	Description
identityProviderType	Enter the value “OpenIdConnect”.
action	Enter the value “RefreshPublicKeys”.

Response

The status code 200 is returned.

Web Server Certificates

The **Web Server Certificates** page displays the SSL Certificate used by the Cluster UI, along with common name, certificate issue date, and validity period. From this tab, you can perform the following actions:

- View Details
- Show Certificate
- Update Certificate

You can hover over ⓘ next to a certificate to view the function of that specific certificate.

View Details

From the **Web Server Certificates** page, click a certificate or click Actions menu (⋮) > **View Certificate** to view the complete web server certificate details.

The page displays certificate attributes and the certificate trust chain. In the trust chain section, the details of the intermediate CA certificates and the root CA certificate are shown.

The following details for the default certificate and trust chain certificates are displayed:

- Issued To
- Issued By
- Issued On
- Valid Till
- SHA Fingerprint

> **Note:** If the issuing CA certificates are not uploaded, the application displays a message indicating that the certificate chain is incomplete.

To view the certificate or any certificate in the trust chain, click **Show Certificate** available in that particular certificate section.

Click **Back to Web Server Certificates** to return to the **Web Server Certificates** page.

Show Certificate

From the **Web Server Certificates** page, click Actions menu (⋮) > **Show Certificate** to view the cluster certificate.

Update Certificate

The **Update Certificate** option allows you to update the cluster certificates. Perform the following steps to update the certificate:

1. From the **Web Server Certificates** tab, select Actions menu (⋮) > **Update Certificate**. The **Update Certificate** pop-up appears.
2. Click **Select File** under **Cluster Web Server Certificate**, browse to the file location, and select the cluster certificate file.
3. Click **Select File** under **Private Key**, browse to the file location, and select the private key file.
4. Choose one of the following options to update the certificate:
    - **Update & Restart Later** – Updates the certificate once you restart the cluster later.
    - **Update & Restart Now** – Updates the certificate immediately. A **Restart Services** pop-up appears to confirm the restart of the cluster services. Enter **YES** in the pop-up and click **Proceed** to restart the cluster immediately.

Support Toolbox Overview

The Support Toolbox is a utility provided by Cohesity that assists you with various support and maintenance tasks on Cohesity clusters.

The key features of the Support Toolbox are as follows:

- **Health Checks**: The toolbox can run various health checks to diagnose issues within the cluster. For instance, it can check network node reachability and other critical parameters before an upgrade.
- **Diagnostics**: It provides tools for diagnosing and troubleshooting issues within the cluster, helping you to identify and resolve problems quickly.
- **Reporting**: The toolbox can generate reports on the health and status of the cluster, which can be useful for both support and maintenance purposes.

You can use Support Toolbox using one of the following options:

- CLI
- Siren Web UI

Access Support Toolbox from CLI

To use Support Toolbox through CLI, you must download the Support Toolbox from the Cohesity downloads portal and then install it on one of the nodes of the cluster.

This version of the Support Toolbox is supported on the Cohesity 6.8.1 and above versions.

Prerequisites

Before you can install the Support Toolbox to your Cohesity cluster, you must enable Host shell access.

To enable host shell access:

1. Log on to the node's secured shell using the support user:

   ```
   $ ssh support@<node_IP>
   ```

   Where node_IP is the IP address of the node in your Cohesity cluster.

2. Log on to the iris_cli and generate a support token using the following command:

   ```
   $ iris_cli support host-access-token
   ```

   By default, the support token is only valid for an hour. You can extend the support token validity for up to 12 hours using the following command:

   ```
   $ iris_cli support host-access-token validity-in-hours=12
   ```

3. Obtain the Cohesity cluster information using the following command:

   ```
   $ iris_cli cluster info
   ```

4. Share the following details with Cohesity Support to generate an access code:
   - Support Token
   - Cluster ID
   - Cohesity Software Version
   - The number of hours you require the host shell access

5. After you receive the access code with Cohesity Support, enable host shell access using the following command in iris_cli:

   ```
   $ iris_cli support host-access ssh-enable=true
   ```

   ```
   access-code=<access_code>
   ```

   Where access_code is the access code you received from Cohesity Support.

   Ensure you enable host shell access before the support access token you generated in Step 2 expires.

6. Access the host shell of the node on port 2222 using the support user. Use the following command:

   ```
   $ ssh support@<node_IP> -p 2222
   ```

> **Note:** Sudo access is disabled by default. For support channel access, you need to enable sudo access for the support user to the host shell using the following command:

```
$ iris_cli user linux-sudo-access enable=true duration-in-hours=12
```

Install Support Toolbox

Perform the following steps to download and install the Support Toolbox:

1. Go to the Cohesity download page, and select **Cohesity Tools** tab.
2. Select the fast link for support-toolbox-installer and download to your local computer.
3. Use `WinSCP` or `SCP` command to upload the support-toolbox-installer file from your local computer to the Cohesity node.
4. SSH to the Cohesity node where you uploaded the package and navigate to the folder where the package was uploaded to.
5. Unpack the tar file:

   ```
   tar -xzvf support-toolbox-installer.tgz
   ```

   The following files will be present in the package.

   - **st-utilities-20241024-aff631b3.x86_64.rpm**
     - Toolbox utilities
   - **st-signatures-20241024-aff631b3.x86_64.rpm**
     - Toolbox health checks
   - **st-brigidServer-20241024-aff631b3.x86_64.rpm**
     - Backend API Server
   - **st-dart-20241024-aff631b3.x86_64.rpm**
     - Detect and resolve tool
   - **st-upgradeSignatures-20241024-aff631b3.x86_64.rpm**
     - Pre and post upgrade checks
   - **st-framework-3.1.0-20240731-cfdadf52.x86_64.rpm**
     - Framework package
   - **st-sirenui-2.0.0-cfdadf52.x86_64.rpm**
     - Siren UI upgrade package
6. Install the package by running the following command:

   ```
   ./install
   ```

Upon successful installation, an output similar to the following will be displayed:

```
[support@cohesity-00056bdf4f1-node-1 ~]$ ./install

Installing Support Toolbox Full Install : 3.1.0
Local Node IP : 10.14.59.91
SSH Port : 2222

Installing :
st-brigidServer-20250731-cfdadf52.x86_64
st-dart-20250731-cfdadf52.x86_64
st-framework-3.1.0-20250731-cfdadf52.x86_64
st-signatures-20250731-cfdadf52.x86_64
st-sirenui-2.0.0-cfdadf52.x86_64
st-upgradeSignatures-20250731-cfdadf52.x86_64
st-utilities-20250731-cfdadf52.x86_64

------------------------------------------------- OK
Copying package to all nodes in the cluster
------------------------------------------------- OK
Deploying package on all nodes in the cluster
------------------------------------------------- OK
Cleaning up install artifacts

Install Complete Logs saved to /cohesity_logs/support-toolbox/support_toolbox_installer.INFO

Environment Variables may not be active To fix run : source ~/.bashrc
```

Upon installation failure, an output similar to the following will be displayed. You can review the l**og support_toolbox_installer.INFO** and resolve the issue on the affected node:

```
============================================================

One or more nodes have reported issues installing

Please run install again after resolving issues on the nodes

Skipping Cleanup!

Log :  /cohesity_users_home/support/support_toolbox_installer.INFO

============================================================

Install Complete, with Errors
```

Support Toolbox CLI Execution Options

The Support Toolbox uses its command-line interface, `hc_cli` to perform health checks and diagnostics on Cohesity clusters. You can run the health checks by logging into one of the Cohesity cluster nodes.

The Support Toolbox supports the following CLI execution options with `hc_cli`:

- hc_cli run
- hc_cli show
- hc_cli debug
- hc_cli version

> **Important:** Cohesity continuously develops and refines this tool. While these tests are set for release on X date, they may be updated and improved over time with subsequent upgrade packages.

hc_cli run

The `hc_cli run` command executes all available health checks for all the modules on the Cohesity cluster, across all nodes. The modules on the Cohesity cluster are as follows:

- Services
- Thirdparty
- Backup
- Vendor
- Software
- Networking
- Hardware
- Archival
- Performance
- Fingerprint
- Upgrade
- Patch

Services

This module verifies the health of all Cohesity services and checks that best practices were adhered to during the cluster configuration process.

If required, you can run the following command to perform only the health check of the services on the Cohesity cluster:

```
hc_cli run -m services
```

Available Service Health Checks

Following are the health checks that will be performed for the Service module:

- Alerts uptime check
- Apollo uptime check
- Athena uptime check
- Bifrost Broker uptime check

- Bridge Proxy uptime check
- Bridge uptime check
- Gandalf uptime check
- Groot uptime check
- Keychain uptime check
- Librarian uptime check
- Magneto uptime check
- Newscribe uptime check
- SMB2 Proxy uptime check
- Stats uptime check
- Storage Proxy uptime check
- Vault Proxy uptime check
- Vault Proxy uptime check
- Yoda uptime check
- Cluster Check if Kerberos is Enabled
- Helios Connectivity Check
- My Cohesity Connectivity Check
- Alert Mail Configuration Check
- Apollo MR Pipelies Gflag Check
- IPMI Permissions Check

Alerts uptime check

Description

This health check verifies whether the alert process has been running for the past 5 minutes.

Data Source

```
hc_cli service_uptime.py alerts
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the alert process has been running for the past 5 minutes.
Fail	Indicates that the alert process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the alert process uptime; the service may be unstable.

Apollo uptime check

Description

This health check verifies whether the Healer and Space Reclamation Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py apollo
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Healer and Space Reclamation Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Healer and Space Reclamation Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Healer and Space Reclamation Services process uptime; the service may be unstable.

Athena uptime check

Description

This health check verifies whether the Marketplace Apps Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py athena
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Marketplace Apps Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Marketplace Apps Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Marketplace Apps Services process uptime; the service may be unstable.

Bifrost Broker uptime check

Description

This health check verifies whether the Hybrid Extender Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py bifrost_broker
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Hybrid Extender Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Hybrid Extender Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Hybrid Extender Services process uptime; the service may be unstable.

Bridge Proxy uptime check

Description

This health check verifies whether the I/O Operations Helper process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py bridge_proxy
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the **I/O Operations Helper** process has been running and stable for the past 5 minutes.
Fail	Indicates that the **I/O Operations Helper** process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the **I/O Operations Helper** process uptime; the service may be unstable.

Bridge uptime check

Description

This health check verifies whether the I/O Operations Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py bridge
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the I/O Operations Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the I/O Operations Servicess processhas not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the I/O Operations Services process uptime; the service may be unstable.

Gandalf uptime check

Description

This health check verifies whether the Cluster Configuration Manager Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py gandalf
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Cluster Configuration Manager Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Cluster Configuration Manager Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Cluster Configuration Manager Services process uptime; the service may be unstable.

Groot uptime check

Description

This health check verifies whether the Reporting Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py groot
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Reporting Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Reporting Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Reporting Services process uptime; the service may be unstable.

Keychain uptime check

Description

This health check verifies whether the Key Management Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py keychain
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Key Management Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Key Management Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Key Management Services process uptime; the service may be unstable.

Librarian uptime check

Description

This health check verifies whether the Index Catalog Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py librarian
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Index Catalog Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Index Catalog Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Index Catalog Services Services process uptime; the service may be unstable.

Magneto uptime check

Description

This health check verifies whether the Backup Manager Service process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py magneto
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Backup Manager Service process has been running and stable for the past 5 minutes.
Fail	Indicates that the Backup Manager Service process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Backup Manager Service process uptime; the service may be unstable.

Newscribe uptime check

Description

This health check verifies whether the Distributed Key Value Store Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py newscribe
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Distributed Key Value Store Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Distributed Key Value Store Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Distributed Key Value Store Services process uptime; the service may be unstable.

SMB2 Proxy uptime check

Description

This health check verifies whether the SMB Backup Helper v2 Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py smb2_proxy
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the SMB Backup Helper v2 Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the SMB Backup Helper v2 Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the SMB Backup Helper v2 Services process uptime; the service may be unstable.

Stats uptime check

Description

This health check verifies whether the Cluster Metrics Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py stats
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the NAS Backup Helper Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the NAS Backup Helper Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the NAS Backup Helper Services process uptime; the service may be unstable.

Storage Proxy uptime check

Description

This health check verifies whether the NAS Backup Helper Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py storage_proxy
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the NAS Backup Helper Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the NAS Backup Helper Services process has not been running consistently or stable for the past 5 minutes.

Status	Description
Error	Indicates an error while retrieving the NAS Backup Helper Services process uptime; the service may be unstable.

Vault Proxy uptime check

Description

This health check verifies whether the Archive Helper Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py vault_proxy
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Archive Helper Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Archive Helper Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Archive Helper Services process process uptime; the service may be unstable.

Yoda uptime check

Description

This health check verifies whether the Indexing Services process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py yoda
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Indexing Services process has been running and stable for the past 5 minutes.
Fail	Indicates that the Indexing Services process has not been running consistently or stable for the past 5 minutes.
Error	Indicates an error while retrieving the Indexing Services process uptime; the service may be unstable.

Cluster Check if Kerberos is Enabled

Description

This health check verifies whether Kerberos authentication is enabled on the cluster.

Data Source

```
nohup cluster_config.sh fetch /tmp/cc_tmp_ng > /dev/null 2>&1; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates Kerberos authentication is enabled on the cluster.
Fail	Indicates Kerberos authentication is not enabled on the cluster.

Helios Connectivity Check

Description

This health check verifies whether Helios is accessible from the cluster.

Data Source

```
curl -Is https://helios.cohesity.com >/dev/null 2>&1 ; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that Helios is accessible from the cluster.
Fail	Indicates that Helios is not accessible from the cluster.

My Cohesity Connectivity Check

Description

This health check verifies whether the mycohesity website is accessible from the cluster.

Data Source

```
curl -Is https://my.cohesity.com/ >/dev/null 2>&1 ; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that mycohesity is accessible from the cluster.
Fail	Indicates that mycohesity is not accessible from the cluster.

Alert Mail Configuration Check

Description

This health check verifies whether any email is configured on the cluster for alert notifications.

Data Source

```
gt.sh fetch alerts_data cohesity.alerts.AlertsDataProto 2>/dev/null | grep -
q email_address ;echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that email is configured on the cluster for alert notifications.

Status	Description
Fail	Indicates that email is not configured on the cluster for alert notifications.

Apollo MR Pipelies Gflag Check

Description

This health check verifies whether the Apollo MR Pipelines Gflag is set correctly

Data Source

```
allssh.sh links http://localhost:24680/flagz | grep apollo_mr_pipelines_to_
run=[0-9].[0-9].[0-9][a-z] | wc -l
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that Gflag is configured correctly (not present).
Fail	Indicates that the Gflag is enabled, disable once agressive garbage collection has been completed.

IPMI Permissions Check

Description

This health check verifies whether the IPMI tool has its permissions set properly.

Data Source

```
ls -l /usr/bin/ipmitool | sed 's/=//g' | awk '{print $3, $4}'
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the permissions for the IPMI tool are set to `root:root`.
Fail	Indicates that the permissions for the IPMI tool are defined incorrectly.

Third-party

This module verifies the health of all the third-party softwares used by the Cohesity cluster, such as active directory.

If required, you can run the following command to perform only the health check of the Thrid-Party software on the Cohesity cluster:

```
hc_cli run -m thirdparty
```

Available Third-party Health Checks

Following are the health checks that will be performed for the Thrid-party module:

- LDAP Server Connectivity Check

LDAP Server Connectivity Check

Description

This health check verifies whether all LDAP servers are accessible from the nodes of the cluster.

Data Source

```
domain_connectivity.py
```

State

The following table lists the output state of this check:

Status	Description
Pass	Indicates that all LDAP servers are available via port 389.
Fail	Indicates that one or more LDAP servers discovered are not accessible via port 389.

Backup

This module verifies the health and status of the Backup Manager Service process and any other related functions regarding data backup and restore.

If required, you can run the following command to perform only the health check of the Backup Manager Service process on the Cohesity cluster:

```
hc_cli run -m backup
```

Available Backup Health Check

Following are the health checks that will be performed for the Magneto module:

- Backup

Backup Manager ServiceMagneto uptime check

Description

This health check verifies whether the Backup Manager Service process has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py magneto
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Backup Manager Service process has been running and stable for the past 5 minutes.
Fail	Indicates that the Backup Manager Service process has not been running consistently or stable for the past 5 minutes.

Vendor

This module verifies the health and status of the OEM Hardware Platforms from Cisco, Dell, HPE, and other qualified vendors.

If required, you can run the following command to perform only the health check of the OEM Hardware Platforms:

```
hc_cli run -m vendor
```

Available Vendor Health Checks

Following are the health checks that will be performed for the vendor module:

- Cisco C240 M5 BIOS 4.0.4p.0.1112191020 Version Check

Cisco C240 M5 BIOS 4.0.4p.0.1112191020 Version Check

Description

This health check verifies if any UCS C240 nodes running BIOS 4.0.4p.0.1112191020.

Data Source

```
sudo dmidecode -t bios | grep Version | cut -d ':' -f2 | grep
'C240M5.4.0.4p.0.1112191020' ; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that none of the UCS C240 nodes are running BIOS 4.0.4p.0.1112191020.
Fail	Indicates that one or or more UCS C240 nodes are running BIOS 4.0.4p.0.1112191020.

Software

This module verifies the software-specific health checks of the Cohesity cluster.

If required, you can run the following command to perform only the software-specific health checks of the Cohesity cluster:

```
hc_cli run -m software
```

Available Software Health Checks

Following are the health checks that will be performed for the Software module:

- One-Minute Average Load Check
- Five-Minute Average Load Check
- Fifteen-Minute Average Load Check
- Root Partition Size Check
- Log Partition Size Check
- Data Partition Size Check
- tmp Partition Size Check
- Outstanding Disk command Backlog Check
- D State Process Check
- Disks stuck in marked for removal check
- Node stuck in marked for removal check

One-Minute Average Load Check

Description

This health check verifies whether the node has been overloaded for the past 1 minute.

Data Source

```
average_load_one.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Average CPU Load is less than 3 times the number of cores.
Fail	Indicates that the Average CPU load is greater than or equal to 3 times the number of cores.
Error	Indicates an error running the command.

Five-Minute Average Load Check

Description

This health check verifies whether the node has been overloaded over the past 5 minutes.

Data Source

```
average_load_five.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the average CPU load is less than 5 times the number of cores.
Fail	Indicates that the average CPU load is greater than or equal to 5 times the number of cores.
Error	Indicates an error running the command.

Fifteen-Minute Average Load Check

Description

This health check verifies if the node has been overloaded over the past 15 minutes.

Data Source

```
average_load_five.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the average CPU load is less than 5 times the number of cores.
Fail	Indicates that the average CPU load is greater than or equal to 5 times the number of cores.
Error	Indicates an error running the command.

Root Partition Size Check

Description:

This health check verifies the usage of the root partition. The root partition is used for the Cohesity node system and OS-related files.

Data Source

```
df -h /| awk '{ print $5 }' | tail -1 | tr -d '%'
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the root partition usage is less than 75%.
Fail	Indicates that the root partition usage is equal to or greater than 85%.
Warning	Indicates that the root partition usage is between 75 and 85%.

Log Partition Size Check

Description:

This health check verifies the usage of the log partition. A log partition is used for storing logs that belong to each of the Cohesity cluster components. It helps monitor and troubleshoot by keeping detailed logs of system activities and errors.

Data Source

```
df -h /cohesity_logs | awk '{ print $5 }' | tail -1 | tr -d '%'
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the log partition usage is less than 75%.
Fail	Indicates that the log partition usage is equal to or greater than 85%.
Warning	Indicates that the log partition usage is between 75 and 85%.

Data Partition Size Check

Description:

This health check verifies usage of the `home_cohesity_data` partition.

The data partition is used by various services to store configuration files and maintain consistency across all nodes in the cluster. This partition has a symlink to `/home/cohesity/data` and holds all the configuration files and data related to each of the Cohesity services.

Data Source

```
df -h /home_cohesity_data | awk '{ print $5 }' | tail -1 | tr -d '%'
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the `home_cohesity_data` partition usage is less than 75%.
Fail	Indicates that the `home_cohesity_data` partition usage is equal to or greater than 85%.

Status	Description
Warning	Indicates that the `home_cohesity_data` partition usage is between 75 and 85%.

tmp Partition Size Check

Description:

This health check verifies the usage of the tmp partition. The tmp partition is used for temporary file storage.

Data Source

```
ps aux | grep 'sg_' |wc -l
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the tmp partition usage is less than 75%.
Fail	Indicates that the tmp partition usage is equal to or greater than 85%.
Warning	Indicates that the tmp partition usage is between 75 and 85%.

Outstanding Disk command Backlog Check

Description:

This health check verifies if there are any excessive disk commands queued up in the backlog.

Data Source

```
ps aux | grep 'sg_' |wc -l
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there is no backlog of disk operations detected.

Status	Description
Fail	Indicates that there are more than 20 outstanding disk commands in the backlog.
Warning	Indicates that there are 10 to 20 disk commands in the backlog.

D State Process Check

Description:

This health check verifies if there are any processes in the D (defunct) state.

Data Source

```
ps aux | awk ' { print $8 } ' | grep D | wc -l
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are no or less than 15 processes in the D (defunct) state.
Fail	Indicates that there are more than 25 processes in the D (defunct) state.
Warning	Indicates that there are between 15 to 25 processes in the D (defunct) state.

Disks stuck in marked for removal check

Description:

This health check verifies if there are any disks marked for removal.

This health check helps Cohesity Support to identify disks awaiting removal and locate any services that have not acknowledged the disk chunks, enabling Cohesity to take necessary remedial actions.

Data Source

```
disk_removal_check.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are disks marked for removal.
Fail	Indicates that no disks are marked for removal.
Warning	Indicates that an exception error occurred during this health check.

Node stuck in marked for removal check

Description:

This health check verifies whether there are any nodes in the cluster stuck in marked for removal state.

This checks helps the Cohesity support to identify nodes which are marked for removal in cluster and find the services which have not ACKed the disks chunks and take necessary remedy steps.

Data Source

```
node_removal_check.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are nodes marked for removal.
Fail	Indicates that no nodes are marked for removal.
Warning	Indicates that an exception error occurred during this health check.

Networking

This module verifies the health state and configuration of the networking and interfaces of the individual nodes within a Cohesity cluster.

If required, you can run the following command to perform only the networking and interfaces health checks of the Cohesity cluster:

```
hc_cli run -m network
```

Available Networking Health Checks

Following are the health checks that will be performed for the Networking module:

- Firewall Check running state
- Bond mode configuration check
- Network interface MTU check
- Default gateway check
- Firewalld Status Check
- Primary interface health check
- VIP on Bond check
- Helios external server connectivity test
- Microsoft o365 external server connectivity test
- Support Channel external server connectivity test
- GCP (Google cloud) External Server Connectivity Test
- AWS (Amazon Web Services) External Server Connectivity Test

Firewall Check running state

Description

This health check verifies whether the Firewalld Daemon is running.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

```
sudo firewall-cmd --state | grep -qx running; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the firewall is currently running.
Fail	Indicates that the firewall is currently not running.

Bond mode configuration check

Description

This health check verifies whether the bonding mode configured for the Cohesity cluster is 1 or 4.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

The following data source finds the names of all the bonds on the node:

- `grep bond /proc/net/dev | awk '{print $1}' | tr -d ':'`

The following data source evaluates the bond mode

- `cat /etc/sysconfig/network-scripts/ifcfg-<bond_name> | grep BONDING_OPTS| cut -d '\"' -f2 | cut -d ' ' -f1 | cut -d '=' -f2`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the bonding mode configured for the cluster is either 1 (active-backup) or 4 (LACP).
Fail	Indicates that the bonding mode configured for the cluster is 0,2,3,5 or 6.

Network interface MTU check

Description

This health check verifies whether the maximum transmission unit (MTU) is set to 1500.

MTU defines the largest packet size that can be transmitted as a single entity in a network connection. The size of the MTU determines the amount of data that can be transmitted in bytes over a network.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

- To gather all interface names:
  - `list_all_nic_ports.sh | grep UP | awk '{print $(NF-2)}'`
- To pull MTU configuration:
  - `ip link show dev <innterface_name> | cut -d ' ' -f5 | head -1`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the MTU is set to 1500.
Warning	Indicates that the MTU is set to a value greater than 1500.
Fail	Indicates that the MTU is set to a value less than 1500.

Default gateway check

Description

This health check verifies the status of the Address Resolution Protocol (ARP) for the default gateway.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

- Collect default gateway IP:
  - `for g in `route -n | grep -w 0.0.0.0 | grep UG| awk ' { print $2 } '`; do ping -c 1 -W 3 $g > /dev/null 2>&1; echo $g;done`

- Evaluate if the default gateways had a successful ARP
  - `arp -an | grep '<gateway_ip>)' | grep incomplete; echo $?`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the ARP for the default gateway is successful.
Fail	Indicates that the ARP for the default gateway is incomplete.

Firewalld Status Check

Description

This health check verifies if the daemon for firewalld is active.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

```
sudo systemctl is-active firewalld | grep -qx active; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the firewalld is active.
Fail	Indicates that the firewalld is inactive.

Primary interface health check

Description

This health check verifies whether all the interfaces connected to the primary bond of the cluster are healthy.

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

Data Source

- Collect all interface names on the primary bond:
  - `cat /proc/net/bonding/`primary_interface_name.sh`|grep 'Slave Interface:'| awk '{print $3}'`
- Verify if all the interfaces are in an UP state
  - `ip link | grep <interface_name> | grep -q 'state UP';echo $?`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the interfaces connected to the primary bond of the cluster are healthy.
Fail	Indicates that the interfaces connected to the primary bond of the cluster are not healthy.

VIP on Bond check

Description

This health check verifies whether any VIP is present on the bond and displays VIP assigned on each of the cluster nodes.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

- To grep the bond interfaces having VIP present on them:
  - `ip addr show|  /usr/bin/awk '{print $NF}' | /usr/bin/grep '^bond'`
- To get the current VIP assignment:
  - `gt.sh fetch vip_assignments cohesity.nexus.services.vip.VipAssignmentProto 2> /dev/null`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the bond interface isn't holding any VIP.
Fail	Indicates that the bond interface is holding a VIP.

Helios external server connectivity test

Description

This health check verifies whether the external Helios server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.

This checks helps the Cohesity Support to identify if the external Helios server are reachable and accessible via SSL/TLS.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

`external_server_connection_test.py`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the external Helios server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.
Fail	Indicates that the external Helios server connectivity for SSL/TLS certificate is not working from all the nodes in a cluster.

Microsoft o365 external server connectivity test

Description

This health check verifies whether the external Microsoft o365 server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.

This checks helps the Cohesity Support to identify if the external Microsoft o365 server are reachable and accessible via SSL/TLS.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

```
external_server_connection_test.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the external Microsoft o365 server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.
Fail	Indicates that the external Microsoft o365 server connectivity for SSL/TLS certificate is not working from all the nodes in a cluster.

Support Channel external server connectivity test

This health check verifies whether the external Support Channel server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.

This checks helps the Cohesity Support to identify if the external Support Channel server is reachable and accessible via SSL/TLS.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

```
external_server_connection_test.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the external Support Channel server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.
Fail	Indicates that the external Support Channel server connectivity for SSL/TLS certificate is not working from all the nodes in a cluster.

GCP (Google cloud) External Server Connectivity Test

This health check verifies whether the external GCP server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.

This checks helps the Cohesity Support to identify if the external GCP server is reachable and accessible via SSL/TLS.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

```
external_server_connection_test.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the external GCP server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.
Fail	Indicates that the external GCP server connectivity for SSL/TLS certificate is not working from all the nodes in a cluster.

AWS (Amazon Web Services) External Server Connectivity Test

This health check verifies whether the external AWS server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.

This checks helps the Cohesity Support to identify if the external AWS server is reachable and accessible via SSL/TLS.

Supported Cluster Types

You can run this health check on the following cluster types:

- Virtual Edition
- Cloud Edition
- Physical Edition

Data Source

```
external_server_connection_test.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the external AWS server connectivity for SSL/TLS certificate is working from all the nodes in a cluster.
Fail	Indicates that the external AWS server connectivity for SSL/TLS certificate is not working from all the nodes in a cluster.

Hardware

This module verifies the health state of the hardware of the individual nodes in a Cohesity cluster.

If required, you can run the following command to perform only the hardware-related health checks of the Cohesity cluster:

```
hc_cli run -m hardware
```

Available Archival Health Checks

Following are the health checks that will be performed for the archival module:

- SMARTCTL Status Check
- IO Stat SD Disk Usage Status Check
- HBA Health Check
- NVME Error Check
- HDD Read Error Check
- HDD Write Error Check
- HDD Verify Error Check
- CPU Throttle State Check
- PCIE Bus Error Check
- Node Uptime Check
- HDD Disk availability Check
- SSD Read only check
- Missing NVME Disk Check
- Missing System Disk Check
- Missing HDD Disk Check
- Missing DIMM Disk Check
- Correctable DIMM error check
- Cohesity Firmware Check
- Product Model Name Check
- Disk Serial Number Check
- NVME slot check

SMARTCTL Status Check

Description

This health check verifies the smartctl status of all disks in the cluster.

Data Source

This health check uses the following commands to check the status:

- To gather disk names :
  - `ls /dev | awk '/^sd[a-z]$/ { print $1 }'`
- To evaluate status ($data is replaced with disk name):
  - `sudo smartctl -H --json /dev/<disk_name> | grep passed | cut -d ':' -f2 | tr -d ' '`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the smartctl is reporting all the disks in the cluster as healthy.
Fail	Indicates that the smartctl is reporting disk failure in the cluster.

IO Stat SD Disk Usage Status Check

Description

This health check verifies whether the disks in the cluster are utilized at 80% or greater of their IO capacity.

Data Source

This health check uses the following commands to check the status:

- To gather disk names:
  - `ls /dev | awk '/^sd[a-z]$/ { print $1 }'`
- To pull IOStat data ($data is replaced with the disk name from above:
  - `iostat -dx <disk_name> | awk '{ print $14 } ' | tr -d ' '| tail -2| head -1`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the disk usage is below 80% of their IO capacity.
Fail	Indicates that the disk usage is equal to or more than 80% of their IO capacity.

HBA Health Check

Description

This health check verifies whether there are any DWORD errors on the disk HBA.

Data Source

This health check uses the following commands to check the status:

- To gather disk names:
  - `ls /dev | awk '/^sd[a-z]$/ { print $1 }'`
- To pull sg_logs data:
  - `sudo sg_logs -a /dev/<disk_name> |grep -qi 'invalid dword' ; echo $?`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are no DWORD errors on the disk HBA.
Fail	Indicates that there are one or more DWORD errors on the disk HBA.

NVME Error Check

Description

This health check verifies whether there are any NVME drive data and media integrity errors.

Data Source

This health check uses the following commands to check the status:

- Precheck to evaluate if the node has NVME drives:
  - `ls /dev | grep -q nvme ; echo $?`
- To gather disk names:
  - `ls /dev | awk '/^nvme[0-99]$/ {print $1}'`
- To pull sg_logs data:
  - `sudo smartctl -a /dev/<disk_name> | grep Media | cut -d ':' -f 2| tr -d ' ,'`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are no errors on the drive.
Fail	Indicates that there are errors on the drive.

HDD Read Error Check

Description

This health check verifies whether the spinning disks in the cluster are encountering any read errors.

Data Source

This health check uses the following commands to check the status:

- To gather disk names:
  - `ls /dev | awk '/^sd[a-z]$/ { print $1 }'`
- To pull sg_logs data:

- `sudo smartctl -l error /dev/<disk_name> | grep -q 'read:' ; echo $?`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the spinning disks in the cluster are not encountering any read errors.
Fail	Indicates that the spinning disks in the cluster are encountering one or more read errors.

HDD Write Error Check

Description

This health check verifies whether the spinning disks in the cluster are encountering any write errors.

Data Source

This health check uses the following commands to check the status:

- To gather disk names:
  - `ls /dev | awk '/^sd[a-z]$/ { print $1 }'`
- To pull sg_logs data:
  - `sudo smartctl -l error /dev/<disk_name> | grep -q 'read:' ; echo $?`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the spinning disks in the cluster are not encountering any write errors.
Fail	Indicates that the spinning disks in the cluster are encountering one or more write errors.

HDD Verify Error Check

Description

This health check verifies whether the spinning disks in the cluster are encountering any verify errors.

Data Source

This health check uses the following commands to check the status:

- To gather disk names:
  - `ls /dev | awk '/^sd[a-z]$/ { print $1 }'`
- To pull sg_logs data:
  - `sudo smartctl -l error /dev/<disk_name> | grep -q 'read:' ; echo $?`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the spinning disks in the cluster are not encountering any verify errors.
Fail	Indicates that the spinning disks in the cluster are encountering one or more verify errors.

CPU Throttle State Check

Description

This health check verifies whether the CPU in the node is running at the right frequency and is not throttled down.

Data Source

This health check uses the following command to check the status:

```
lscpu | grep 'CPU MHz' | cut -d ':' -f 2| tr -d ' '
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the CPU in the node is running at a frequency greater than 1GHZ.
Fail	Indicates that the CPU in the node is running at a frequency less than 1GHZ.

PCIE Bus Error Check

Description

This health check verifies whether the PCI Bus is reporting any errors.

Data Source

This health check uses the following command to check the status:

```
dmesg |grep -q 'PCIe Bus Error';echo $?
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the PCI Bus is not reporting any errors.
Fail	Indicates that the PCI Bus is reporting one or more errors.

Node Uptime Check

Description

This health check verifies how long the node has been up.

Data Source

This health check uses the following command to check the status:

```
cat /proc/uptime | cut -d ' ' -f1
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition
- Virtual Editon

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the node has been up for a day or more.
Warning	Indicates that the node is up for more than 5 minutes but not longer than a day.
Fail	Indicates that the node is up for less than 5 minutes.

HDD Disk availability Check

Description

This health check verifies whether at least two or more disks are mounted on the node.

Data Source

This health check uses the following command to check the status:

```
smartctl --scan | wc -l
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Cloud Edition
- Physical Edition
- Virtual Editon

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that more than two disks are mounted on the node.
Fail	Indicates that less than two disks are mounted on the node.

SSD Read only check

Description

This health check verifies whether there are any SSD(s) that are currently set to read-only.

Data Source

This health check uses the following command to check the status:

```
grep ' ro,' /proc/mounts |awk '{print $1}' | grep nvme
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are no SSD(s) that are currently set to read-only.
Fail	Indicates that there are SSD(s) that are currently set to read-only.

Missing NVME Disk Check

Description

This health check verifies whether all expected NVME disks are present and mounted.

Data Source

This health check uses the following command to check the status:

```
platform.py missing_disks:NVME_Disk
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all expected NVME disks are present and mounted.
Fail	Indicates that not all expected NVME disks are present or mounted.

Missing System Disk Check

Description

This health check verifies whether all expected system disks are present and mounted.

Data Source

This health check uses the following command to check the status:

```
platform.py missing_disks:SystemDisk
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all expected system disks are present and mounted.
Fail	Indicates that not all expected system disks are present or mounted.

Missing SATA SSD Disk Check

Description

This health check verifies whether all expected SATA SSD disks are present and mounted.

Data Source

This health check uses the following command to check the status:

```
platform.py missing_disks:SataSsdDisk
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all expected SATA SSD disks are present and mounted.
Fail	Indicates that not all expected SATA SSD disks are present or mounted.

Missing HDD Disk Check

Description

This health check verifies whether all expected HDD disks are present and mounted.

Data Source

This health check uses the following command to check the status:

```
platform.py missing_disks:HddDisk
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all expected HDD disks are present and mounted.
Fail	Indicates that not all expected HDD disks are present or mounted.

Missing DIMM Disk Check

Description

This health check verifies whether the node has all the expected DIMM slots populated with working memory.

Data Source

This health check uses the following command to check the status:

```
platform.py missing_dimm:none
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the node has all the expected DIMM slots populated with working memory.
Fail	Indicates that the has one or more DIMMs missing.

Correctable DIMM error check

Description

This health check verifies whether there are any correctable DIMM errors present on the node.

Data Source

This health check uses the following command to check the status:

```
platform.py correctable_dimm_error:none
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are no correctable DIMM errors present on the node.
Fail	Indicates that are are one or more correctable DIMM errors present on the node.

Cohesity Firmware Check

Description

This health check verifies whether the firmwares of the cluster nodes are up to date with the latest Cohesity release version .

Data Source

This health check uses the following command to check the status:

```
platform.py check_firmware:none
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there all firmware is up to date with the latest Cohesity releases version.
Fail	Indicates that are are one or more devices in the node that needs a firmware update.

Product Model Name Check

Description

This health check verifies whether the product model of the nodes of the cluster is unknown.

Data Source

This health check uses the following command to check the status:

```
grep -i productModel\\\" /home_cohesity_data/nexus/hardware.json | grep -q
unknown ; echo $?
```

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the product model of the nodes of the cluster is not unknown.
Fail	Indicates that the product model of the nodes of the cluster is unknown.

Disk Serial Number Check

Description

This health check verifies whether all disks mounted on the node match a serial number available in the cluster config.

Data Source

This health check uses the following command to check the status:

- To gather disk serial:
  - `list_all_disks.sh | grep -v '=\\|Disk\\|SLOT\\|^[[:space:]]*$'|awk '{print$2}'`
- To validate serial numbers:
  - `grep -q $<disk_serial> /tmp/cc_tmp_ng; echo $?`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all disks mounted on the node match a serial number available in the cluster config.

Status	Description
Fail	Indicates that one or moew disks mounted on the node does not match a serial number available in the cluster config.

NVME slot check

Description

This health check verifies whether current mapping of the NVMe disk slot to the physical slot matches the expected configuration.

Data Source

This health check uses the following command to check the status:

- To get the expected mapping:
  - `/home/cohesity/software/crux/bin/product_helper --op=GET_DISK_SLOT_TO_PHY_MAP --diskType=NvmeSsdDisk`
- To get the current mapping of the NVME slot to the phy slot:
  - `/usr/sbin/lspci -D  | grep 'Non-Volatile' | awk '{print $1}'`

Supported Cluster Types

You can run this health check on the following cluster types:

- Physical Edition

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the current mapping of the NVMe disk slot to the physical slot matches the expected configuration – No slot mismatch detected.
Fail	Indicates that the current mapping of the NVMe disk slot to the physical slot does not matches the expected configuration – Slot mismatch detected.

Archival

This module verifies the health state of the I/O Operations Services and Archive Services.

If required, you can run the following command to perform only the archival health checks of the Cohesity cluster:

```
hc_cli run -m archival
```

Available Archival Health Checks

Following are the health checks that will be performed for the archival module:

- Bridge uptime check
- Vault Connectivity Check

Bridge uptime check

Description

This health check verifies whether the I/O Operations Services is up on all cluster nodes and has been running and stable for the past 5 minutes.

Data Source

```
service_uptime.py bridge
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the I/O Operations Services has been running and stable for the past 5 minutes.
Fail	Indicates that the I/O Operations Services is not stable for the past 5 minutes.

Vault Connectivity Check

Description

This health check verifies whether all configured vaults are reachable via port 3000.

Data Source

```
vault_connectivity.py
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all configured vaults are reachable via port 3000.
Fail	Indicates that all configured vaults are not reachable via port 3000.

Performance

This module checks performance indicators on the cluster, to verify if it is operating within expected parameters.

If required, you can run the following command to perform only the performance-related health checks of the Cohesity cluster:

```
hc_cli run -m perform
```

Available Performance Health Checks

Following are the health checks that will be performed for the Performance module.

External IO Percentage Usage Check

Description

This health check verifies the I/O consumption on the I/O Services.

Data Source

```
elinks -dump http://0:11111 | grep -i 'external io' | awk '{ print $6 }' |
head -1
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the I/O consumption on the I/O services is less than 75%.
Fail	Indicates that the I/O consumption on the I/O services is 100%.
Warning	Indicates that the I/O consumption on the I/O services is less than 75%.

Upgrade

This module verifies the upgrade-specific health checks of the Cohesity cluster.

If required, you can run the following command to perform only the upgrade-specific health checks of the Cohesity cluster:

```
hc_cli run -m preUpgrade
```

Available Upgrade Health Checks

Following are the health checks that will be performed for the archival module:

- Boot Partition Size Check
- CPU Throttle State Check
- Gandalf Stability Check
- Bond Configuration Check
- Downloads Site Connectivity Check
- /home_cohesity_data Read Only Check
- Nexus Stability Check
- SMARTCTL Status Check
- NVME Error Check
- Missing NVME Disk
- Missing System Disk
- Missing SATA SSD Disk
- Missing HDD Disk
- /home_cohesity_data Space Check
- Virtualization Enabled Check
- Disk(s) marked for removal Check
- Stale Mount Check
- Athena app subnet overlap check

Boot Partition Size Check

Description

This health check verifies the usage of the boot partition.

Data Source

```
df -h /boot | awk '{ print $5}' | tail -1 | cut -d '%' -f1
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the boot partition usage is less than 80%.
Fail	Indicates that the boot partition usage is equal to or greater than 80%.

CPU Throttle State Check

Description

This health check verifies whether the CPU's of the cluster are running at their default frequency and not throttled slower

Data Source

```
lscpu | grep 'CPU MHz' | cut -d ':' -f 2| tr -d ' '
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the CPU is running at frequency greater than 1GHz.
Fail	Indicates that the CPU is throttled below 1GHz.

Gandalf Stability Check

Description

This health check verifies whether the gandalf service has been stable for the past 5 minutes.

Data Source

```
service_uptime.py gandalf
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the gandalf service has been stable for the past 5 minutes.
Fail	Indicates that the gandalf service has not been stable or down for the past 5 minutes.

Node Connectivity Check

Description

This health check verifies whether a node can ping all other nodes in a cluster.

Data Source

```
nping -c 1 -p 22 $data | grep attempts | cut -d '|' -f3 | cut -d ':' -f2 |
cut -d '(' -f1 | tr -d ' '
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the node is able to be reach over the network.
Fail	Indicates that the node is unable to be reach over the network.

Bond Configuration Check

Description

This health check verifies whether the bonding mode configured for the cluster is either 1 (active-backup) or 4 (LACP).

Data Source

- The following data source finds the names of all the bonds on the node:
  - `grep bond /proc/net/dev | awk '{print $1}' | tr -d ':'`
- The following data source evaluates the bond mode
  - `cat /etc/sysconfig/network-scripts/ifcfg-<bond_name> | grep BONDING_OPTS| cut -d '\"' -f2 | cut -d ' ' -f1 | cut -d '=' -f2`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the bonding mode configured for the cluster is either 1 (active-backup) or 4 (LACP).
Fail	Indicates that the bonding mode configured for the cluster is 0,2,3,5 or 6.

Downloads Site Connectivity Check

Description

This health check verifies whether the downloads.cohesity.com site is reachable from the cluster node.

Data Source

```
curl -Is https://downloads.cohesity.com | head -1 | awk '{print $2}'
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the downloads.cohesity.com site is reachable from the cluster node.
Fail	Indicates that the downloads.cohesity.com site is not reachable from the cluster node.

SSD Lifetime Write Limit Check

Description

This health check verifies whether the SSD of the cluster has reached its End of Life.

Data Source

```
service_uptime.py gandalf
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the SSD of the cluster has not reached its End of Life.
Fail	Indicates that the SSD of the cluster is approaching its End of Life.

/home_cohesity_data Read Only Check

Description

This health check verifies whether the /home_cohesity_data partition is mounted as read-write or read-only.

Data Source

```
mount | grep home_cohesity_data | cut -d '(' -f2 | cut -d ',' -f1
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the /home_cohesity_data partition is mounted as read-write.
Fail	Indicates that the /home_cohesity_data partition is mounted asread-write or read-only.

Nexus Stability Check

Description

This health check verifies whether the Cluster Manager Service has been stable for the past 5 minutes.

Data Source

```
service_uptime.py nexus
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the Cluster Manager Service has been stable for the past 5e has been stable for the past 5 minutes.
Fail	Indicates that the Cluster Manager Service has not been stable or down for the past 5 minutes.
Error	Indicates an error while retrieving the Cluster Manager Service uptime; the service may be unstable.

SMARTCTL Status Check

Description

This health check verifies the smartctl status of all disks in the cluster.

Data Source

```
sudo smartctl -H --json /dev/$data | grep -io \"\\\"passed\\\": [a-zA-Z\\!]*\\|USB\" | cut -d ':' -f2 | tr -d ' '
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the smartctl is reporting all the disks in the cluster as healthy.
Fail	Indicates that the smartctl is reporting disk failure in the cluster.

NVME Error Check

Description

This health check verifies health of NVME disks on the cluster.

Data Source

- To collect NVME Names
  - `ls /dev | awk '/^nvme[0-99]$/ {print $1}'`
- To evaluate NVME Error counts:
  - `sudo smartctl -a /dev/<nvme_name> | grep Media | cut -d ':' -f 2| tr -d ' ,'`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that no errors on the NVME disks.
Fail	Indicates that there are errors on the NVME disks.

Missing NVME Disk

Description

This health check verifies whether all the expceted NVME disks are available and mounted.

Data Source

```
platform.py missing_disks:NVME_Disk
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all the expceted NVME disks are available and mounted.
Fail	Indicates that one ore more NVME disks are missing.

Missing System Disk

Description

This health check verifies whether all the expceted systemE disks are available and mounted.

Data Source

```
platform.py missing_disks:SystemDisk
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all the expceted system disks are available and mounted.
Fail	Indicates that one ore more system disks are missing.

Missing SATA SSD Disk

Description

This health check verifies whether all the expceted SATA SSD disks are available and mounted.

Data Source

```
platform.py missing_disks:SATA_SSD_Disk
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all the expceted SATA SSD disks are available and mounted.
Fail	Indicates that one ore more SATA SSD disks are missing.

Missing HDD Disk

Description

This health check verifies whether all the expceted HDD disks are available and mounted.

Data Source

```
platform.py missing_disks:HDDDisk
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that all the expceted HDD disks are available and mounted.
Fail	Indicates that one ore more HDD disks are missing.

/home_cohesity_data Space Check

Description

This health check verifies whether the /home_cohesity_data partition has 14 GiB or more free space.

Data Source

```
df -h /home_cohesity_data/ | awk '{print $4}'|tail -1|tr -d 'G'
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the node has sufficient space for upgrade.
Fail	Indicates that the node has insufficient space for upgrade.

Virtualization Enabled Check

Description

This health check verifies whether virtualization technology is enabled on the CPU.

Data Source

This health check uses the following command to check the status:

```
sudo virt-host-validate qemu | grep -q FAIL;echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the node has hardware virtualization enabled.
Fail	Indicates that the node does not have hardware virtualization enabled.

Disk(s) marked for removal Check

Description

This health check verifies whether there are any disks marked for removal.

Data Source

This health check uses the following command to check the status:

```
grep -q MarkedForRemoval /tmp/cc_tmp_ng ; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that there are no disks marked for removal.
Fail	Indicates that one ore more disks marked for removal.

Stale Mount Check

Description

This health check verifies whether there are any stale Indexing Service mounts on the cluster.

Data Source

This health check uses the following commands to check the status:

- To collect list of any Indexing Service mounts
  - `mount | grep yoda_clone_backup | awk '{print $3}'`
- To evaluate mount time
  - `find $<index_service_mount>/clone_backup/ -maxdepth 1 -mtime +365`

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that the node has no old or stale backup views present.
Fail	Indicates that there are one or more stale backup views present.

Athena app subnet overlap check

Description

This health check verifies whether the Marketplace Apps Service subnets overlaps with any internal subnet.

Data Source

This health check uses the following command to check the status:

```
grep 'is_athena_subnet_clash' /tmp/cc_tmp_ng | awk '{print $2}'
```

State

The following table lists the status of this check:

Status	Description
Pass	Marketplace Apps Service subnets do not overlap with any internal subnet.
Fail	Marketplace Apps Service subnets overlap with an internal subnet.

Vendor

This module verifies the health and status of the OEM Hardware Platforms from Cisco, Dell, HPE, and other qualified vendors.

If required, you can run the following command to perform only the health check of the OEM Hardware Platforms:

```
hc_cli run -m vendor
```

Available Vendor Health Checks

Following are the health checks that will be performed for the vendor module:

- Cisco C240 M5 BIOS 4.0.4p.0.1112191020 Version Check

Cisco C240 M5 BIOS 4.0.4p.0.1112191020 Version Check

Description

This health check verifies if any UCS C240 nodes running BIOS 4.0.4p.0.1112191020.

Data Source

```
sudo dmidecode -t bios | grep Version | cut -d ':' -f2 | grep
'C240M5.4.0.4p.0.1112191020' ; echo $?
```

State

The following table lists the status of this check:

Status	Description
Pass	Indicates that none of the UCS C240 nodes are running BIOS 4.0.4p.0.1112191020.
Fail	Indicates that one or or more UCS C240 nodes are running BIOS 4.0.4p.0.1112191020.

hc_cli show

The Show method enables you to review previous runs, and regenerate the reporting.

You can run the hc_cli show command to view the results of recent health checks and other diagnostic tests that are run on the cluster.

To view the result of a specific health check, run the following command:

```
hc_cli show <health_check_test_ID>
```

Options

You can use the following options with the hc_cli show command using the format,

```
hc_cli show <option>:
```

Option	Description
`-h or --help`	Displays help output for the hc_cli show command.
`-C or --current-node`	Displays the health check output from only the current node.
`-F or --failed`	Displays only failed health checks.
`-V or --verbose`	Displays detailed output of all checks across all nodes.
`--warning`	Displays detailed output of health checks that show warnings.
`--skipped`	Displays detailed output of health checks that were skipped.
`hc_cli show --all`	Displays detailed output of all health checks.

hc_cli debug

The hc_cli debug command provides a detailed logging or verbose output to help diagnose issues within the health check process.

You can use the debug method within the CLI shell or outside of it.

You can run the following command to enter the debug shell:

```
$ hc_cli debug
```

Running the above command will drop you into the debug shell where you can execute various maintenance commands.

Following are the modules available with in the debug shell:

- Apps
- Archival

- Backup
- Cert
- Data Management
- Data Sources
- Gflag
- Indexing
- Log Utilities
- Metadata
- Networking
- OS U
- Platform

hc_cli version

The Version method prints the current version of the support toolbox installed on the Cohesity cluster.

You can run the following command to know the version of the support toolbox installed:

```
$ hc_cli version
```

Access Support Toolbox from Siren Web UI

Siren is a web-based tool that provides access to logs and other serviceability information on a Cohesity cluster. You can access Support Toolbox through the Siren Web UI and perform health checks and resolve cluster issues.

Perform the following steps to access Support Toolbox:

1. Login to the cluster's Siren UI. The user logging into the cluster must have either admin credentials, or cluster support roles defined for their user.

   ```
   <cluster_ip_address>/ siren/v2
   ```

   Where, <`cluster_ip_address`> is the IP address of your cluster.
2. From the Siren Landing page, select **Support Toolbox**.

Support Toolbox Operations

The Support Toolbox in the siren UI provides diagnostic and maintenance utilities such as:

- Health Check
- Results
- Maintenance

Health Check

From the **Health Check** Tab, you can perform the health checks for all the modules across all nodes in the Cohesity cluster. The modules on the Cohesity cluster are as follows:

- fingerprint
- services
- vaas
- hardware
- performance
- preUpgrade
- thirdparty
- networking
- archival
- vendor
- software
- Backup

To perform the health check:

1. From the Siren UI of your cluster, navigate to Support Toolbox.
2. Click **Health Check**.

   The page displaying the list of modules for which you can perform the heath check appears.
3. Select the modules for run.

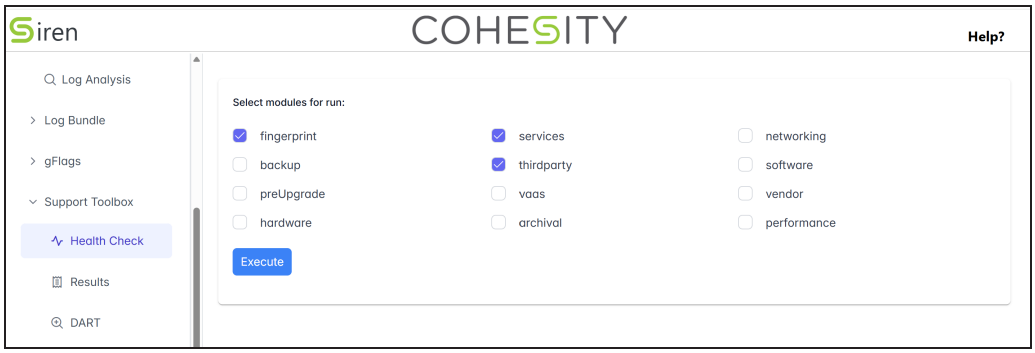

4. Click **Execute**.

   The health check has been initiated. The execution may take a few minutes to complete. Once finished, a View Results link will appear, directing you to the Results page.

Results

The **Results** page allows you to review the most recent health check results, as well as a limited history of past runs. You can:

- Sort results (ascending or descending) by **Severity**, **Node Name**, or **Test Name**.
- Download the results in JSON format to attach to a support case.

Result Categories

Health check results are organized into the following categories:

- **Passed** – No failures detected.
- **Warning** – An issue was found, but it has not reached a full failure state. Consider it an early indication of a potential failure.
- **Fail** – The test has failed and remediation is required.
- **Skip** – The test is irrelevant to the current platform, hardware, or codebase.
- **Error** – The test itself encountered an error.

Severity Levels

Each test is assigned a severity level from 1 (Imminently Critical) to 4 (Informational).

A **Severity 1 Failure** is the most urgent and must be addressed immediately.

A **Severity 4 Failure** is the least critical and primarily informational.

Result Details

Expanding an individual test result displays the following details:

Result – Human-readable description of the issue.

Test Name – The command or script executed.

State – The outcome of the test (Pass, Fail, or Warning).

Severity – The criticality of the issue, from 1 (Most Critical) to 4 (Least Critical).

Node IP – IP address of the reporting node.

Node ID – Cohesity Cluster Config ID of the reporting node.

Node Name – Hostname of the reporting node.

Test KB Link – Link to a Cohesity Support KB article explaining the test and possible resolutions.

View Health Check Results

1. To view results of the health checks performed:
2. From the siren UI of your cluster, navigate to **Support Toolbox**.
3. Click **Results**.

4. From the **Select the Run Id** drop-down list, choose the health check executed on the desired date and time.
5. The results of the checks performed for the selected run are displayed.

You can also download the results in JSON format by clicking the **Download** button located at the top-right of the page.

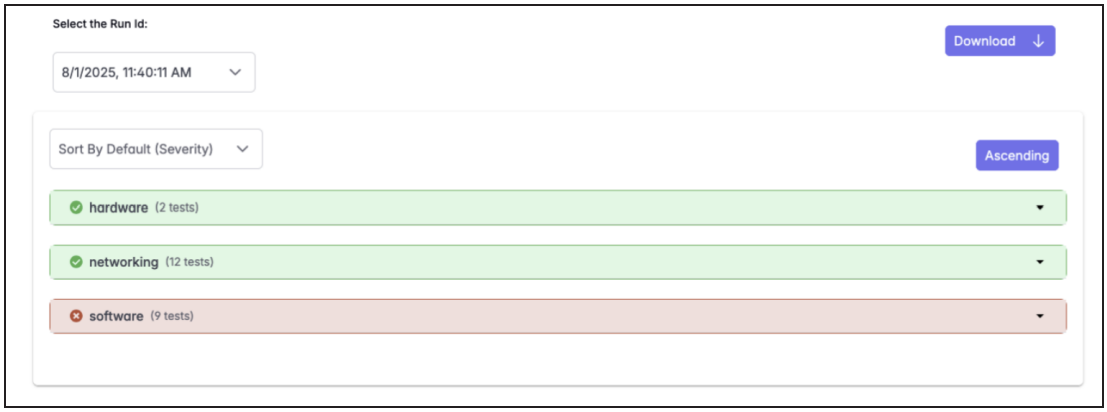

You can expand health check results by click on them.

Any module presented in red has a failure, yellow a warning, green is all passed.

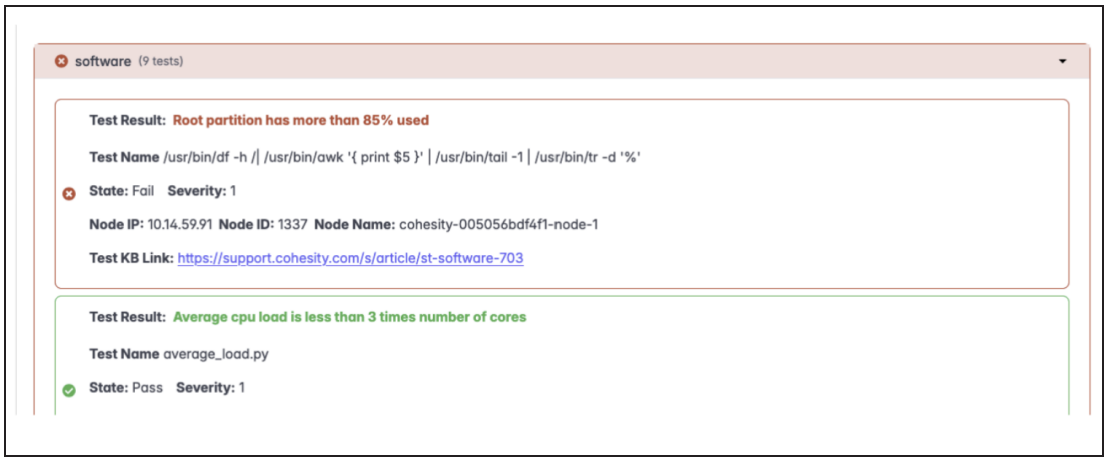

Maintenance

From the **Maintenance** tab, you can perform the following actions:

- Restart Brigid API Server
- Update Support Toolbox
- Rolling Reboot

Restart Brigid API Server

By selecting the Restart Brigid API Servicer button, reports the status of the underlying API server.

Update Support Toolbox

Using the **Update Support Toolbox** option, you can upgrade the Support Toolbox to the latest version.

To upgrade the Support Toolbox to the latest version, perform the following steps:

1. Go to the Cohesity download page, and select **Cohesity Tools** tab.
2. Select the fast link for support-toolbox-installer and download to your local computer.
3. From the Siren UI of you cluster, navigate to **Support Toolbox** > **Maintenance** > **Update Support Toolbox.**.
4. A window will pop up, prompting you to provide either a single **.TGZ** bundle, or a single RPM. We recommend using the **.TGZ** bundle from the downloads portal for simplicity. Click the choose button and select your **.TGZ** file.

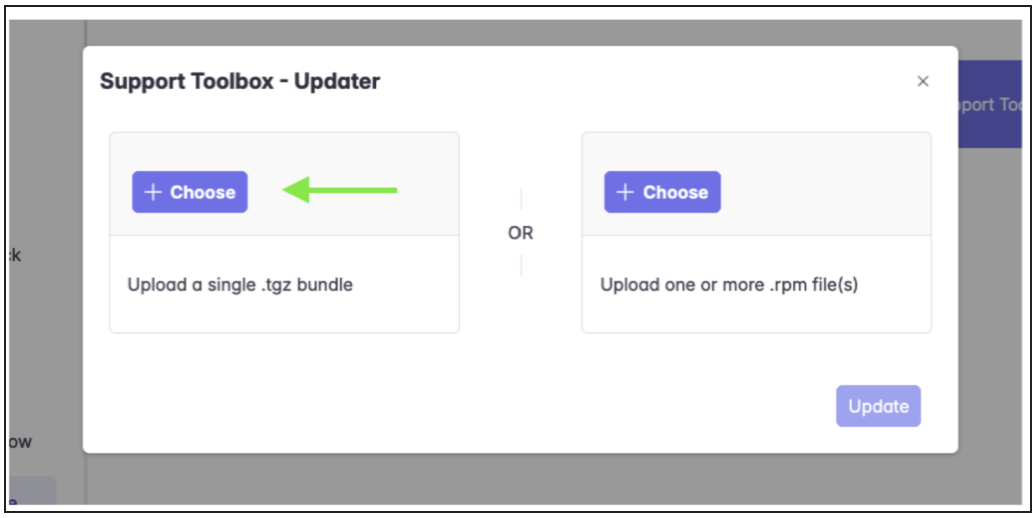

5. Click **Update**.

   A progress screen will be displayed, showing the update status for each package.

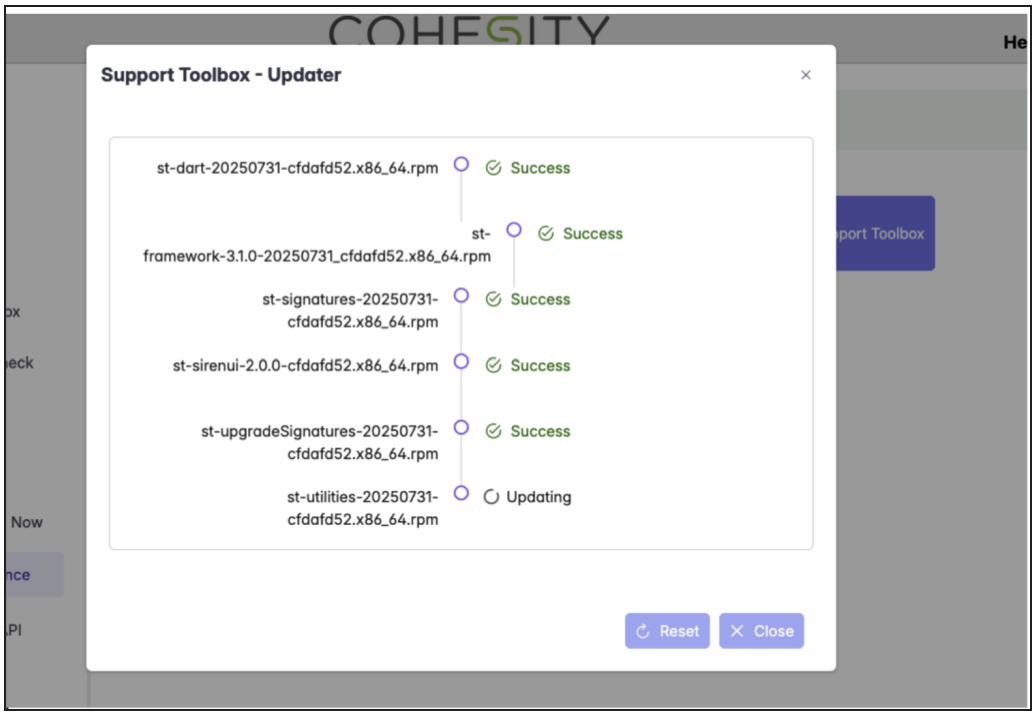

6. When the update is complete, a confirmation message appears indicating that the package update was successful.

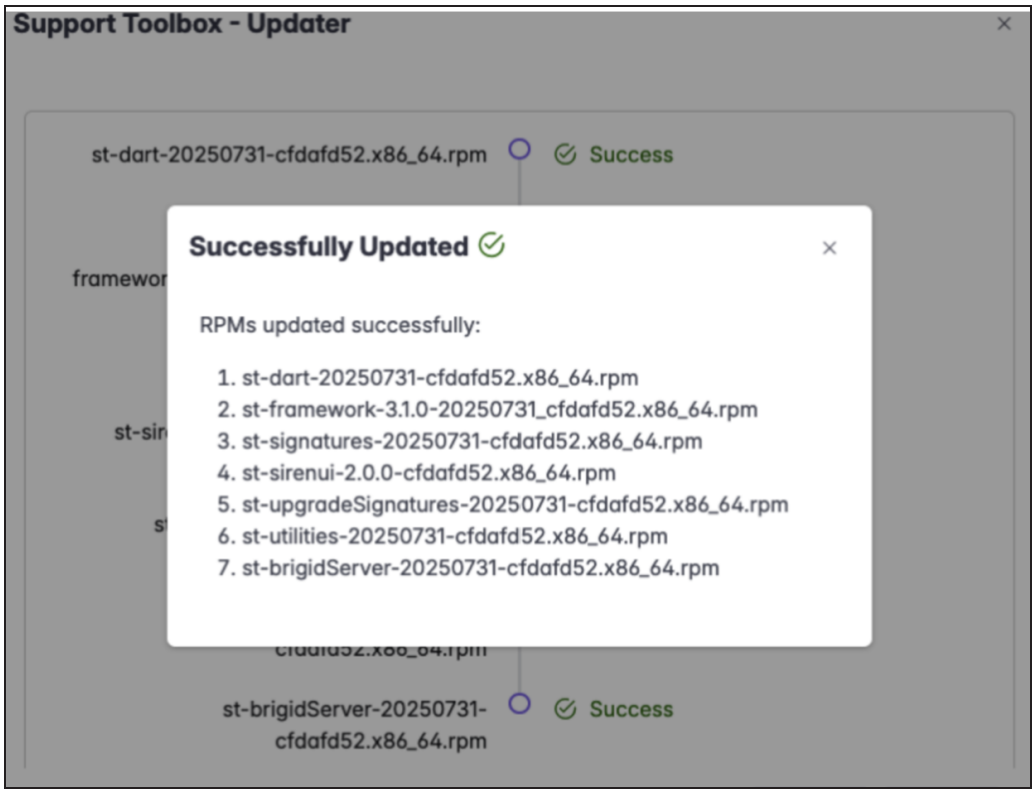

Rolling Reboot

The **Rolling Reboot** option allows you to reboot cluster nodes sequentially. This process helps identify disks that spin down during the reboot but fail to spin up again, potentially causing upgrade operations to hang.

> **Note:** Cluster nodes are rebooted sequentially, not in parallel.

To reboot the cluster nodes, perform the following steps:

1. Select one or more nodes on which you want to perform the rolling reboot.

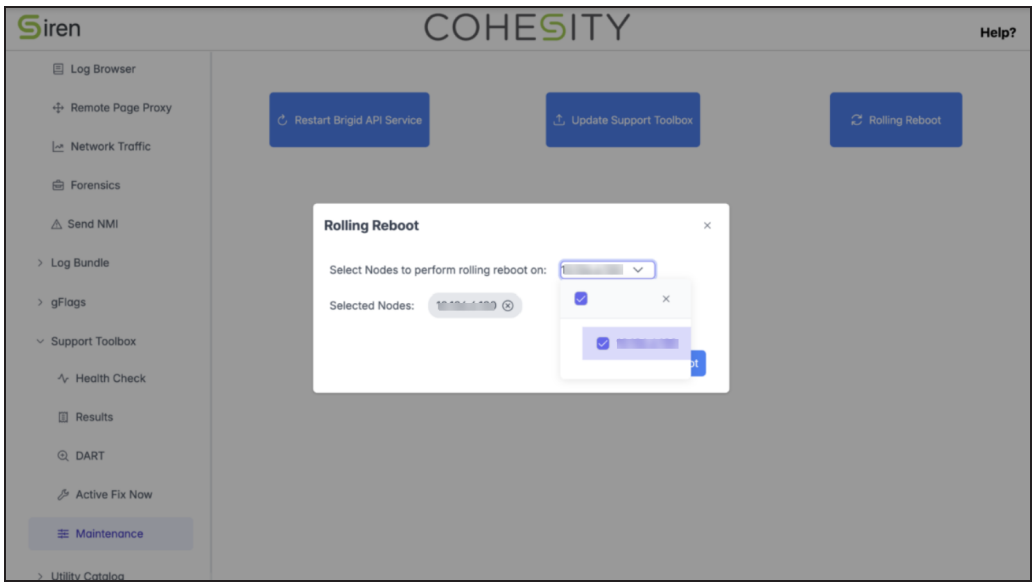

2. Click **Reboot** to initiate the process.
3. A confirmation dialog appears. Click **Yes** to proceed.

4. The rolling reboot begins. A timeline displays the status of each node during the process.

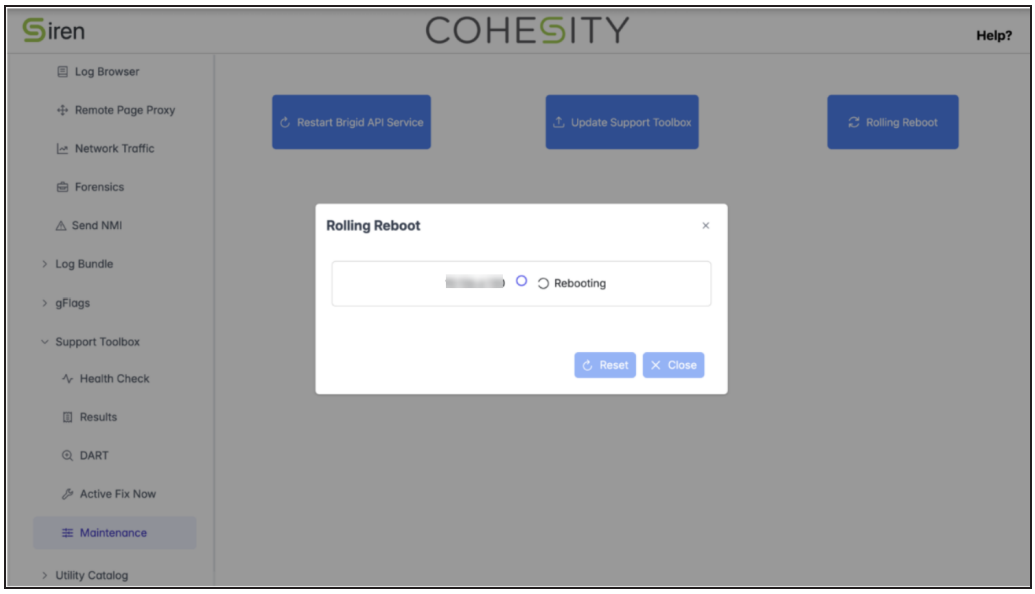

5. After completion, a status dialog appears showing the outcome.

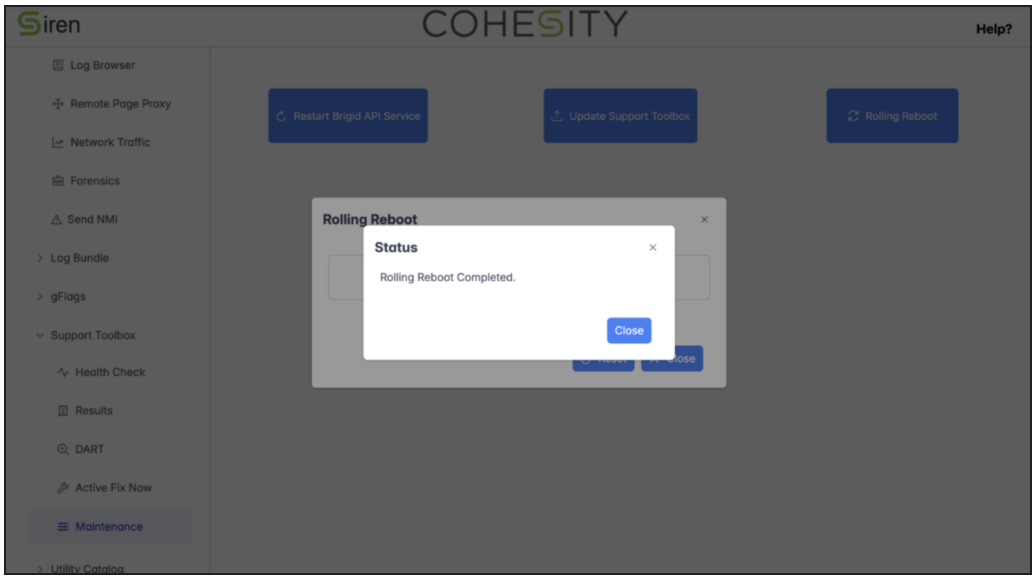

Each node's timeline indicates whether the reboot:

- Completed successfully
- Timed out
- Ended with an error

Environment Setup

This section provides information about setting up environments that interface with the Cohesity cluster.

Backup Manager Service Replica

A Cohesity multi-node cluster consists of a primary instance of the Backup Manager Service. The Backup Manager Service is responsible for handling all the requests from the Cohesity and APIs.

A multi-node cluster environment with a large amount of data increases the load on the primary instance of the Backup Manager Service. Heavy workloads on the primary instance of the Backup Manager Service increase the response time of Cohesity and API calls.

Enabling the Backup Manager Service Replica allows the primary instance of the Backup Manager Service to serve a high volume of traffic. Each Cohesity cluster can have only one Backup Manager Service Replica.

Benefits of enabling Backup Manager Service Replica on a multi-node Cohesity cluster are:

- Offloads the read requests on the primary instance of the Backup Manager Service.
- Improves performance of Cohesity.
- Lower API latency.

Cohesity recommends enabling Backup Manager Service Replica on a multi-node cluster only.

> **Note:** Backup Manager Service Replica reads the primary instance of the Backup Manager Service once in 15 seconds.

Ensure Adequate Privileges for Cohesity on the Source

The Cohesity cluster must perform various actions on the source. For the Cohesity cluster to perform these actions, the user specified to connect to the source (the one used to register the source) must have adequate privileges. For information about each source type, refer to the appropriate section:

- VMware vCenter
- Standalone ESXi
- Microsoft Hyper-V
- vCloud Director
- Nutanix Acropolis
- NetApp
- Dell EMC Isilon
- Storage Snapshot Integration
- Active Directory

- Linux Servers
- Windows Servers

Create a Role with Necessary vCenter Server Privileges

The Cohesity cluster performs various VMware actions such as creating VMware snapshots, restoring VMware snapshots to ESXi hosts using vSphere Storage vMotion and creating new VMs by mounting the new VMDKs. For the Cohesity cluster to have adequate vCenter server privileges, the VirtualCenter user specified to connect to the source (vCenter Server) must have the role privileges listed in Role Privileges for vCenter Server and that VirtualCenter user must be associated with vCenter objects that the Cohesity cluster interfaces with.

> **Note:** By default, the **VMware Consolidated Backup user (sample)** role does not have adequate privileges for the VMware activities performed by the Cohesity cluster.

Create a Role with adequate vCenter Server Privileges

1. Using the vSphere Web Client, log in to the vCenter server that contains the VM to protect.
2. Create a new role.
3. Select the appropriate privileges for the new role as listed below.
4. Assign this role to the vCenter users that will access the Cohesity cluster. Assign this role to users at the top vCenter object level and enable **Propagate to children** check box.
5. When registering a source on the Cohesity cluster, specify a VirtualCenter user assigned to a role with the required Cohesity privileges (defined below).

For more information on VMware vCenter permissions, refer the knowledge base article, Required Cohesity Permissions for VMware vCenter Environments.

Role Privileges for vCenter Server

The following table lists all the privileges that you need to use a Cohesity cluster with VMware vCenter.

Note:

- The Session view and stop sessions privilege is applied to all the tasks to close any sessions after a process finishes. The Cryptographic privileges for Add Disk and Direct Access are only needed for encrypted VMs. The Configure Datastore privilege is only needed for throttling or storage snapshot integration.
- Starting with VMware vSphere 8.0 version, the **Profile-driven Storage** privilege level is replaced with **VM Storage Policies**.

Privilege Level	Permissions Required When Using	
	VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
Cryptographic operations	• Add Disk • Direct Access • Encrypt new • Migrate • Manage key servers **Note:** Cryptographic privileges are only needed for encrypted VMs.	• Add Disk • Direct Access • Encrypt new • Migrate • Manage key servers **Note:** Cryptographic privileges are only needed for encrypted VMs.

Privilege Level	Permissions Required When Using	
	VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
Datastore	• Allocate space • Browse datastore • Configure datastore • Low level file operations • Move datastore • Remove file • Configure datastore IO management **Note:** Configure Datastore privilege is only needed for throttling.	• Allocate space • Browse datastore • Configure datastore • Low level file operations • Move datastore • Rename datastore • Remove file • Configure datastore IO management **Note:** Configure Datastore privilege is only needed for throttling.
Folder	• Create folder • Delete folder	• Create folder • Delete folder

Privilege Level		Permissions Required When Using	
		VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
Global		• Disable Methods • Enable Methods • Licenses • Log event • Manage custom attributes • Set custom attribute	• Disable Methods • Enable Methods • Licenses • Log event • Manage custom attributes • Set custom attribute
Host	Configuration	• Maintenance • Query Patch • Storage partition configuration	Storage partition configuration
	Local operations	N/A	Delete virtual machine
Network		Assign network	Assign network
Profile-driven Storage		• Profile-driven storage update • Profile-driven storage view	

Privilege Level	Permissions Required When Using	
	VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
Resource	• Assign virtual machine to resource pool • Migrate powered off virtual machine • Migrate powered on virtual machine	• Assign virtual machine to the resource pool • Migrate powered off virtual machine • Migrate powered on virtual machine
Session	View and stop sessions	View and stop sessions

Privilege Level		Permissions Required When Using	
		VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
Virtual Machine	Change Configuration	• Acquire disk lease • Add existing disk • Add new disk • Add or remove device • Advanced configuration • Change Settings • Change Swapfile placement • Configure Raw device • Display connection settings • Modify device settings • Remove disk • Rename* • Set Annotation • Toggle disk change tracking • Upgrade Virtual machine compatibility	• Acquire disk lease • Add existing disk • Add new disk • Add or remove device • Advanced configuration • Change Settings • Change Swapfile placement • Configure Raw device • Modify device settings • Remove disk • Rename* • Set Annotation • Toggle disk change tracking • Upgrade Virtual machine compatibility

Privilege Level		Permissions Required When Using	
		VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
	Change Configuration (For Runbook)	• Change CPU count • Change Memory • Change Settings • Change resource • Modify device settings • Rename	• Change CPU count • Change Memory • Change Settings • Change resource • Modify device settings • Rename
	Edit Inventory	• Create new • Register • Remove • Unregister	• Create new • Register • Remove • Unregister
	Guest Operations	• Guest operation modifications • Guest operation program execution • Guest operation queries	• Guest operation modifications • Guest operation program execution • Guest operation queries
	Interaction	• Connect devices • Guest operating system management by VIX API • Power off • Power on	• Connect devices • Guest operating system management by VIX API • Power off • Power on

Privilege Level		Permissions Required When Using	
		VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
	Provisioning	• Allow disk access • Allow read-only disk access • Allow virtual machine download • Customize guest (For Runbook) • Mark as template	• Allow disk access • Allow read-only disk access • Allow virtual machine download • Customize guest (For Runbook) • Mark as template
	Snapshot management	• Create snapshot • Remove snapshot • Revert snapshot	• Create snapshot • Remove snapshot • Revert snapshot
vApp		• Add virtual machine • Assign resource pool • Unregister	• Add virtual machine • Assign resource pool • Unregister

Privilege Level	Permissions Required When Using	
	VMware Snapshot Only	Storage Snapshot Integration (IBM, Pure, Nimble, NetApp)
vSphere Tagging	• Assign or Unassign vSphere Tag • Assign or Unassign vSphere Tag on Object • Create vSphere Tag • Create vSphere Tag Category • Delete vSphere Tag • Delete vSphere Tag Category • Edit vSphere Tag • Edit vSphere Tag Category	• Assign or unassign tag • Create vSphere Tag • Create vSphere Tag Category • Delete vSphere Tag • Delete vSphere Tag Category • Edit vSphere Tag • Edit vSphere Tag Category
VM storage policies	• Apply VM storage policies • Update VM storage policies • View VM storage policies	• Apply VM storage policies • Update VM storage policies • View VM storage policies

*Rename permission is required for a copy recovery.

Role Privileges for Standalone ESXi

The role privileges required by the Cohesity cluster for standalone ESXi are listed below.

Privilege Level		Required Permissions
Datastore		AllocateSpace Browse Config–Required if Source Datastore throttling is enabled Delete–Required if Source Datastore throttling is enabled DeleteFile FileManagement Move–Required if Source Datastore throttling is enabled Rename–Required if Source Datastore throttling is enabled UpdateVirtualMachineFiles–Required if Source Datastore throttling is enabled UpdateVirtualMachineMetadata–Required if Source Datastore throttling is enabled
Folder		Create Delete
Global		DisableMethods EnableMethods Licenses LogEvent Manage custom attributes Set custom attribute
Host	Configuration	Storage
	Local operations	Delete virtual machine
Network		Assign
Resource		AssignVMToPool ColdMigrate HotMigrate
System		Anonymous Read View
vApp		AssignResourcePool AssignVM Unregister

Privilege Level	Required Permissions
Session	View and stop sessions

Privilege Level		Required Permissions
Virtual machine	Configuration	AddExistingDisk AddNewDisk AddRemoveDevice AdvancedConfig Annotation CPUCount ChangeTracking DiskLease EditDevice HostUSBDevice Memory RawDevice ReloadFromPath RemoveDisk Rename ResetGuestInfo Resource Settings SwapPlacement Upgrade Virtual machine compatibility
	GuestOperations	Execute Modify Query
	Interact	GuestControl PowerOff PowerOn
	Inventory	Create Delete Register Unregister
	Provisioning	DiskRandomRead GetVmFiles
	State	Create snapshot Remove snapshot Revert to snapshot

Privilege Level	Required Permissions
Cryptographic Operations	Add Disk Direct Access Encrypt Migrate

Microsoft Hyper-V

For Hyper-V SCVMM, specify the domain user with administrative privileges when registering the SCVMM application. Cohesity recommends you to specify a user with the LOCAL SYSTEM account privileges when installing the Cohesity Agent. For more information, see Download and Install the Agent on the SCVMM Server.

For Hyper-V Standalone and Hyper-V Failover clusters, perform the following actions:

- Add all hosts' machine accounts:
  1. Start the **Failover Cluster Manager**.
  2. Right-click the standalone cluster from the clusters list and select **Properties > Cluster Permissions > Add > Object Types > Check Computers > OK**.
  3. Type <the hostname>$ in **"Enter the object names to select"** and select **Check Names > OK > Allow "Read" > OK**.
- Add all machine accounts to the administrators group of each host in the standalone cluster.

vCloud Director

The following table lists all the privileges that you need to use a Cohesity cluster with vCD.

Permission	Application	Privilege	
		View	Manage
Access Control	Organization	Allow Access to All Organization VDCs	Change the Owner of the Content Hub External Source
		View the Content Hub External Source	Edit the External Source in Content Hub
		View Content Hub External Source ACL	Share the Content Hub External Source

Permission	Application	Privilege	
		View	**Manage**
		View Catalog ACL	Manage Quotas of Organization
		View the LDAP settings	Implicitly Import User/Group from IdP while Editing VDC ACL
		View Organization Networks	Administer and traverse into other organizations
		View Organization Administrative Details	Use Provider LDAP as Tenant
		*View Access Control List of Organization VDCs	Edit Organization Properties
		*View Organization Associations	
		*View OAuth Settings	
		View vApp ACL	
	User	View Users and Groups	Manage user's own API token
Administration	General	Administrator View	Administrator Control
		View Error Details	Resume / Cancel / Fail a Task
		Perform administrator queries	
		View Quota Policy Capabilities	
		Test Connection	
		View task information	
Compute	Organization VDC	View Alternate Admin version of Compute Policies for an Organization VDC	Open Organization VDC Network in vSphere

Permission	Application	Privilege	
		View	Manage
		View Compute Policies for an Organization VDC	Edit Disk IOPS
		View Disk IOPS	Change Owner
		View Disk Encryption Status	Create a Disk
		View Disk Properties	Delete a Disk
		View Organization VDC Resource Pool	Edit Disk Properties
		View Organization VDC Storage Policy Capabilities	Move a Disk
		View Organization VDC Extended Properties	Create a Shared Disk
		View Organization VDCs	Edit Organization VDC Storage Policy
		View Organization VDC Memory and CPU Reservation	Disable Organization VDC Storage Policy
		View stranded items	Open Organization VDC Storage Policy in vSphere
		Limited Admin View	Remove Organization VDC Storage Policy
			Set Default Storage Policy
			Create Organization VDC
			Delete Organization VDCs
			Edit Organization VDC Name and Description

Permission	Application	Privilege	
		View	Manage
			Enable / Disable Organization VDC
			Edit Organization VDC Extended Properties
			Manage stranded items
	Provider VDC	View Compute Policies for a Provider VDC	Manage Compute Policies for a Provider VDC
		View Provider VDC Resource Pool	Migrate Provider VDC Resource Pool VMs
		View Provider VDC Storage Policy	Open Provider VDC Resource Pool in vSphere
		View Provider VDC	Edit Provider VDC Storage Policy
		Limited Provider VDC Storage Policy View	Enable / Disable Provider VDC Storage Policy
		Limited Provider VDC View	Open Provider VDC Storage Policy in vSphere
			Remove Provider VDC Storage Policy
	vApp	View vApp Shadow VMs	Change vApp Template Owner
		View Encryption Status of VMs and VM's disks	Force storage lease expiration for vApp templates
		View VM metrics	Import vApp Template
			Open vApp Template in vSphere
			Change Owner

Permission	Application	Privilege	
		View	Manage
			Copy a vApp
			Create / Reconfigure a vApp
			Delete a vApp
			Download a vApp
			Edit vApp Properties
			Edit VM Compute Policy
			Edit VM CPU
			Edit VM Hard Disk
			Edit VM Memory
			Edit VM Network
			Edit VM Properties
			Enter / Exit vApp Maintenance Mode
			Force runtime lease expiration for vApps
			Force storage lease expiration for vApps
			Import vApp
			Manage maintenance mode for vApps
			Manage VM Password Settings
			Open vApp in vSphere

Permission	Application	Privilege	
		View	Manage
			Start / Stop / Suspend / Reset a vApp
			Share a vApp
			Create / Revert / Remove a Snapshot
			Upload a vApp
			Access to VM Console
			Edit / View VM Boot Options
			View Compliance of vApp VMs
			Migrate / Force Undeploy / Relocate / Consolidate vApp VMs
			Allow metadata mapping domain to vCenter
Extensions	Metadata		Metadata File Entry: Create/Modify
	VMware Cloud Director Extension	View Tenant Portal Plugin Information	Create / Update / Delete/ Upload / Associate / Disassociate Tenant Portal Plugin
		View VMware Cloud Director Extension	Register / Unregister / Refresh / Associate / Disassociate Extension
Infrastructure	Datastore	View Datastore	Delete Datastore
			Edit Datastore

Permission	Application	Privilege	
		View	Manage
			Enable/Disable Datastore
			Open Datastore in vSphere
	Host	View Host	Open a Host in vSphere
	Resource Pool	Open Resource Pool	
		View Resource Pool	
	vCenter	View vCenter	Reload a VM from vSphere
		View vCenter server	Open vCenter Server in vSphere
			Refresh vCenter
			Manage vCenter server
Libraries	Catalog	View Catalog vApp Shadow VMs	Add a vApp from My Cloud
		View Private and Shared Catalogs within Current Organization	Change Owner
		View Shared Catalogs from Other Organizations	Manage Multi vCenter VM
		View the content library system settings	
	Catalog Item	View vApp Templates / Media	Copy / Move a vApp Template / Media
			Create / Upload a vApp Template / Media
			Edit vApp Template / Media Properties
			Add to My Cloud

Permission	Application	Privilege	
		View	Manage
			Enable vApp Template / Media Download
Networking	Organization VDC Network	View Properties	Manage Direct Organization VDC Network
			Edit Properties
			Import Network
			Create Shared Organization VDC Network
	Provider Gateway Services	View Limited Provider Gateway	
	Provider Network	View Provider Network	Create / Delete Provider Network
			Edit Provider Network
			Open Provider Network in vSphere

*These permissions apply only to users who access the plugin through a tenant login. Users performing backup and recovery operations directly on the Cohesity platform do not require them. Additionally, these permissions are needed in addition to the existing permissions when enabling tenant users to access and interact with the plugin.

Nutanix Acropolis

Specify either the local Prism admin user or a service account that has admin privileges to ensure that the user has adequate privileges to perform actions on the source.

Set up a service account

1. In Prism create a link to the Active Directory (AD) domain that you want to use for authentication.
    a. Click the settings icon and select **Authentication**.
    b. Click **New Directory**.

   c. Complete the Authentication Configuration fields and click **Save**.
2. Map the AD account to the Cluster Admin role.
   a. Click the settings icon and select **Role Mapping**.
   b. Click **New Mapping**.
   c. Complete the Create Role Mapping fields.
      - For **LDAP Type**, select User.
      - For **Role** select Cluster Admin.
   d. Click **Save**.
3. Through the Prism self-service portal, assign SSP administrator privileges to the user.

AHV Prism Element

The minimum privileges required for a custom user role are **Cluster Admin**.

> **Note:** Cluster Administrators are automatically granted Backup Administrator privileges.

AHV Prism Central

Starting with version 7.2.2_u2, the minimum privileges required for AHV Prism Central are listed in the table below.

Entity type	Operations
AHV VM	Allow Virtual Machine Power On Allow Virtual Machine Power Off Update Virtual Machine Update Virtual Machine Boot Config Update Virtual Machine Categories Update Virtual Machine CPU Update Virtual Machine Description Update Virtual Machine Disk List Update Virtual Machine Power State View Virtual Machine
Category	View Name Category View Value Category
Cluster	View Cluster
Domain Manager (Prism Central)	View Prism Central

NetApp

To protect NetApp SMB, specify a user that has adequate privileges to perform actions on the source.

Set up a user to protect SMB

1. Log in to the NetApp cluster as an administrator.
2. If registering a cluster, add a cluster role that has the following attributes:

Command	Access Level	Details
export-policy (optional)	all	Adds the Cohesity cluster subnet to the export-policy so that Cohesity can mount volumes
volume snapshot	all	Allows fetching, creating and deleting snapshots for volumes
cifs	readonly	Fetches information about CIFS/SMB shares for volumes
cluster identity	readonly	Fetches information about the cluster
network interface	readonly	Fetches information about network interfaces that Cohesity connects to for mounting volumes
volume	readonly	Fetches information about volumes
vserver	readonly	Fetches information about SVM

3. If registering a SVM, add a SVM role that has the following attributes:

Command	Access Level	Details
vserver export-policy	all	Adds the Cohesity cluster subnet to the export-policy so that Cohesity can mount volumes
volume snapshot	all	Allows fetching, creating and deleting snapshots for volumes
vserver cifs	readonly	Fetches information about CIFS/SMB shares for volumes
network interface	readonly	Fetches information about network interfaces that Cohesity connects to for mounting volumes

Command	Access Level	Details
volume	readonly	Fetches information about volumes
vserver	readonly	Fetches information about SVM

4. Create a cluster/SVM user and assign the role you created previously. The user must have ONTAPI allowed as one of the User Login methods.

> **Note:** The username used to register the NetApp cluster and SVM is case sensitive.

5. To enable the Cohesity cluster to back up volumes on a share, do one of the following:
   - Set permissions on a share so that any necessary Active Directory principals have at least Read access.
   - Add an Active Directory principal that has at least read access for the c$ share (if this share on the SVM is available)

To protect NetApp NFS, specify a user that has privileges to perform actions on the source.

Set up a user to protect NFS

1. Log in to the NetApp cluster as an administrator.
2. If registering a cluster, add a cluster role that has the following attributes:

Command	Access Level	Details
export-policy	all	Adds the Cohesity cluster subnet to the export-policy so that Cohesity can mount volumes
volume snapshot	all	Allows fetching, creating and deleting snapshots for volumes
cluster identity	readonly	Fetches information about the cluster
network interface	readonly	Fetches information about network interfaces that Cohesity connects to for mounting volumes
volume	readonly	Fetches information about volumes
vserver	readonly	Fetches information about SVM

3. If registering a SVM, add a SVM role that has the following attributes:

Command	Access Level	Details
vserver export-policy	all	Adds the Cohesity cluster subnet to the export-policy so that Cohesity can mount volumes
volume snapshot	all	Allows fetching, creating and deleting snapshots for volumes
network interface	readonly	Fetches information about network interfaces that Cohesity connects to for mounting volumes
volume	readonly	Fetches information about volumes
vserver	readonly	Fetches information about SVM

4. Create a cluster/SVM user and assign the role you created previously. The user must have ONTAPI allowed as one of the User Login methods.

> **Note:** The username used to register the NetApp cluster and SVM is case sensitive.

> **Note:** For the applicable NetApp volumes, Cohesity adds the node IPs in the export rule that has the **Allow Superuser Access** option enabled for NFS.

Dell EMC Isilon

User Permissions to Register Isilon

To register an Isilon NAS with Cohesity Isilon Adapter, ensure that the following permissions are configured in the Isilon cluster:

Name	Description	Access
Platform API	To have access to Isilon's APIs	Read-Only
Auth	To verify users and passwords	Read-Only
Cluster	To obtain cluster identity and settings	Read-Only

Name	Description	Access
Job Engine	To read and write ChangeList jobs	Read/Write
Network	To obtain the network interfaces	Read-Only
SMB	To read the settings in SMB server	Read-Only
Snapshot	To fetch, create, and delete snapshots for shares and exports	Read/Write
NFS	To read and write settings to and from the NFS server. **Note:** This setting modifies the NFS export used to mount, such as /ifs.	Read/Write **Note:** If NFS backup is not required, the user must have **Read-Only** permission.

Snapshot IQ License

You must have the SnapshotIQ license to take a snapshot on Isilon.

Storage Snapshot Integration

The Cohesity cluster integrates natively with IBM, Pure Storage or Nimble Storage to improve snapshot-based data protection. To leverage storage snapshot integration, Cohesity cluster needs to perform the following tasks:

1. Discover, configure, and register iSCSI LUN.
2. Perform VMFS datastore updates and configure storage I/O control on datastore.
3. Rename, mount/unmount, and delete the datastore.
4. Remove iSCSI LUN.
5. Delete VM and folder.

User Permission

For storage snapshot integration, the following additional privileges are required along with the permissions described in Role Privileges for vCenter Server .

Privilege Level		Required Permissions
Virtual Machine	Change Configuration	Acquire disk lease Add or remove device Change Settings
	Provisioning	Allow disk access

Active Directory

For more information, see Active Directory Protection Requirements.

Linux Servers

For more information on source privileges for Linux servers, see Install and Manage the Agent on Linux Servers.

Windows Servers

Specify the local or AD user with administrative privileges on the server. For more information, see Install and Manage the Agent on Windows Servers.

For information on minimum permissions, prerequisites and considerations to protect physical servers with Cohesity cluster, see Plan and Prepare for Physical Server Protection.

Antivirus Software Exclusions

To prevent antivirus software from interfering with backup and recovery operations involving Cohesity Agent on your Windows or Linux servers, Cohesity recommends adding the following files to the antivirus software exclusion list.

The following utilities can be added to the exception list.

For Windows:

For **Backups**: Provide the agent deployment path.

Examples:

- `C:\Program Files\Cohesity\cohesity_windows_agent_service.exe` (This is the installed Agent.)
- `C:\_Cohesity_\windows_agent\cohesity_windows_agent.exe` (This is the auto-deployed Agent.)

For **Recovery**:

- `*restoreFilesUtil.exe` (the * represents the variable task ID used as a file name prefix)

- `volumes_util.exe C:\Cohesity\windows_agent\cohesity_windows_agent.exe` (This is the auto-deployed Agent.)

The above utilities may be deployed during the recovery process and copied to a temporary directory such as `C:\Users\<Username>\AppData\Local\Temp\Cohesityvmware`*`nnnnnn`*, where *`nnnnnn`* is a variable task ID.

For Linux:

The following table provides the list of files that need to be added to the antivirus software exclusion list:

Installation Type	Base directory	Files
RPM	/opt/cohesity/agent/software/crux/bin/	/opt/cohesity/agent/software/crux/bin/linux_agent_exec
Script	/home/cohesityagent/cohesityagent/software/crux/bin/	/home/cohesityagent/cohesityagent/software/crux/bin/linux_agent_exec

> **Note:** Add the paths above to the antivirus exclusion list to prevent the antivirus software from interfering with backup and recovery operations.

Access Cohesity Host OS

You can access the host OS on a node through the Intelligent Platform Management Interface (IPMI) remote console and Secure Shell by using **cohesity_console** and **support user** accounts respectively.

Cohesity restricts the access to host OS through multi-layer of security that requires login credentials, and a challenge and response workflow.

About Challenge and Response Text

A challenge is an alphanumeric text that is displayed after you attempt to log in to the Cohesity host OS with your credentials through an IPMI remote console or a VE node's console. The challenge appears in the following format.

-----Challenge String-------

<Alphanumeric_Value>

-----Challenge String End-------

This challenge is an 81-character text. It is valid for 60 minutes if you have accessed the Cohesity Platform through an IPMI console.

To proceed further, you have to copy this challenge and share it with Cohesity Support. The Support team verifies the request and shares a 16-character alphanumeric response.

On the IPMI remote console or the Secure Shell, a prefix is also displayed along with the challenge. The prefix should be combined with the response you have received from the Cohesity Support team to generate the login password and give you access to the host OS.

> **Note:** The prefix should not be shared with Cohesity.

Access Host OS through IPMI remote console

Follow the steps below to access the host OS on a node via IPMI.

1. Create a support ticket to access the host OS. For more information, see Cohesity Support.
2. Launch your preferred IPMI utility tool IPMI remote console.
3. Enter the IP address of the node with the required command. For more information on using Secure Shell and IPMI remote console, see Using the Secure Shell and Using the Cohesity CLI.

   The message **Welcome to Cohesity OS!** is displayed with details about accessing the host OS. Towards the end of the message, a challenge text is displayed in the format mentioned in the About Challenge and Response Text topic.
4. Copy the challenge text and share it with Cohesity Support through the support ticket.
5. After receiving the challenge, Cohesity Support verifies the request and shares the response text.
6. Enter the **Prefix** followed by the response text from Cohesity Support together as the password.

The console should now display a live session to the host OS on the node.

Hardware Replacement

Hardware replacement procedures (PDFs) are available at: Hardware Universe.

Chapter 5: FortKnox Self-Managed

FortKnox Self-Managed enables you to vault a copy of your backup data in a securely isolated Cohesity cluster using pull based method. It allows you to maintain an isolated, immutable copy of your data within your own self-managed infrastructure. This also ensures protection against cyber threats or primary cluster failures while enabling full control, security, and compliance over the vaulted data.

Key Concepts

Vault Cluster

A vault cluster stores another copy of data from the Cohesity cluster and recovers data to the original or a new location. You can convert a Cohesity cluster into a vault cluster to vault the data. The vault cluster provides isolated copy, immutable, WORM (Write Once, Read Many) storage, ensuring protection against unauthorized modifications, deletions or a cyber event.

Air Gapping

The Cohesity vault cluster operates on a pull-based model, leveraging outbound-only connections from an isolated network and ensuring no unauthorized external access. Connections are strictly established during designated vaulting windows, reinforcing air-gapped data protection.

Storage Immutability & WORM Support

The vaulted data is stored with Write Once, Read Many (WORM), ensuring immutability and protection against unauthorized modifications or deletions.

Quorum-Based Access Control

Helios Self-Managed, running within the secure Vault network, enforces quorum control, requiring multiple approvals for any critical vault operations, enhancing security and governance.

Multi-Tenancy Support

FortKnox Self-Managed supports multi-tenancy, allowing the service provider to manage the mapping of storage domains and vaulting windows for the tenants in a secure infrastructure.

Low Recovery Time Objective (RTO)

FortKnox Self-Managed is optimized for fast data recovery by maintaining a local copy within your premises, minimizing downtime in the event of a cyber-attack.

Vaulting Window

The Vaulting Window in FortKnox Self-Managed is an air-gapping mechanism that defines when the vault cluster can connect to the primary cluster. During the configured time window, connections are opened to pull data. Once the window closes, connections are dropped and active jobs are paused. These jobs automatically resume when the window opens again.

Supported Workloads

FortKnox Self-Managed supports the vaulting workflow for all workloads that are currently qualified for the Replication workflow. The workloads listed in the table below have been qualified for vaulting, while qualification for the remaining workloads is in progress. This list will be updated periodically as additional workloads are qualified.

Qualified Workload Name
VMware ESXi VM Servers
VMware vCloud Director
Physical Files

Qualified Workload Name
Network Attached Storage (NAS) • NetApp NFS • Isilon NFS and SMB • Generic NAS
File-based and VDI-based SQL backups
Cohesity View S3 1.0
Physical
Oracle Database
Azure VMs
Microsoft 365
Nutanix Acropolis
MongoDB
Cassandra Database

Get Started

The following topics provide guidance that helps you to get started with FortKnox Self-Managed.

Prerequisites

Ensure the following prerequisites are met before enabling FortKnox Self-Managed:

- Encrypt Storage Domains
- Enable Cluster Level Encryption
- Enable Node to Node Encryption
- Configure Firewall
- Enable MFA
- Disable Archival and Replication
- Disable Support Channel

Encrypt Storage Domains

Storage Domains must be encryption enabled. Delete the unencrypted storage domains before converting the cluster to the vault cluster.

Encryption can be enabled only when you are creating the Storage Domain. Once the Storage Domain is created, the options to enable encryption or modify its encryption settings are disabled in the Cohesity cluster.

To create a storage domain and enable encryption, see Create or Edit Storage Domains.

Enable Cluster Level Encryption

Enable the software encryption at the cluster level to convert the cluster to vault cluster.

> **Note:** Cluster-level encryption can be enabled only during the cluster creation, and it cannot be disabled or modified later.

To enable the encryption,

1. Navigate to **Settings** > **Summary** and click **Configure**.
2. Turn on the **Encryption** toggle.
3. Scroll down to the end of the page and click **Save**.

> **Note:** Once encryption is enabled for a cluster, it cannot be disabled. Additionally, any storage domains created without encryption before cluster-level encryption was enabled must be deleted for this pre-requisite to succeed.

Enable Node to Node Encryption

Enable Node-to-node encryption for all services.

Node-to-node encryption is enabled by default starting from, 7.1.X and later for Physical Clusters

For Clusters upgraded from earlier versions, run the following command to enable node-to-node encryption manually,

```
$ iris_cli cluster tpenc enable-encryption=true
```

After enabling, run the following command to restart all services to stabilize the cluster,

```
$ iris_cli cluster restart-services all-services=true
```

Configure Firewall

You need to configure firewall, before converting a Cohesity cluster to a vault cluster.

To configure the firewall:

1. Navigate to **Settings** > **Networking** and select the **Firewall** tab.
2. Click edit icon ( ) to edit the incoming firewall rules.
3. Select **Allow** action for **Management**.
4. Select **Deny** action for all other applications.
    1. To deny **SSH**, users must manually select each interface group individually from the drop-down when configuring the deny rule.
    2. To deny other applications, no selection is required from the drop-down, the rule automatically applies to all interface groups by default.
5. After configuring the firewall to deny all firewall rules except management, follow these steps to add a single trusted IP address to the **Management** firewall rule to enforce least security measure.
    1. Select **Custom IP/Subnet** from the **Source** drop-down menu.
    2. Enter any one IP address to access the vault cluster.

> **Important:** Deny action should be set for all applications other than the Management inbound traffic.

Manage Firewall Ports

You need to open the following replication ports on the primary and vault cluster.

Primary Cluster to Vault cluster Traffic

For replication between the Primary cluster and the Vault cluster, the vault cluster's outbound replication ports must be open, and the primary cluster's inbound and outbound ports must be open.

Outgoing Traffic

Source	Destination	Port	Protocol	Usage Notes
Vault cluster	Cohesity cluster	23335	TCP	Required for vaulting scenario.
Vault cluster	Cohesity cluster	23336	TCP	Required for vaulting scenario.

Cohesity recommends keeping IPMI ports turned off.

For more information, see Manage Firewall Ports.

Enable MFA

Before you convert a Cohesity cluster into a vault cluster, you need to enable MFA to provide an additional authentication factor to access the vault cluster.

To enable MFA for local and active directory users:

1. Navigate to **Settings** > **Access Management** and click the **Secure Login** tab.
2. Turn on the **MFA** toggle, select one or both authentication methods, and click **Save**. For more information, see Enable MFA for Local and Active Directory Users.

To enable MFA for the support user account:

1. Navigate to **Settings > Access Management> Support**.
2. In the **Support** tab, click **Unlock to make changes**.
3. In the **Unlock Support User Settings** pop-up window, enter the **support** user password and click **Unlock**.
4. In the **Support User** section:
    - Enter the email address of the **support** user and click **Update Email**.

      > **Note:** A verification code will be sent to the current email to update the email address if you enable MFA with Email. The update completes only after successful verification.

    - To update the Linux Shell password, click **Update Password** and update the password.
5. Turn on the **Multi-factor Authentication (MFA)** toggle, select one or both authentication methods, and click **Save**. For more information, see Enable MFA for Support User Account.

Disable Archival and Replication

Ensure that the Vault cluster does not have any active policies with outgoing Archive or Replication workflows enabled, and do not configure it as an Archive or Replication target for any other Cohesity cluster.

To verify:

Navigate to **Data Protection** > **Policies** and check the existing policies.

Delete all existing replication and archival policies before proceeding if any active policies are present.

> **Note:** Snapshot data does not need to be removed and can be retained for recovery purposes.

Disable Support Channel

You must disable the support channel to convert the cluster to vault cluster.

To disable the support channel:

1. Navigate to **Settings** > **Summary** and locate the **Support Channel** option in the Cluster Summary list.
2. Click edit icon ( ) and select **Off** option from the dropdown.
3. Click **Save** to disable the Support Channel.

Create a Local User with Vault Admin Role on Primary and Vault Clusters

Vault admin role is required to configure and manage vault operations. Create a local user with the **FortKnox Self Managed Vault Admin** role on both the primary and vault clusters.

To create a local user with vault admin role, see Add or Edit a Local User.

Convert a Cluster to Vault Cluster

To vault the data from the primary Cohesity cluster to an isolated Cohesity cluster, first you need to convert the cluster to vault cluster.

To convert a cluster to vault cluster:

1. Log in as the user with FortKnox Self-Managed Vault Admin role.
2. Navigate to **Settings** > **Summary** and select the **Vault** tab.
3. In the **Vault** page, click **Get Started**.
4. Ensure that the eligibility list is ticked on the **Convert to Vault Cluster** dialog box.

   > **Note:** Refer to Prerequisites section if any eligibility criteria is not met.

5. Click **Continue**.
6. Type YES on the confirmation field and click **Complete**.

> **Important:** Once a cluster is converted into a vault cluster, it cannot be converted back into the normal cluster. To convert it back, you need to perform cluster reset operation in which all the data will be deleted from the cluster, and the cluster will be a new cluster.

Now, you will see a confirmation message that the cluster is converted into a vault cluster.

Physical Air-Gap

The Physical Air-Gap setting controls whether the interface group configured for vaulting can be brought down outside the vaulting window. You can enable or disable this setting directly from the Paired Clusters page on the vault cluster.

This setting applies only to the interface group used for vaulting, and only when a non-default interface group is configured for vaulting.

> **Important:** Cohesity recommends keeping Physical Air-Gap enabled at all times, except when the same interface group is used to perform recovery or replication operations from the vault cluster.

Physical Air-Gap States

State	Behavior
Enabled	Recovery or replication operations that use the interface group configured for vaulting are not permitted outside the vaulting window.
Disabled	The interface group configured for vaulting remains available outside the vaulting window, allowing recovery or replication operations from the vault cluster to proceed outside the vaulting window.

Enable or Disable Physical Air-Gap

To change the Physical Air-Gap setting on the vault cluster:

1. On the vault cluster, navigate to **Vault** > **Paired Clusters**.
2. Locate the Physical Air-Gap toggle. The toggle reflects the current state: **Enabled** or **Disabled**.
3. Click the toggle to open the **Physical Air-Gap Setting** dialog. The dialog shows the current state and describes the impact of the change.
4. Type **YES** in the confirmation field.
5. Click **Enable Physical Air-Gap** or **Disable Physical Air-Gap** to apply the change, or click **Cancel** to keep the current setting.

> **Note:** The Physical Air-Gap setting is configured at the vault cluster level. It is not configurable per pairing and applies to all paired clusters simultaneously.

> **Note:** For Physical Air-Gap to take effect, the interface group used for vaulting must be dedicated to vaulting. Do not use the same interface group for other operations, such as ongoing replication. If the vaulting interface group is also used for replication, the interface group remains enabled at all times, and Physical Air-Gap cannot disable it outside the vaulting window.

For more information, see Pair Primary Cluster to Vault cluster.

Pair Primary Cluster to Vault cluster

Pair the primary cluster with vault cluster to enable the connection between the clusters.

Primary cluster can be a brand new cluster or an existing cluster you use for backup and restore operations. For more information on initial cluster setup, see Initial Cluster Setup.

Connect Vault Cluster to the Primary Cluster

To connect the vault cluster to primary cluster:

1. On the vault cluster, navigate to **Vault** > **Paired Clusters** and click **Pair Cluster**.

   > **Note:** Only users with the FortKnox Self Managed Vault Admin role can perform cluster pairing and update the existing pairing configurations including those previously created by Admin role users. For more details on vault admin role, See Manage Roles.

2. Enter all the Virtual IPs or Node IP Addresses of the cluster that you need to pair with the vault cluster in the **VIP or Node IP Addresses** column.
3. Enter the Username and Password of the user with the FortKnox Self Managed Vault Admin role on the primary cluster.
4. Select the required Interface Group **Automatic** or **Manual**.

   > **Note:** If you select Manual, you can choose from the list of Interface Groups.

5. Click **Continue: Storage Domain Pairing**.

Pair Storage Domain

To pair the primary and vault cluster storage domains:

1. Select the source storage domain on the **Source** column. Source is the storage domain on the primary cluster that is paired with the destination storage domain on the vault cluster.
2. Select the destination storage domain on the **Destination** column. Destination is the storage domain on the vault cluster.

   > **Note:** Vault cluster is the destination cluster.

3. Select the **Vaulting window**. The vaulting window is the time you choose to vault the data from the primary cluster to the vault cluster.

> **Note:** Cohesity recommends configuring the vaulting window only after the initial full and a few incremental vaultings are successfully completed. Ensure that the vaulting window is appropriately sized to allow vaulting to complete.

4. Click **Add Domain Pair**. You can add multiple domain pairs.

> **Note:** This pairing connects the storage domains on the primary cluster and vault cluster that includes the vaulting window configuration. Data replication to the vault cluster occurs only during the defined vaulting window, when communication between the clusters is enabled.

5. **Optional**: Toggle on the **Data Transfer Throttle** if required.

   In **Throttle to** column, enter the throttling speed in Megabits per second (Mbps).

6. Click **Continue: Summary** to view the cluster pairing summary.
7. Click **Complete Pairing** to pair the cluster to vault cluster.

Download Paired Clusters Report

You can download the report that lists all the clusters that are paired to the vault cluster, including pairing information, vaulting window, and protection group.

To download paired clusters report:

1. On the vault cluster, navigate to **Vault** > **Paired Clusters**.
2. Click **Download CSV** to download the paired clusters report.

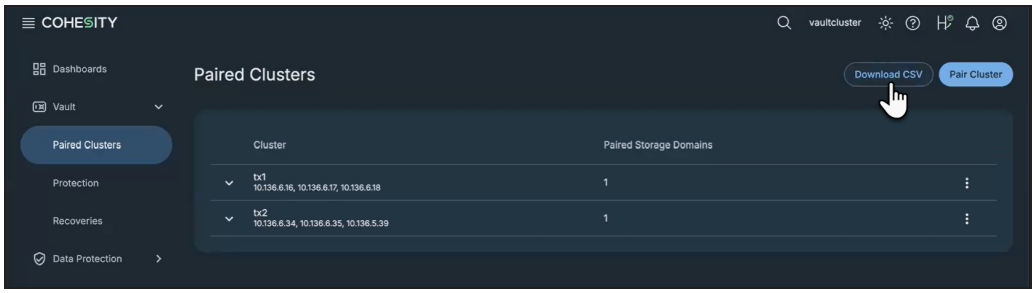

Refresh API Keys

The Refresh API Key option allows the vault administrators to re-establish communication between paired clusters when API keys expire. It addresses scenarios where a primary cluster is reinstalled without first unpairing from the vault cluster. In such cases, the API keys associated with the original pairing can expire, resulting in communication failures when the new primary cluster attempts to connect.

To refresh the API Keys:

1. On the vault cluster, navigate to **Vault** > **Paired Clusters**.

   > **Important:** Log in to the vault cluster with the Vault Admin role to perform the refresh API Keys operation.

2. Hover over the paired cluster that you want to refresh the API Key, click the action icon (⋮), and select **Refresh API Key**.
3. In the **Refresh API Key** dialog, enter the Vault Admin remote cluster credentials.
4. Click **Refresh**.

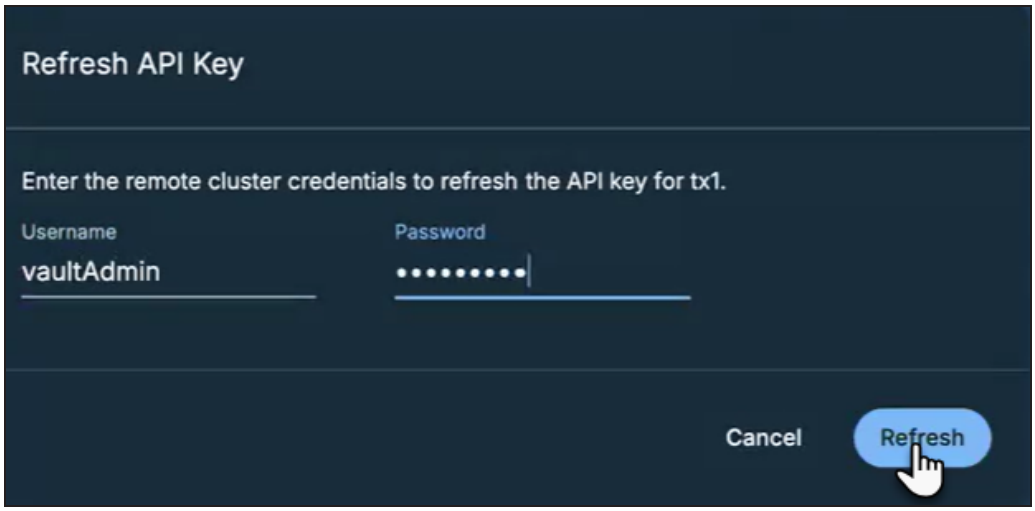

Enter Multiple Node IPs During Cluster Pairing

Entering node IP addresses during cluster pairing now supports comma-separated input. You can type or paste multiple IP addresses in a single entry into the VIP or Node IP Addresses field. The field resolves all entered addresses automatically without requiring any additional steps between entries.

Previously, each IP address had to be entered individually, which was time-consuming for clusters with multiple nodes.

> **Tip:** You can copy a comma-separated list of IP addresses from a spreadsheet or text file and paste it directly into the field. For example: 192.168.1.185, 192.168.1.186, 192.168.1.187

Enter Multiple Node IPs

To pair a cluster using multiple node IPs:

1. On the vault cluster, navigate to **Vault** > **Paired Clusters** and click **Pair Cluster**.
2. In the **VIP** or **Node IP Addresses** field, type or paste IP addresses separated by commas. The field resolves each address automatically when focus leaves the field.
3. Verify that all IP addresses are displayed correctly before proceeding.
4. Continue with the remaining cluster pairing steps.

For more information, see Connect Vault Cluster to the Primary Cluster.

Prevent Duplicate Cluster Pairing

FortKnox Self-Managed now validates cluster registrations during pairing and blocks any attempt to pair a cluster that is already registered as a remote or vault-paired cluster. A clear error message is displayed when a duplicate pairing is attempted, and the operation does not proceed.

This validation prevents accidental duplicate registrations in environments where the same cluster may be visible under multiple roles.

Error Behavior

- The pairing operation is blocked before any changes are made.
- An error message is displayed indicating that the cluster has already been registered.
- To pair the cluster in a new role, remove the existing pairing or registration first, then initiate the new pairing.

> **Note:** This validation applies to all pairing attempts, including attempts to register the same cluster as both a remote cluster and a vault-paired cluster.

Create Protection Policy for Vault Cluster from Primary Cluster

To create a protection policy for the vault cluster from the primary cluster:

1. On the primary cluster, navigate to **Data Protection** > **Policies**.
2. Click **Create Policy** located at the top right of the page.
3. **Policy Name**: Specify a protection policy name. The name can contain the alphanumerics, underscores, hyphens, periods, and spaces. This field is required and can be changed later.
4. Specify when backups are captured:
    1. **Backup every xx**: Specify the minutes, hours, weeks, months, or years to set the frequency for how often backups are captured by the Protection Group. This defines the protection schedule for capturing backups.

       The policy schedule can be configured to run on specific days. From the **Every** drop-down, select **Week** and in the **On** field, select the days of the week you want to run the backup. For example, you can schedule a backup to run on Monday, Wednesday, and Friday, and not on Tuesday and Thursday.

       The policy schedule can be configured to run on specific dates as well (calendar-based scheduling). From the **Every** drop-down, select **Month**, and in the **On** field, select the **Date**. You can then choose the date of the month you want to run the backup. For example, you can schedule a backup to run on the 5th of every month.
5. Under **Primary Copy**, specify the following:
    1. **Keep On**: Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster.

       If you are selecting any external targets, then this policy can be only used for performing CloudArchive Direct wherein the primary copy of the data is stored on the external targets. For more information, see CloudArchive Direct.
    2. **Retain for**: Specify how many days, weeks, months, or years backups are retained on the Cohesity cluster before they are deleted.
        1. The **Retain** option can be set to a maximum value of 365000.
        2. Retentions at 7-day boundaries turn into weeks.
    3. **Lock**: Specify how many days, weeks, months, or years the retained log backups should be locked. You cannot delete the snapshots during the specified period.

       > **Note:** This option is available only if you have enabled the DataLock option at the policy level.
6. Click **More Options** to configure the additional policy settings.

7. **DataLock**: Toggle on to lock the policy and any snapshots that belong to Protection Groups associated with the policy. DataLock is typically used for compliance and regulatory purposes. For more information, see DataLock.

   If you have enabled the **DataLock** option at the policy level, then you have the option to lock the backup retained on the Cohesity cluster. In the **Lock** option field available in the **Primary Copy** pane, specify how many days, weeks, months, or years the retained backups must be locked. You cannot delete the snapshots during the specified lock period.
8. Select **Add Vault** to add the vault cluster for the backups to the protection policy. Specify the following:
   1. **Transfer to**: Select the Vault cluster from the dropdown menu to transfer the data from the source cluster.
   2. **Every**: Specify the frequency of data transfer to the vault.
   3. **Retain for**: Specify the duration of the vaulted copy to be retained on the vault cluster.

      The default minimum retention is set to the primary copy retention value. The vault retention value can be greater than or equal to the primary copy retention.

      > **Note:** Data transferred from the source cluster to the vault cluster is automatically locked using DataLock. This lock is permanently applied to vaulted snapshots, cannot be disabled, and remains active throughout the retention period. DataLock ensures immutability and compliance by preventing any modification or deletion of data during its lifecycle in the vault cluster.
9. **Optional**: Enable **Transfer only successful runs** if required.

   Transfer only successful runs: With this option, the data from the source cluster is transferred to the vault cluster only after the successful protection run.
10. If required, you can add additional policy settings. For more information, see Create or Edit a Standard Policy.

    > **Note:** Only the Data Security user can edit the vaulting protection policy.
11. Click **Create**.

Flexible Vault Retention Configuration

The vault retention period in a protection policy can now be set to any value, including a value lower than the primary backup retention. Previously, the system enforced a minimum retention equal to the primary copy retention and prevented saving the policy if the vault retention was set lower.

This change supports use cases where a shorter retention period is appropriate for vaulted copies, such as cost optimization or regulatory requirements that call for a reduced vault retention independent of the primary backup timeline.

> **Note:** Removing the minimum retention restriction does not affect DataLock. Data transferred to the vault cluster is automatically locked and cannot be modified or deleted during the retention period.

Set Vault Retention in a Protection Policy

To configure vault retention during policy creation:

1. On the primary cluster, navigate to **Data Protection** > **Policies** and click **Create Policy**.
2. Click **Add Vault** to add the vault cluster to the policy.
3. In the **Retain for** field under the vault configuration, specify the desired retention period.

   > **Note:** You can now enter a value lower than the primary copy retention.

4. Complete the remaining policy settings and click **Create**.

Exclude Paired Clusters from Replication Targets

When creating a protection policy on the vault cluster, clusters that are already paired to the vault cluster as primary clusters are no longer displayed in the replication target dropdown. Only clusters registered as independent remote clusters appear in the list.

Previously, vault-paired primary clusters appeared in the same replication target list as independent remote clusters, which could lead to unintended selection and misconfiguration.

To view the paired clusters that are excluded from replication targets:

1. On the vault cluster, navigate to **Data Protection** > **Policies** and click **Create Policy**.
2. Under the replication configuration, click **Replicate to** drop down and select **Remote Cluster**.

The replication target dropdown now shows only independent remote clusters. Vault-paired primary clusters are automatically excluded and do not appear in the list.

> **Note:** No configuration change is required. Vault-paired primary clusters are excluded from the replication target list automatically.

Register a Source on Primary Cluster

Registering a source makes it eligible to be protected or backed up on the Cohesity cluster.

To register a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** > select any source.

For more information on registering a source, see Register a Source.

Create a Protection Group on Primary Cluster

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group.

To create a protection group,

1. On the primary cluster, navigate to **Data Protection** > **Protection** and do one of the following:
    1. To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
    2. To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. If creating a new Protection Group, select **Protect** >select a source from the list.
3. In the **New Protection** pop-up, do one of the following:
    1. Protect the source by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
        1. **Add Objects**: Select the registered source or click **Register Source** to register a source. From the **Objects** list, select the volumes you want to back up.
        2. **Protection Group**: Specify a name for the Protection Group.

3. **Policy**: Select an existing vault protection policy or click **Create Policy** to create a new protection policy. For more information, see Create Protection Policy for Vault Cluster from Primary Cluster.
4. **Storage Domain**: Select the appropriate storage domain.

   > **Note:** This option is visible on the New Protection pop-up only if there is more than one storage domain.

   > **Important:** Select the storage domain used during pairing to ensure vaulting runs successfully.

5. Click **Protect**.

2. To protect the source by specifying all the options, click More Options and select the required settings.

   The **Vault Enabled** tag displayed next to the Protection Group name indicates that vaulting is enabled for the Protection Group, clearly distinguishing it from local Protection Groups.

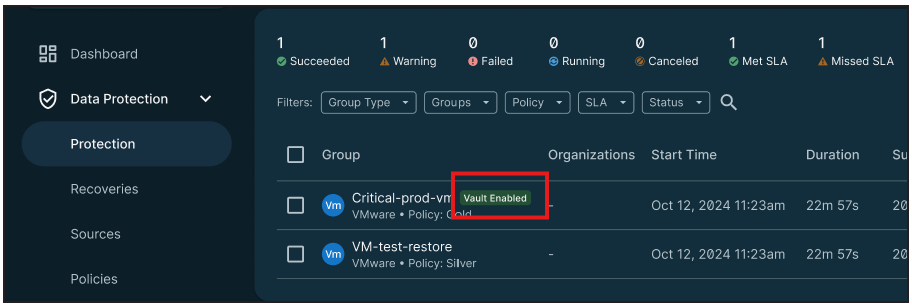

For more information on creating a protection group, see Create a Protection Group.

Edit a Protection Group Run to Vault Specific Objects

You can vault only specific objects from an existing vault-enabled protection group during an edit run. It allows you to re-vault only selected objects within a protection group, rather than vaulting the entire group.

Example use case

If a protection group is configured to vault three virtual machines and all three were successfully vaulted in an earlier run, you can now edit the run and choose to vault:

- A single virtual machine
- A subset of the virtual machines
- All virtual machines again, if required

Only the selected objects are included in the edit run and re-vaulted.

To vault specific objects using **Edit Run**:

1. On the primary cluster, navigate to **Data Protection** > **Protection**.
2. Select the vault-enabled protection group.
3. Hover over the protection run that you want to modify, click the action icon (⋮), and select **Edit Run**.

   > **Note:** You can click the protection run to view the backup snapshots, and then select the **Vault** tab to verify the objects vaulted in the earlier run.

4. In the **Edit Run** dialog, click **Vault**.
5. Click the **Vault** dropdown and select one of the below options:
   a. Select **All objects in the protection group** to vault all objects in the protection group.
   b. Select **Selected objects in the protection run** to view the objects originally included in the protection group.
      i. Select the objects that you want to vault in this edit run.
      ii. Click **Continue**.

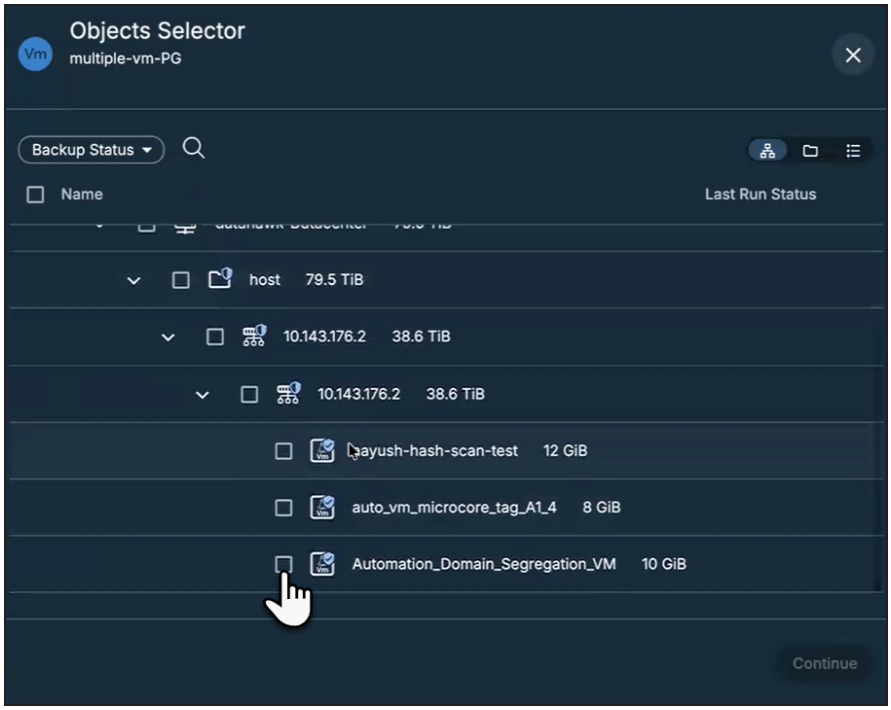

> **Note:** You can select one, multiple, or all available objects. Only the selected objects are included in the new vaulting operation.

Click **Run**.

Cohesity now updates the protection run and vaults only the selected objects. Objects that were vaulted previously but are not selected are skipped.

Prevent Duplicate Targets in Edit Run

When editing a protection run, replication and vault targets that are already associated with that run are now displayed as disabled in the target selection list. This prevents the same target from being added more than once to a single run.

Behavior by Cluster Context

Cluster Context	Available Target Types in Edit Run
Primary cluster	Replication, Archival, and Vault. Targets already selected for the run are grayed out and cannot be selected again.
Vault cluster	Replication and Archival only. Vault is a primary cluster workflow and is not available in the vault cluster context. Already-selected targets are grayed out.

Edit a Run

To edit a protection run and add a new target:

1. Navigate to **Data Protection** > **Protection** and select the vault-enabled protection group.
2. Hover over the protection run, click the action icon, and select **Edit Run**.
3. In the **Edit Run** dialog, click the + icon to add a target. Targets already associated with the run appear grayed out in the list.
4. Select an available target and click **Run** to apply the change.

> **Note:** This behavior applies to both replication targets and vault targets. No configuration change is required.

For more information, see Edit a Protection Group Run to Vault Specific Objects.

Vaulting Status

You can view the vaulting status once the protection group is created from the primary cluster. The status can be viewed from the primary cluster and the vault cluster.

> **Note:** The vaulting status will appear in the vault cluster only after a successful protection run.

View Vaulting status from the Primary Cluster

To view the vaulting status from the primary cluster:

1. On the primary cluster, navigate to **Data Protection** > **Protection**.
2. Select the protection enabled source.
3. Click the **Vault** tab to see the vaulting status.

The **Paused** status means the protection run could not finish within the set vaulting window. The run stays visible in the protection group and will automatically resume when the window opens again.

View Vaulting status from the Vault Cluster

To view the vaulting status from the **Data Protection** tab:

1. On the vault cluster, navigate to **Data Protection** > **Protection**.
2. Select the successful protection run.
3. Click the **Vault** tab to see the protection run details.

Use the **Status** filter to view the specific protection status.

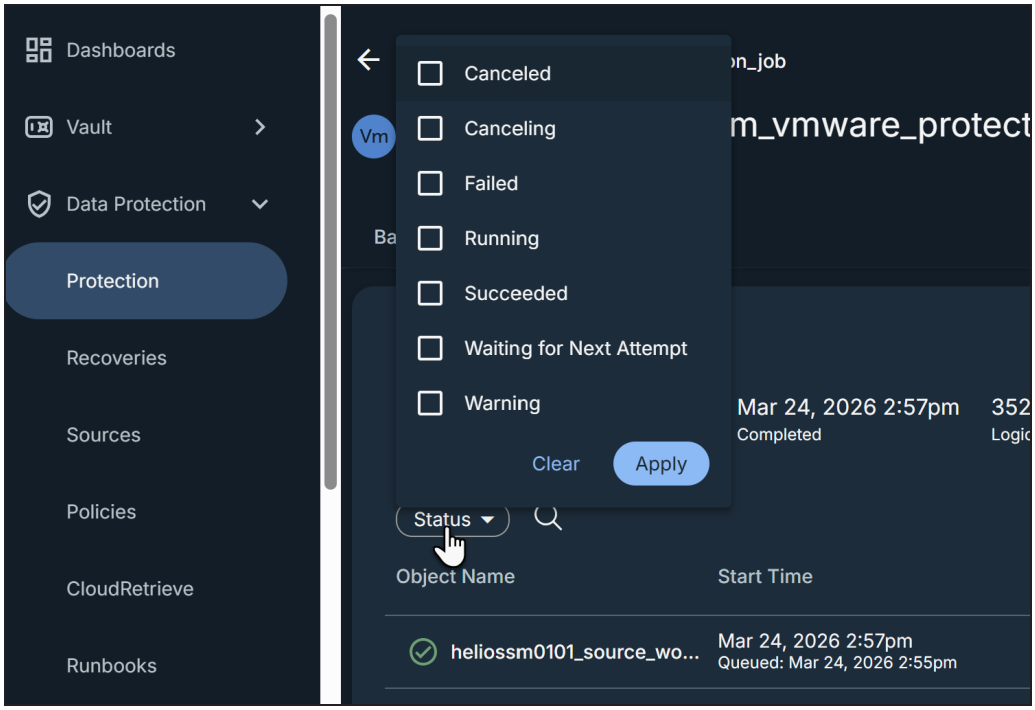

To view the vaulting status from the **Vault** tab:

1. On the vault cluster, navigate to **Vault** > **Protection**.

   Use the **Status** filter to view the specific protection status.

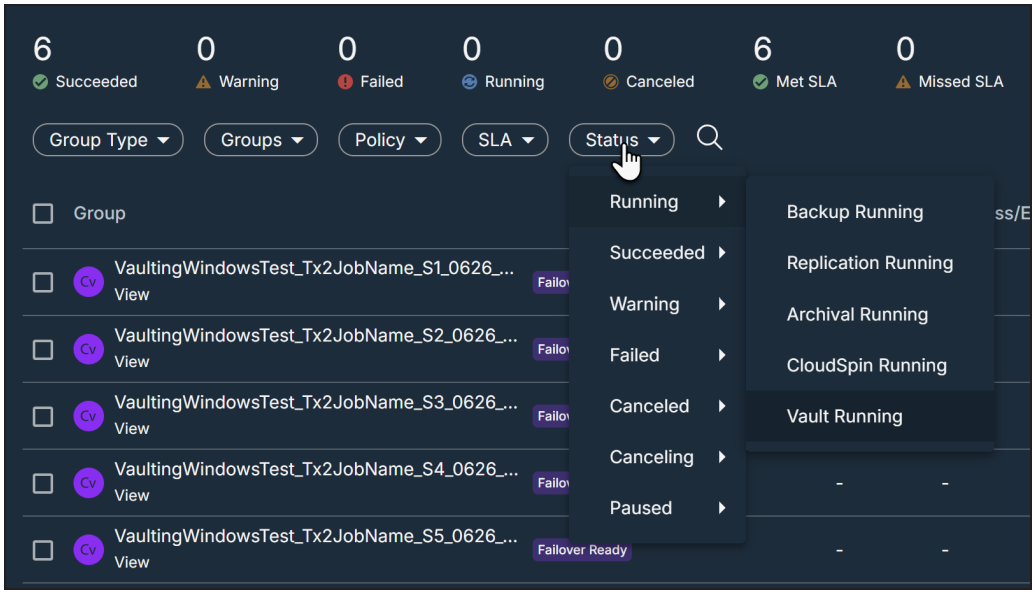

2. Select a vault-enabled protection group and select a successful protection run.
3. Click the **Vault** tab to see the protection run details.

Recover from Vault Cluster

You can recover the sources that are vaulted to the vault cluster to the primary cluster or a newly specified location.

Refer to Considerations section before you initiate a recovery.

To recover the vaulted sources from the vault cluster:

1. On the vault cluster, navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** located at the top right of the page.

For more information on the recovery, see Recovery.

To recover the files or folders of a source, refer to the recovery section of the specific source.

To manage a recovery, see Manage Recovery.

Supported Vaulting and Recovery Configurations

FortKnox Self-Managed supports multiple configurations for both vaulting and recovery workflows to enable deployment flexibility.

Configuration 1

The primary cluster vaults data to a designated vault cluster. All recovery operations are performed directly from the vault cluster.

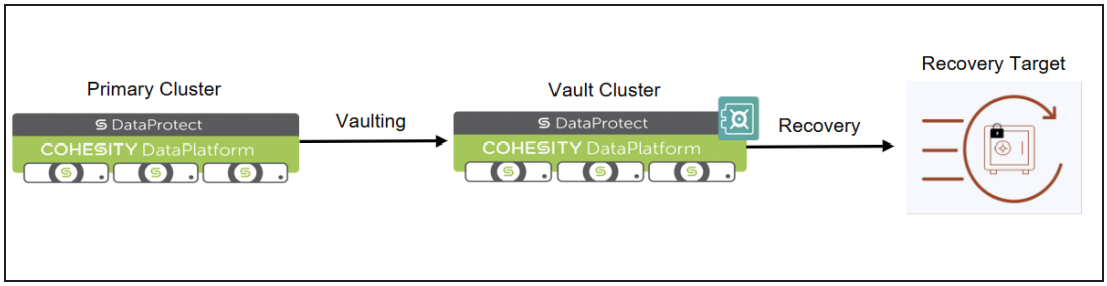

To vault the backup data directly from the primary cluster to the vault cluster:

1. From the vault cluster, initiate pairing with the primary cluster.
2. On the primary cluster,
    1. Create a vaulting policy. For more information, see Vaulting Policy.
    2. Configure a protection job and assign the vaulting policy to it. For more information, see Protection Group.

Configuration 2

The primary cluster vaults data to a designated vault cluster. The vault cluster then performs policy-based replication to a secondary Cohesity cluster, which is used for

recovery operations. This workflow is supported through APIs. Data protection policies configured on the primary cluster control both vaulting and replication.

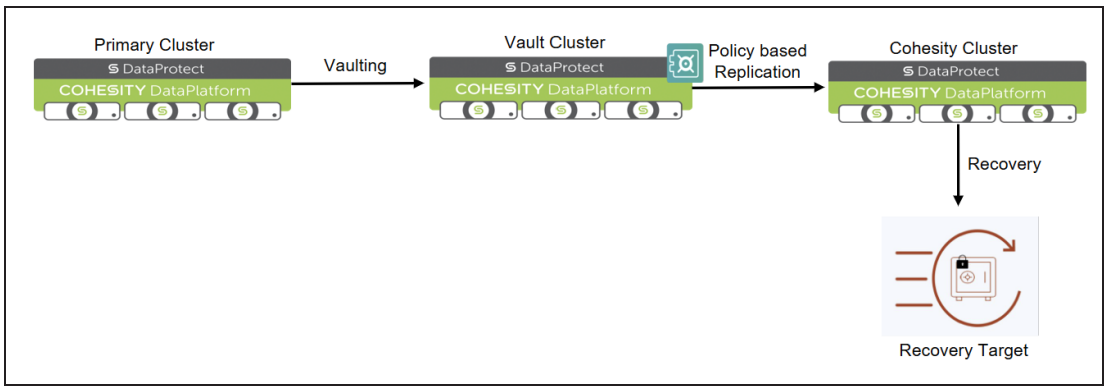

To vault and recover the backup data:

1. From the vault cluster, initiate remote cluster pairing with the secondary Cohesity cluster.
2. On the primary cluster,
    1. Use the API to configure a cascading policy that vaults the data to the vault cluster and then replicates from the vault cluster to the secondary Cohesity cluster.
    2. Provide the cluster ID and name for both the vault cluster and the secondary Cohesity cluster in the cascading policy.
    3. Create and assign a protection job that uses this cascading policy.

    For more information, see Cascaded Replication, Archival, and Vaulting Policy.

Configuration 3

The Primary cluster vaults data to a designated vault cluster. When recovery is needed, snapshots are manually replicated from the vault cluster to a secondary Cohesity cluster. This secondary cluster is used to perform recovery operations.

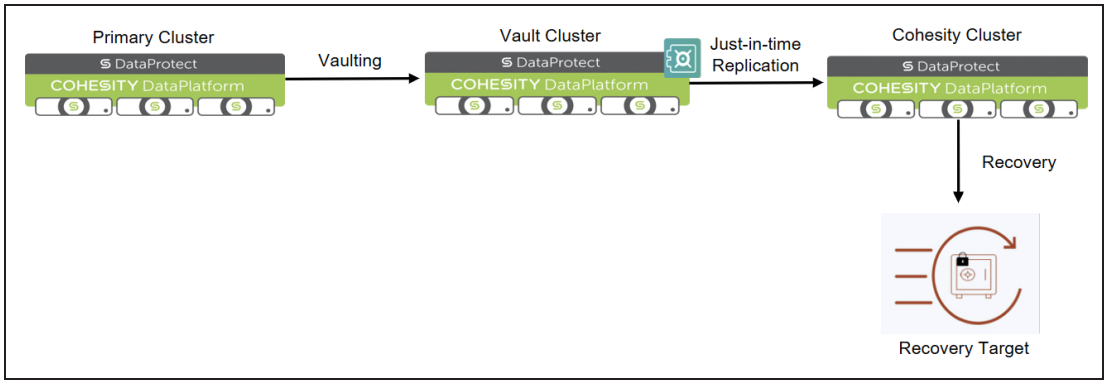

To perform just-in-time replication of snapshots to secondary Cohesity cluster:

1. On the vault cluster, navigate to **Vault** > **Protection**.
2. Click the action menu (⋮) next to the successful protection run and select **Edit Run**
3. On the **Edit Run** dialog, click (+) icon and select **Replication**.
4. Select the intermediate recovery cluster in **Replication Target** drop-down as a new replication target for the specific snapshot and click **Run**.
5. On the intermediate recovery cluster,
    1. Navigate to **Data Protection** > **Protection**.
    2. Verify that the replication is successful.
6. Initiate recovery from the intermediate recovery cluster to the required recovery target.

    You can also add a remote cluster as a replication target after registering the remote cluster in both primary and vault cluster. For more information on registering a remote cluster, see Create a Connection to the Remote Cluster.

Configuration 4

The primary cluster replicates data to a replica Cohesity cluster using policy-based replication. The replica cluster then vaults the replicated data to the vault cluster. This workflow is supported through APIs. Policies configured on the primary cluster control both replication and vaulting operations.

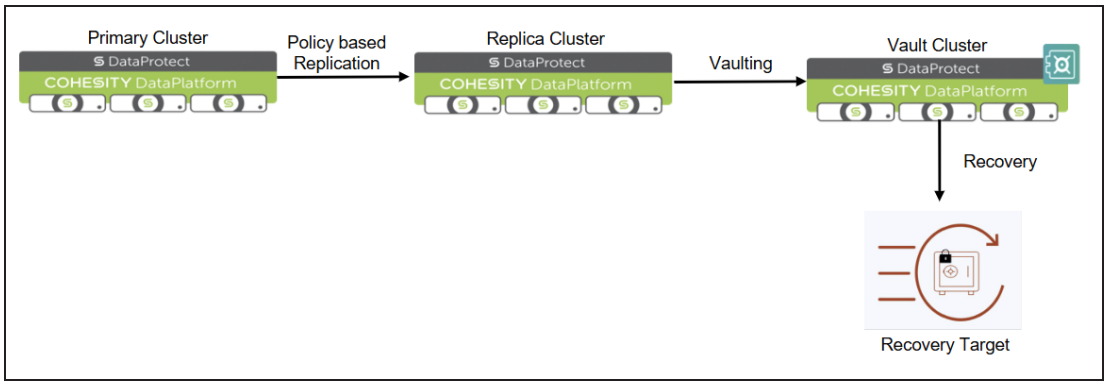

To replicate the backup data and recover:

1. From the vault cluster, initiate vault pairing with the Replica Cohesity cluster.
2. On the primary cluster,
    1. Use the API to configure a cascading policy that replicates the data to the replica Cohesity cluster and then vaults from the replica cluster to the vault cluster.
    2. Provide the cluster ID and name for both the replica cluster and the vault cluster in the cascading policy.
    3. Create and assign a protection job that uses this cascading policy.

Limitation for Cascaded Policy

The Primary cluster UI does not display the vault cluster.

For more information, see Cascaded Replication, Archival, and Vaulting Policy.

Configuration 5

FortKnox Self-Managed supports multiple primary clusters vaulting to one vault cluster.

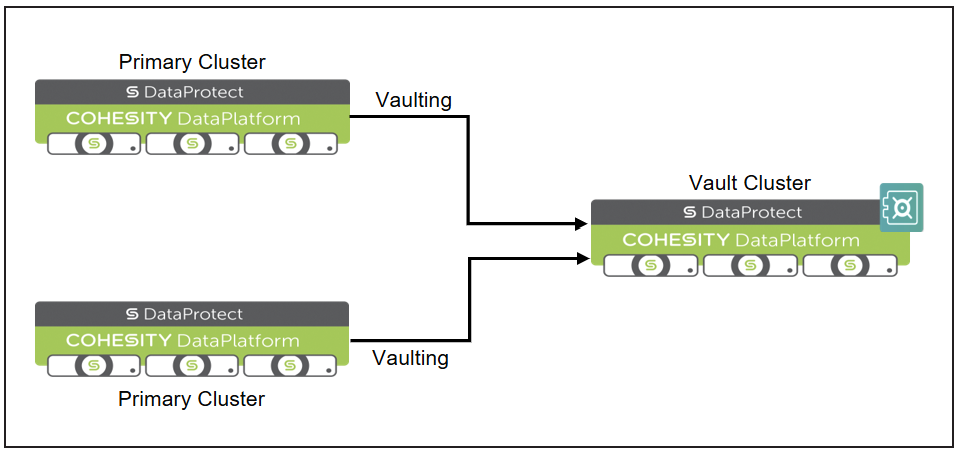

For more information, see Configuration 1.

Configuration 6

FortKnox Self-Managed supports one primary cluster vaulting to multiple vault clusters using different policies.

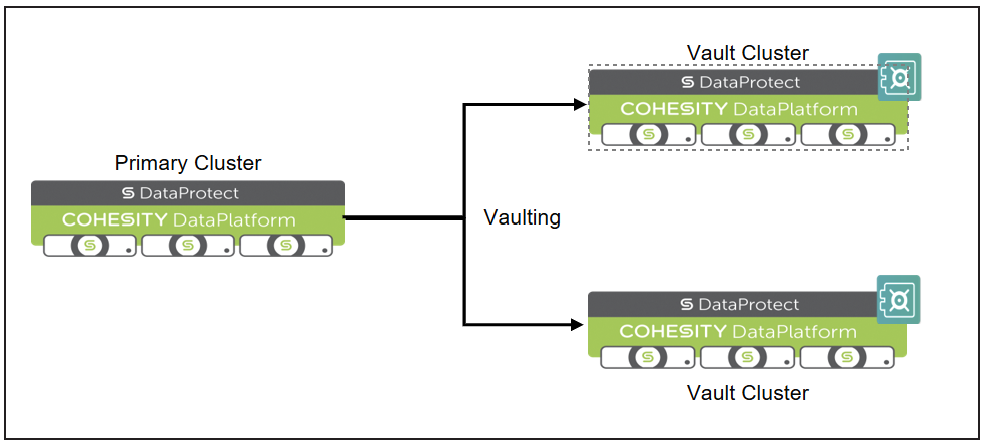

For more information, see Configuration 1.

Clean Room Recovery using FortKnox Self-Managed

FortKnox Self-Managed enables you to vault a copy of your backup data in a securely isolated Cohesity cluster within your own infrastructure. The Clean Room Recovery workflow for FortKnox Self-Managed provides a secure mechanism to recover data

Snapshots and the Digital Jump Bag from the Vault Cohesity to an Isolated Clean Room environment using Just-In-Time replication, governed by Quorum approval. The vault cluster does not accept any inbound connection from the Clean Room environment during this process.

Digital Jump Bag

A Digital Jump Bag is a secure software repository that stores the critical assets you need to rebuild your infrastructure during a cyberattack, including OS images, software packages, configuration files, and documentation templates for dial-tone workloads such as Active Directory, DNS, security tools, contact lists, and recovery procedures. With Cohesity, you create a Digital Jump Bag using a SmartFiles View with data lock and object lock enabled and secure it using policy-based replication and FortKnox vaulting so that a trusted, immutable copy is always available when you need it.

Goals

- Securely recover the Digital Jump Bag and data Snapshots from the FortKnox Vault Cluster to the Isolated Clean Room environment.
- Maintain WORM and immutability SLAs for all backup data stored in the FortKnox Vault Cluster throughout the recovery process.
- Ensure no inbound communication is established into the FortKnox Vault Cluster during the CyberAttack Recovery process.

Non-Goals

- Setting up the Clean Room environment, network isolation, or micro-segmentation configuration.
- Performing forensic analysis on the data recovered from the vault cluster to the Clean Room.
- Recovering sanitized data from the Clean Room to an isolated Sandbox environment after forensic analysis is complete.
- Restoring data from the Sandbox environment to the Production environment, or reconfiguring primary backup, replication, and vaulting.

Assumptions

- The Digital Jump Bag, which contains essential OS images, configuration templates, and installers, is stored in a SmartFiles View on the Primary Cohesity Cluster and vaulted to the FortKnox Self-Managed Vault Cluster.

- All Primary Cohesity clusters are inaccessible. The FortKnox Vault Cluster is the sole source for data Snapshots and the Digital Jump Bag during Clean Room Recovery.
- Your Active Directory servers and all online identity providers are compromised and cannot be trusted as part of the attack.

> **Note:** The Intermediate Cluster must reside in a dedicated, isolated network segment (for example, a VLAN with firewall rules) during the replication phase. This prevents lateral movement and preserves the integrity of the FortKnox Vault Cluster.

Prerequisites

Ensure the following are configured before a cyber-recovery event occurs.

All prerequisite tasks must be completed before a cyber-attack occurs. These tasks cannot be performed after primary clusters become inaccessible.

- Set Up a Quorum Group in Helios Self-Managed
- Configure Vault Pairing
- Set Up the Digital Jump Bag
- Configure Vaulting to the FortKnox Vault Cluster

Set Up a Quorum Group in Helios Self-Managed

In Helios Self-Managed, deployed in the isolated Vault network, set up a Quorum Group for the following Cluster Operation:

Update runs for a Protection Group on the FortKnox Vault Cluster.

For more information, see Set Up a Quorum Group.

Configure Vault Pairing

Set up Vault pairing between the Primary Cluster and the FortKnox Vault Cluster. For more information, see Pair Primary Cluster to Vault cluster.

Set Up the Digital Jump Bag

To set up the Digital Jump Bag:

1. Create a SmartFiles View on the Primary Cohesity Cluster.
2. Copy all Digital Jump Bag contents (OS images, configuration templates, and installers) into the View.
3. Configure a local backup policy for the View.

Configure Vaulting to the FortKnox Vault Cluster

To configure vaulting:

1. Create a protection policy that targets the FortKnox Vault Cluster.
2. Create Protection Groups for both the data Snapshots and the Digital Jump Bag SmartFiles View, and assign the vaulting policy to each.

For more information, see Create a Protection Group on Primary Cluster.

Recovering to the Clean Room

Complete the following steps in order during a cyber-attack recovery event:

- Close the Vaulting Window on the FortKnox Vault Cluster
- Pair the Intermediate Cluster to the FortKnox Vault Cluster
- Trigger Just-In-Time Replication
- Un-pair the Intermediate Cluster and Recover Data Snapshots

Close the Vaulting Window on the FortKnox Vault Cluster

Close the Vaulting Window to stop all active and scheduled vaulting from Primary Clusters to the Vault Cluster. This prevents any compromised cluster from writing to the Vault Cluster during the recovery process.

To close the Vaulting Window:

1. On the FortKnox Vault Cluster, navigate to **Vault** > **Paired Clusters**.
2. Click the action icon next to the paired Primary Cluster and select **Edit Pairing**.
3. Set the Vaulting Window to Closed and click **Save**.
4. Repeat for all paired Primary Clusters.

> **Note:** Closing the Vaulting Window drops active vaulting connections and pauses in-progress jobs. Any paused jobs will not resume because the window remains closed.

Pair the Intermediate Cluster to the FortKnox Vault Cluster

Pair a new Cohesity cluster to the FortKnox Vault Cluster. This Intermediate Cluster serves as the staging target for all data and is used for recovery into the Clean Room environment.

To pair the Intermediate Cluster:

1. On the FortKnox Vault Cluster, navigate to **Vault** > **Paired Clusters** and click **Pair Cluster**.

2. Enter the Virtual IPs or Node IP addresses of the Intermediate Cluster in the **VIP** or **Node IP Addresses** field.
3. Enter the username and password of a user with the FortKnox Self-Managed Vault Admin role on the Intermediate Cluster.
4. Select the required Interface Group and click **Continue: Storage Domain Pairing**.
5. Select the source and destination storage domains for the pairing.
6. Click **Complete Pairing**.

> **Note:** Do not configure a Vaulting Window for this pairing. The Intermediate Cluster connection is used for Just-In-Time replication only and should not receive recurring vaulting jobs.

Trigger Just-In-Time Replication

Initiate a one-time, Quorum-gated replication of the required data Snapshots and the Digital Jump Bag from the FortKnox Vault Cluster to the Intermediate Cluster.

> **Note:** Just-In-Time replication is triggered through the API only. This operation is not available from the Cohesity UI.

To trigger Just-In-Time replication:

1. Using the Cohesity API, send an Update Protection Group Run request for the Protection Groups containing the desired data Snapshots, targeted at the Intermediate Cluster.
2. Send a second Update Protection Group Run request for the Digital Jump Bag SmartFiles View Protection Group, targeted at the Intermediate Cluster.
3. A Quorum request is generated automatically for each operation.
4. Notify the designated Quorum Group members to review and approve each request in Helios Self-Managed.
5. After Quorum approval, monitor the replication jobs until they complete.
6. Verify that all required Snapshots and the Digital Jump Bag are present on the Intermediate Cluster.

Un-pair the Intermediate Cluster and Recover Data Snapshots

After replication is complete, un-pair the Intermediate Cluster from the FortKnox Vault Cluster to sever all network connections, connect it to the Clean Room environment, and recover the required data Snapshots to the Recovery Targets.

To un-pair the Intermediate Cluster, bootstrap the Clean Room environment, and recover data Snapshots:

1. On the FortKnox Vault Cluster, navigate to **Vault** > **Paired Clusters**.
2. Click the action icon next to the Intermediate Cluster and select **Unpair**.
3. Confirm the unpairing when prompted.
4. Verify that the Intermediate Cluster no longer appears in the **Paired Clusters** list.

   > **Important:** Do not connect the Intermediate Cluster to the Clean Room environment until the unpairing is confirmed. Connecting before unpairing could expose the Vault Cluster's network path to the Clean Room environment.

5. Connect the Intermediate Cluster to the Clean Room network.
6. Mount the replicated Digital Jump Bag SmartFiles View on the Intermediate Cluster.
7. Deploy the required OS images, apply configuration templates, and install the required software from the Digital Jump Bag to bootstrap the Clean Room environment.
8. Trigger the recovery of the desired data Snapshots from the Intermediate Cluster, specifying the Recovery Targets in the Clean Room environment.
9. Monitor the recovery job until it completes and verify that all data Snapshots are available on the specified Recovery Targets.

> **Note:** After completing recovery, all data Snapshots and the Digital Jump Bag are available in the Clean Room environment, with the FortKnox Vault Cluster remaining fully isolated and its WORM data intact. Proceed with forensic analysis on the recovered data.

Vaulting Scenario for Organizations

FortKnox Self-Managed supports the Organizations (multitenancy) feature. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. You can use a single Cohesity cluster to provide service to multiple tenants.

To configure FortKnox Self-Managed for Organizations:

1. **Enable Organizations**

   Enable the Organizations feature on both the primary and vault clusters. For more information, see Enable Organizations (Multitenancy) on a Cluster.

2. **Create Tenant Storage Domains**

Create one or more storage domains for the tenant on both primary and vault clusters.

> **Note:** Only Service Provider (SP) admin can create the storage domains.

3. **Establish Cluster Pairing**

   On the vault cluster, create a pairing between the primary and vault clusters.

   During the pairing process, map the tenant's storage domain on the primary cluster to the corresponding storage domain on the vault cluster.

4. **Create Vaulting Policy**

   On the primary cluster, create a vaulting policy that references the vault cluster as the target.

   > **Note:** Only Service Provider (SP) admin can create the vaulting policy.

5. **Assign Vaulting Policy to Tenant Organization**

   On the primary cluster, within the Service Provider management interface, assign the newly created vaulting policy to the tenant's Organization.

   > **Note:** Only Service Provider (SP) admin can assign vaulting policy to tenant organization.

6. **Create Protection Group**

   On the Primary cluster, create a protection group and select the storage domain designated for vaulting.

   > **Note:** Only Service Provider (SP) admin or Organization Admin can create protection group.

If the tenant requires changes to the vaulting window schedule, they must request the service provider to update the vault pairing configuration with the new time intervals.

For more information on end-to-end Organizations setup, see Organizations (Multitenancy)

Considerations

Refer to these considerations before deploying FortKnox Self-Managed to ensure proper configuration and compliance.

Vaulting Window

Cohesity limits the number of parallel tasks using a slot allocation algorithm. With the current implementation, vault replication slots are not released if vaulting does not finish before the vaulting window closes. The paused vaulting task holds the slot until the vaulting window opens again. This can lead to slot starvation for new vaulting or regular replication tasks, preventing them from running, and impacting cluster scalability.

As a workaround, Cohesity recommends that the vaulting window duration be long enough to allow vaulting to complete.

Agent-Based Recovery from Vault Cluster

To recover agented data sources from a vault cluster, perform a manual registration.

Follow the procedure documented in Generate and Upload Cluster Certificates to register the source with the vault.

This process enables successful recovery without requiring a Cohesity Agent reinstallation.

Inbound Connection Requirements for Certain Recoveries

Certain recovery operations require the recovery target to establish inbound connections to the vault cluster.

Modify firewall configurations to allow the required inbound traffic.

> **Important:**
> - For vCenter VM recoveries, set the inbound NFS firewall rule to Allow in the vault cluster firewall settings.
> - For MongoDB recoveries, set the inbound Data Protection firewall rule to Allow in the vault cluster firewall settings.

Node Addition and Removal

When a new node added or the existing node removed from the primary cluster, make sure to update the vault pairing configuration.

To update the pairing configuration:

1. On the vault cluster, navigate to **Vault** > **Paired Clusters**.
2. Click the **Actions** menu (⋮) next to the paired cluster and select **Edit**.
3. Add/Remove the Node IP/VIP that was added/removed from the primary cluster.
4. Complete the cluster pairing procedure.

If a node is removed from the vault cluster, the vaulting process will get stuck and will need a restart of services on the primary cluster. Contact Cohesity Support for details.

Best Practices

This page outlines best practices for deploying and managing a Cohesity vault, focusing on physical and logical isolation, access control, firewall configuration, foundational services, and recovery strategies. Refer to the Prerequisites before configuring FortKnox Self-Managed.

Vault Deployment

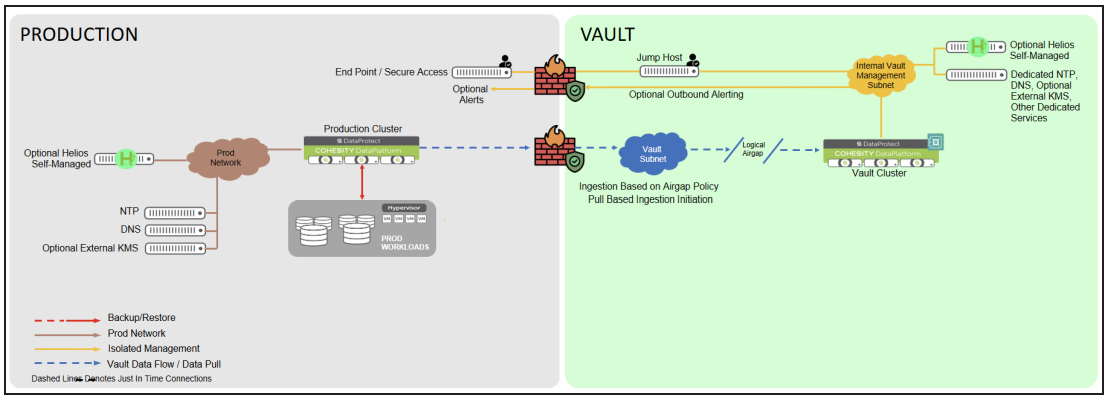

Follow these best practices when deploying the vault.

Physical Isolation

- Deploy the vault in a physically isolated environment, separate from production systems.
- Use dedicated, locked cages or racks to enhance physical security.
- Avoid sharing physical hardware with production infrastructure.

Logical Isolation

- Place the vault in an isolated network protected by a firewall.
- Use separate network interfaces for management and vaulting traffic to prevent cross-contamination.

Vault Access and Management

- Limit access to a minimal number of authorized users.
- Enforce Least Privilege Access and Role-Based Access Control (RBAC).
- Allow management access only via a secure jump host within the isolated network.

- Physically disconnect IPMI ports to prevent out-of-band access.
- Do not register the vault with Cohesity managed Helios.

Vault Firewall Configuration

Block all inbound traffic to the vault network except for controlled access to the jump host.

Allow only the following outbound traffic from the vault:

Destination	Port	Protocol	Usage Notes
Primary cluster	23335	TCP	Required for vaulting scenario.
Primary cluster	23336	TCP	Required for vaulting scenario.

Restrict all other outbound traffic to essential services only.

Foundational Services

Do not share foundational services (e.g., NTP, DNS, AD/LDAP, MFA) between vault, production, and recovery zones.

Data Recovery Options

- While recovery directly from the vault is supported, it is not recommended.
- Replicate data from the vault to a Cohesity cluster in a Trusted Recovery Zone.
- Perform recovery from the Recovery Zone cluster back to production using one of the approved recovery patterns.

Recovery 1

Replicate Vault Snapshots to Recovery Zone, Restore in Recovery Zone, Validate Clean Restore Point, and Restore from Recovery Zone to Trusted Prod Target.

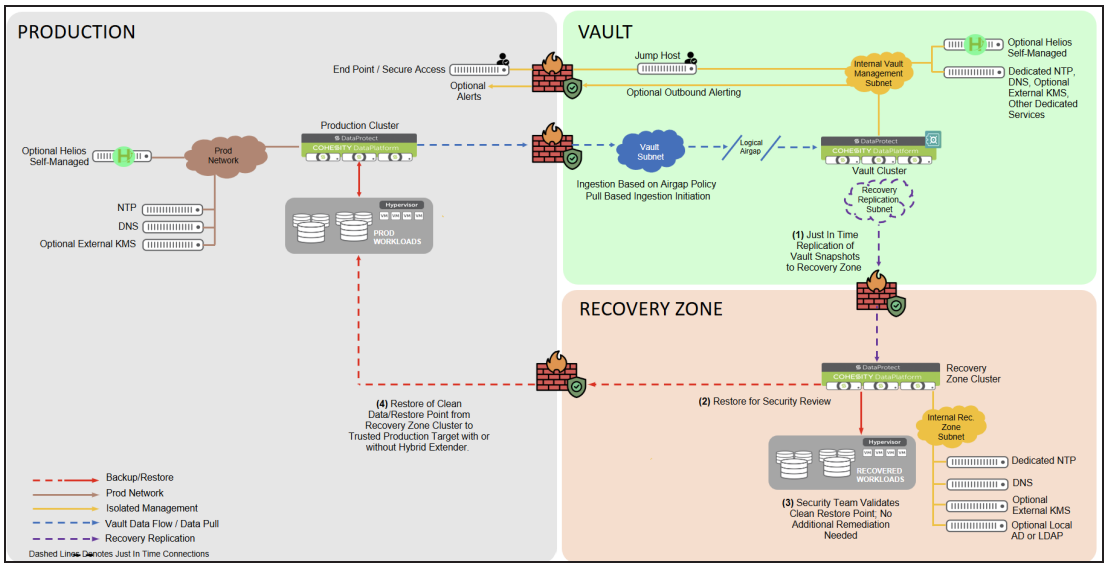

Recovery 2

Replicate Vault Snapshots to Recovery Zone, Restore in Recovery Zone, Validate Clean Restore Point, Backup to Recovery Zone Cluster, and Restore to Trusted Production Target.

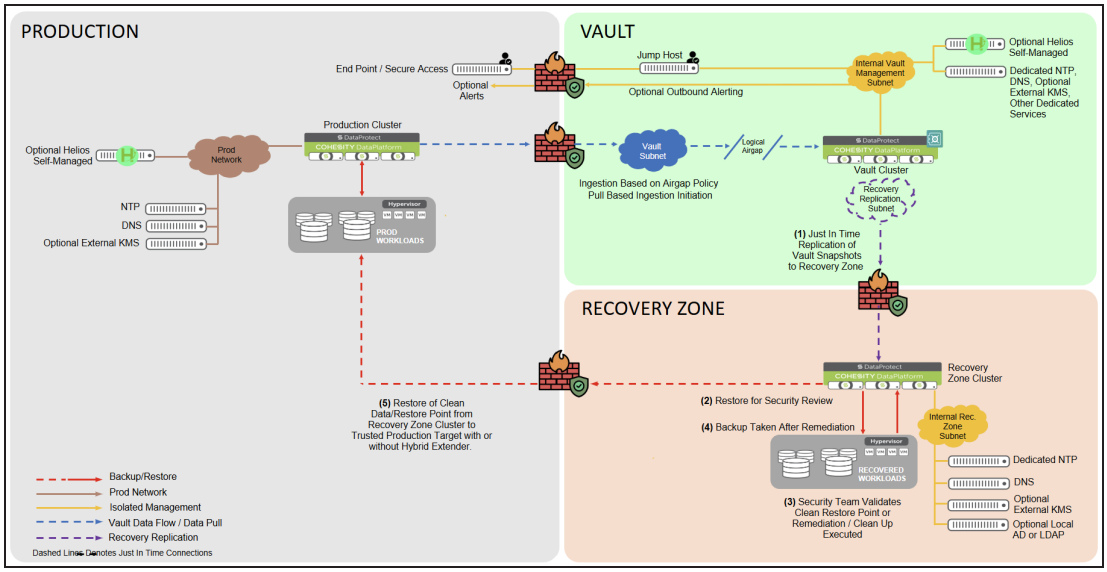

Recovery 3

Replicate Snapshots to Recovery Zone, Validate Clean Restore Point, and Replicate Clean Workload Data to Trusted Target with Workload Tools.

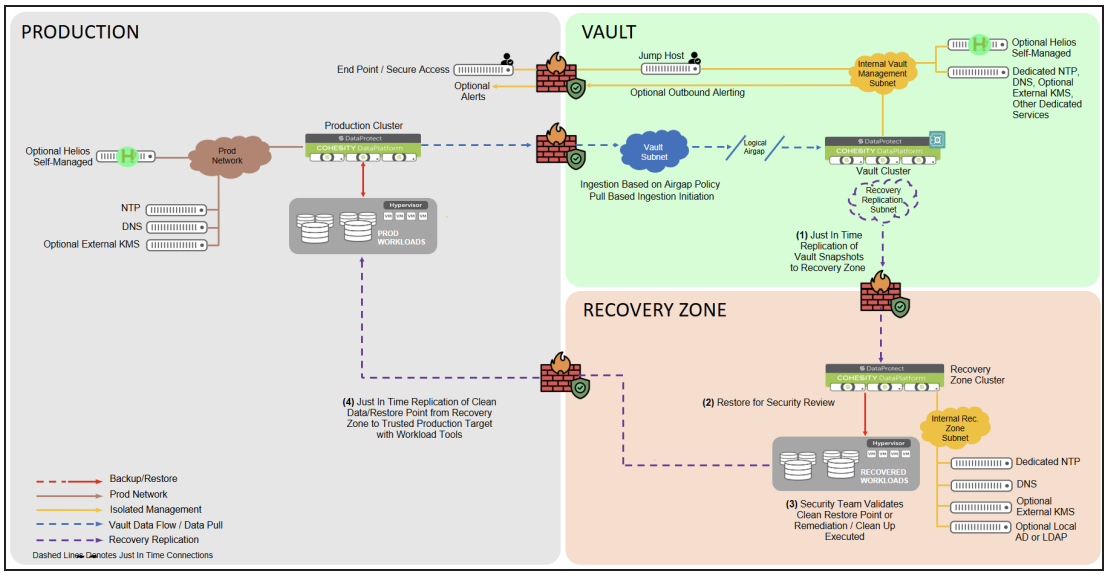

Additional Best Practices

- Use physical clusters for all production vault deployments.
- Avoid using the vault as a Primary backup cluster, even though it is technically supported.
- Implement quorum-based approval via Helios Self-Managed for critical operations (e.g., recovery, firewall changes).
- Use a dedicated Helios Self-Managed instance for vault operations.

Frequently Asked Questions

1. Can previously replicated data still be used for restore?

   Yes. Data replicated before switching to a vault remains fully available for restore operations. However, the cluster cannot be converted to a vault cluster unless storage domain encryption or software encryption at the cluster level is enabled.

2. Does existing replicated data expire?

   Yes. Replicated snapshots will expire according to their original retention policies.

3. Does transitioning to a vault policy or protection group require a new full backup?

   No. Once the replication cluster is converted to a vault target (removed as a replica and added as a vault target, and policies updated accordingly), the next vaulting run will transfer incremental data.

> **Note:** Existing replicated snapshots are immutable if they were data-locked. If they were not data-locked, they are not immutable.

4. Are there version compatibility requirements between the primary backup cluster and FortKnox Self-Managed?

   Yes. The primary backup cluster must be running version 7.2.2_u2 or later to be compatible with FortKnox Self-Managed.

5. Does Proof of Concept (POC) require Multi-Factor Authentication (MFA)?

   Yes. MFA is required. You can use either email-based MFA or an authenticator app for verification.

6. Can data be vaulted from a FortKnox Self-Managed cluster to FortKnox SaaS?

   No. Vaulting from FortKnox Self-Managed to FortKnox SaaS is not supported.

7. Can CloudArchive be used for workloads in a FortKnox Self-Managed cluster?

   No. CloudArchive is not supported for workloads in FortKnox Self-Managed clusters.

8. Can the vault cluster also be used as a replication target?

   No. The vault cluster cannot be used as a replication target.

9. Does Helios Self-Managed require managing the vault cluster?

   No. Helios Self-Managed is optional and needed only for quorum controls. Vault cluster does not support SaaS.

10. What are the requirements when converting a primary cluster to a vault cluster?

    Before converting a primary cluster to a vault cluster,

    1. Enable Multi-Factor Authentication (MFA) for local users.
    2. Encrypt either the entire cluster or all storage domains.
    3. Enable node-to-node encryption across the cluster.
    4. Configure the firewall to deny all traffic except for the “Application - Management” rule.
    5. Remove any active policies for replication or cloud archival.

    After conversion, you can change the firewall settings as they are no longer enforced.

11. Is Cloud Edition/NGCE supported as a FortKnox Self-Managed cluster?

    No. Cloud Edition/NGCE is not supported as a FortKnox Self-Managed cluster. For production deployments, any physical hardware supported for Cohesity DataProtect is also supported for FortKnox Self-Managed. For Proof of Concept (POC) purposes, Cohesity Virtual Edition can be used.

Chapter 6: Cohesity Data Cloud for NetBackup

Cohesity Data Cloud for NetBackup is Cohesity's unified data protection solution that modernizes NetBackup by bringing it onto a single, scalable Cohesity Data Platform. It simplifies operations by replacing fragmented architectures with a streamlined, hyper-converged design that delivers faster recovery, seamless upgrades, and uninterrupted protection.

With centralized management through Cohesity Helios and cyber resilience built in, Cohesity Data Cloud for NetBackup helps organizations protect mission-critical workloads while reducing complexity and operational overhead. The solution is designed to support long-term growth, security, and innovation with a consistent, future-ready data protection experience.

Data Cloud for NetBackup leverages Cohesity cluster to run NetBackup applications. You can manage Data Cloud for NetBackup clusters directly from Helios. For more information on managing Data Cloud for NetBackup clusters directly from Helios, see Manage Data Cloud for NetBackup Clusters from Helios.

Deployment Modes

When you set up a cluster, you choose one of two deployment modes. Your choice determines which services are installed on the cluster nodes and what steps are required during configuration.

- With media servers: Media server services are installed on the cluster nodes. The cluster is registered as a storage server with the NetBackup primary server during setup, and NetBackup assigns backup and restore jobs directly to the media servers running on the cluster. This mode follows the standard Data Cloud for NetBackup configuration workflow.
- Without media servers:Media server services are not installed on the cluster nodes. The cluster operates as a DirectIO storage target and uses external media servers that are registered with the NetBackup primary server. Storage server registration and logical storage unit creation are not performed during initial configuration.

Deployment Overview

Stage 2 - Assemble the nodes

Carefully evaluate all the software and hardware requirements. The requirements cover high level areas such as power and cooling needs, networking infrastructure, and rack sizing, to more specific needs such as system requirements, IP addresses, storage, and security. It is critical that your IT environment meets all the required infrastructure needs to ensure a smoother deployment and operational experience.

At a high level, deploying the cluster involves the following stages:
Stage 1 - Verify the deployment requirements

Assemble the supported hardware and mount the nodes on to a rack in your data center. Connect all the power and network cables as per the instructions provided. After setting up the target systems, procure the installation media and install the softtware on all the nodes.

Stage 3 - Configure the NetBackup network

Configure the NetBackup external primary server and media servers.

> **Note:** This step has to be performed only if you have opted to build the cluster with media servers.

Stage 4 - Configure the cluster

From a web browser, connect to one of the nodes using a public IP address and run the cluster configuration workflow to configure all the nodes into a cluster. During the cluster configuration, you will configure the infrastructure components as well as the core NetBackup services.

Stage 5 - Sign in to the NetBackup UI

Once the configuration is successful, you sign in to the NetBackup web UI and create protection plans and start protecting desired workloads. You can also monitor the infrastructure components and the general health of all the configured services.

Topologies and use cases

Cohesity Data Cloud for NetBackup supports multiple deployment topologies to address different customer requirements such as infrastructure consolidation, phased migration, and storage modernization.

Topology 1: External primary with cluster media servers

In this topology, the NetBackup primary server is deployed externally, while NetBackup media server functionality runs as containerized services on the Cohesity cluster.

Use cases:

- Reduce media server sprawl by adopting a hyperconverged architecture that eliminates external media servers.
- Replace existing media server infrastructure during hardware refresh.

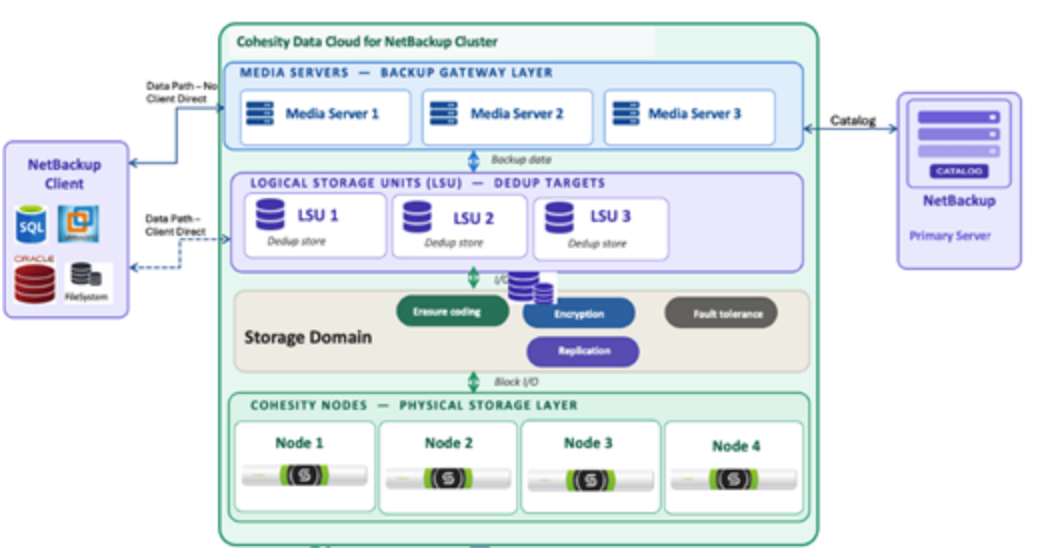

Topology 2: Hybrid deployment

External NetBackup primary server + Cohesity Data Cloud for NetBackup with embedded media servers + External NetBackup media server

In this hybrid model, the NetBackup primary server is deployed externally, while NetBackup media server functionality runs both as containerized services on the Cohesity cluster and on external media servers.

Use cases:

- Support phased transition, where selected workloads are migrated to cluster media servers while existing workloads continue to use external media servers.
- Support environments that require capabilities like Fiber channel initiator workflows, malware scanning, and other related features.

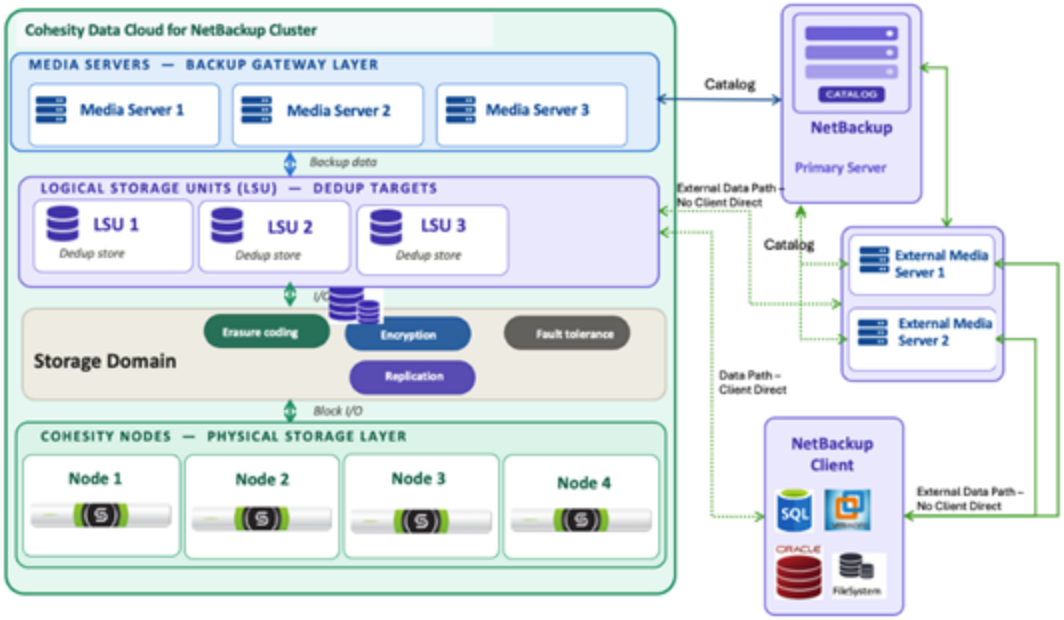

Topology 3: SpanFS as Storage Target

External NetBackup primary server + External media servers + Cohesity Data Cloud for NetBackup cluster

In this topology, the NetBackup primary server and media servers are deployed externally, while Cohesity Data Cloud is used as the storage target using SpanFS.

Use cases:

- Replace or refresh third-party storage with SpanFS without modifying the existing NetBackup environment.
- Support phased transition to a fully integrated, hyperconverged deployment model.

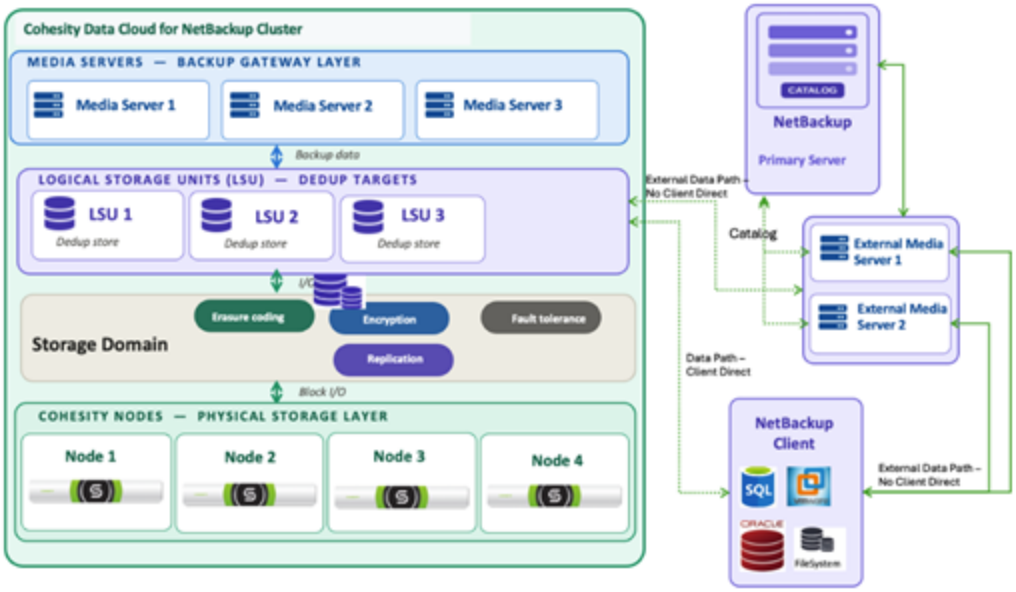

For detailed information about Topology 1 and Topology 2, see this document. For information about Topology 3, see the Cohesity Data Cloud for NetBackup Administrator's Guide.

Plan and Prepare for Data Cloud with NetBackup

Before you configure the cluster, ensure that:

- Cohesity cluster supports the cluster version you plan to configure. For more information, see Required Cohesity Software Version
- You meet the network requirements to configure the cluster. For more information, see Network Requirements
- The firewall ports required for the communication between the cluster and NetBackup are open. For more information, see Manage Firewall Ports
- You meet the prerequisites to configure and manage your cluster.
- Ensure that an external NetBackup Primary Server is available and running. You can use any existing primary server running NetBackup version 11.2 or later.

  This release supports only an external NetBackup Primary Server.

  For more information, see the *NetBackup Installation Guide* on SORT.

Required Cohesity Software Version

You must use Cohesity Data Protect software version 7.4.1 or later.

Network Requirements

You require a minimum of three nodes to configure a cluster. Verify the following cluster network requirements for a three-node cluster.

- Reserve IP addresses for the nodes in the cluster on a single subnet (the cluster subnet).
- Reserve two sets of virtual IP addresses (VIPs). One set should resolve to a single fully qualified domain name (FQDN) for the storage server. The other set should resolve to different FQDNs for media server containers.

  > **Note:** If you plan to deploy without media servers, only one set of VIPs (for the storage server FQDN) is required.

- Open firewall ports to allow the cluster to transmit and receive data and to allow communication between the client, primary server, and data servers on the cluster.
- Verify that the cluster subnet can communicate with the source's subnet.
- Reserve high speed ports (one for each node) on one switch. Minimum of 10GbE required.
- Reserve high speed ports (one for each node) on another switch. Minimum of 10GbE required.
- Reserve IP addresses for the IPMI interfaces. Most hardware platforms require one IP address for each node. All the IPMI IP addresses must be in the same subnet, either in the IPMI subnet or the cluster subnet.
- Reserve 1GbE ports (one for each node) for the IPMI on a switch.
- Enable multicast traffic on your network (required for Cohesity's auto-discovery function).
- The cluster supports bonding mode 1 (active-backup) and bonding mode 4 (LACP 802.3ad). By default, bonding mode 1 is configured on the cluster bond interfaces.

Limitations

Operational Considerations

- To successfully collect logs, initiate log collection using a job ID. When a job ID is used, the system automatically selects a resolvable host name and collects the logs successfully.

  When you initiate log collection without a job ID, you can select a media host and client host manually. For Data Cloud with NetBackup clusters, the cluster ID is

currently displayed in the media host list instead of a resolvable host name. If you select the cluster ID and initiate log collection, media logs are not collected because the cluster ID cannot be resolved to a reachable host and during log collection, the system skips any host that it cannot resolve or reach. Manually adding media hosts or client hosts is not supported in this release.

- Collecting specific debug logs is not supported for Data Cloud with NetBackup media servers.

Set Up the Cluster using the UI

To set up the cluster, you have to perform the following steps:

- Node Selection
- Cluster Network Configuration
- NetBackup Network Configuration
- Cluster Creation

Node Selection

After you verify that all prerequisites are met, you can set up the nodes that will be part of the cluster.

To select the nodes for the cluster:

1. Power on the nodes you want to add to the cluster.
2. After all the cluster nodes are cabled and powered on, connect your setup laptop directly to the cluster node1 using the RJ45 cable/ USB Crash Cart (KVM) Adapter.
3. On the laptop, open the Safari browser and enter the serial number in the browser address bar:

   ```
   http://chassis-SN-node-1.local/
   ```

   Where SN is the serial number of one of the nodes.
4. When the cluster UI login page appears, log in using the default System Admin account called *admin* and the default password, *admin*.

   If Safari indicates *This Connection is Not Private*, click **Show Details**. In the resulting dialog, click the **Visit this website** link and then click Visit Website in the confirmation. If prompted, type the system administrator password for the laptop.
5. On the **What type of cluster would you like to build?** page, choose the **Build Cohesity Data Cloud for NetBackup Cluster** option. Choose one of the following configurations:

- **Build Cluster with Media Servers**
  Installs media server services on the cluster nodes.
  Registers the cluster as a storage server using media servers running on the cluster.
- **Build Cluster without Media Servers**
  Does not install media server services on the cluster.
  Requires registering the cluster later using the NetBackup Primary Web UI.
  Uses external media servers.

6. Ensure that you have all the necessary information and click **Get Started**.
7. The next page shows the number of chassis and nodes that were discovered.
8. In the **Select Nodes** page, select the nodes you want to set up. Select all the available nodes or select the corresponding check boxes. The node that you are connected to is automatically selected. You need a minimum of three nodes to create a cluster. Click **Continue**.

Cluster Network Configuration

To configure the cluster network

1. On the **Configure Cluster Network** page, provide the subnet and IP addresses for the cluster nodes that you selected. Enter the IP and IPMI addresses and credentials for each node to enable hardware-level management. You can use the first node's IP and IPMI addresses to autofill the corresponding fields for the remaining nodes.
2. Under **Cluster External Access & Network Settings**, provide a FQDN for this cluster. Provide one or more virtual IPs to enable high availability access.
3. Under **Apps Network Settings**, provide a private IP subnet in CIDR format for hosting Cohesity Marketplace applications. Ensure that this subnet does not overlap with any existing cluster or client subnet.
4. Click **Continue**.

NetBackup Network Configuration

> **Note:** Perform the steps in this section only if you have opted to build the cluster with media servers.

To configure the NetBackup network:

1. In the **Configure NetBackup Network** page, provide the required details to connect to the NetBackup external primary server and media servers. Enter the FQDN of the NetBackup external primary server and a valid API key to enable secure

communication between the media servers on the cluster and your NetBackup setup.

2. Under **Media Server Access & Network Settings**, provide a unique FQDN and IP addresses for each media server instance running on the cluster nodes. Ensure that the IPs are reachable and the FQDNs are resolvable within your network environment. You have the option to use the first media server's IP to autofill the remaining media server IP addresses.
3. Under **Media Server Internal Network Settings**, provide the subnet IP and mask to define the private IP range. Ensure that this range does not overlap with any existing cluster, app or client networks.
4. Click **Continue**.

Cluster Creation

To create the cluster:

1. In the **Create Cluster** page, provide the cluster details to create the cluster.

Cluster Name	Specify a unique name for the cluster. Only alphanumeric characters and hyphens are allowed. A hyphen cannot be the first or last character. Length cannot exceed 32 characters. No other characters are allowed.
DNS Servers	The IP addresses of the Domain Name System (DNS) servers that the cluster should use. Separate multiple IPs with commas. Ensure that the Active Directory DNS IP address (if applicable) is listed first. Verify that the specified DNS server can resolve the NTP servers and other entities in the system.
Domain Name	The domain name(s) for the cluster.
Search Domains	Specify a domain search list for hostname lookup.

2. Specify the IP address or the Fully Qualified Domain Name of the NTP server(s). The cluster uses the specified NTP server to synchronize the time on all nodes in the cluster.

   Cohesity recommends using the external Google Public Network Time Protocol (NTP) server and specifying multiple servers (time1.google.com, time2.google.com, time3.google.com, time4.google.com). Avoid using the pool.ntp.org or time.nist.org NTP servers, as they are sometimes unavailable and their IP addresses tend to change. If using an internal NTP server, use only one server (and no external servers).
   Also, toggle on **Use Authentication Key** to secure the communication between the NTP server and the cluster. In the Key ID field, enter the Key ID associated with the SHA-1 key, and in the Key field, enter the SHA-1 key.

> **Note:** It is not mandatory to toggle on **Use Authentication Key** if you are using an internal NTP server.

3. If desired, toggle **Encryption** on. Turning on encryption for the cluster encrypts all data that will be stored on the cluster.

   > **Note:** To encrypt an entire cluster, you must specify the encryption option while creating the cluster. After a cluster has been created, cluster encryption is not editable. If encryption is not enabled for a cluster, you can enable encryption at the Storage Domain level. The clusters use FIPS 140-3 compliant algorithms for encryption. The NetBackup media server is FIPS 140-2 compliant.

4. After Encryption is turned on, the following option is available:
   Rotation Period specifies how often the cluster's encryption key is rotated. After the time period is reached, the old encryption key is replaced by a new key and the data on the cluster remains as it was originally encrypted. The Rotation Period default value is 90 days. You can change this to the desired value.
5. Click **Create**.
6. The **Creating Cluster** page displays the cluster creation progress. You can click View Details to track the progress. Wait for several minutes to allow services to restart.
7. Click **Sign in** and log in to the cluster UI using the default System Admin account called *admin*, and the default password, *admin*.

Create Cluster using CLI

To create cluster using the CLI:

1. Run the following command in the iris_cli to create a cluster:

```
iris_cli cluster create enable-software-encryption=true enable-fips-
mode=true dns-server-ips=<<DNS_IP>> vips=<<VIP_1,VIP_2,VIP_3>>
hostname=<<VIP_Hostname>> domain-names=<<Domain_Name>> ntp-
servers=<<NTP_Server>> name=<<Cohesity_Name>>  subnet-
gateway=<<Gateway>>  node-ips=<<Node_IP1,Node_IP2,Node_IP3>>  subnet-
mask=<<SubnetMask>>  apps-subnet=<<Apps_Private-Subnet>> apps-subnet-
mask=<<apps_subnet_mask>> nb-primary-server-fqdn=<<NBU_Primary_FQDN>>
nb-primary-server-api-key=<<API_Key>> nb-deploy-media-
servers=true/false nb-media-server-vips=<<NBU_Media_server_VIPS1,NBU_
```

```
Media_server_VIPS1,NBU_Media_server_VIPS1>>  nb-media-server-
fqdns=<<NBU_Media_server_Hostname1,NBU_Media_server_Hostname2,NBU_
Media_server_Hostname3>>  nb-container-subnet=<<Apps_Container_Private_
Subnet>> nb-container-subnet-mask=<<Apps_Container_Private_subnet>>
cluster-type=netbackup
```

> **Note:** All IPs and hostname should have both forward and reverse lookup.

Get Started with Data Cloud for NetBackup

In the **Getting Started** page, perform the following steps to begin using your cluster.

1. Click **Getting Started**.
2. Read the EULA License Agreement and click **Accept Terms and Conditions**.
3. In the **Change Password dialog** box, enter and confirm the new password for the System Admin account. The minimum length of the password must be 8 characters. Click Confirm.
4. (If you have opted to build the cluster with media servers) Register the cluster with the NetBackup primary server so that you can use it for backup and restore operations. Click **Register Storage Server**.

   In the **Register Storage Server** dialog box, enter the admin credentials to register the cluster as storage server with the NetBackup primary server. Click **Register**.
5. (If you have opted to build the cluster with media servers) Once the registration is complete, click **Next**.
6. In the **Create Logical Storage Unit** page, click **Create LSU** to create one or more LSUs to allocate space from the cluster for NetBackup disk pools.
7. In the **Create Logical Storage Unit (LSU)** dialog box, enter the LSU name. To add a new Storage domain, select **Add Storage Domain**. Fill all the required details and click **Create**.

   For more details, see

   If Storage domains have already been created, you can select the Storage domain from the drop-down.
8. Select the **LSU Immutability mode**. The following options are available:
   - Non-WORM
     This mode does not support WORM storage and data retention capabilities.
   - Enterprise mode
     In this mode, you can create WORM storage and define the duration for which you want to retain the data, but you cannot override the WORM policy. A user

with Appliance administrator role can remove the retention lock and expire data but cannot reduce the lock duration.

- Compliance
  In this mode, you can create WORM storage and define the duration for which you want to retain the data. However, you cannot expire data before the defined lock duration.You must specify the minimum and maximum lock duration if you choose Compliance or Enterprise mode.

9. Click **Add**.
10. The LSU gets created and gets displayed on the UI. You have the option to edit or delete the LSU. Create additional LSUs if needed. After all the LSUs have been added, click **Done**.
11. After you see the Dashboard page, you can click **Launch NetBackup Web UI** to launch the NetBackup Web UI to configure storage units and disk pools.

> **Note:** This option is available only when the cluster is set up with media servers.

NetBackup Media Servers

The **Summary > Services** page shows the status of the NetBackup media servers that are running on the appliance.

It has details such as the FQDN, status, NetBackup version and IP addresses of the media servers.

You can perform the following operations on the media server:

- Deactivate Media Server
- Stop Media Server
- Restart Media Server
- Start Media Server
- Activate Media Server

> **Note:** The NetBackup media server is FIPS 140-2 compliant.

For clusters deployed without media servers, the cluster must be registered as a storage server using external media servers from the NetBackup Primary web UI. No additional configuration steps are required on the cluster side.

Deactivate Media Server

To deactivate media server:

1. Navigate to **Summary > Services**.
2. Locate the media server that you want to deactivate and click the Actions menu on the right side of the row. Click **Deactivate**.
3. In the **Deactivate Media Server** window, type YES to confirm that you want to deactivate and select **Deactivate**.

> **Note:** Deactivating the media server will prevent NetBackup from assigning new backup or restore jobs to it. Running jobs will continue until completion. Policies and schedules using this server may queue or be redirected based on availability. Existing data and configuration will not be affected.

The selected media server is deactivated.

Stop Media Server

To stop a media server:

1. Navigate to **Summary > Services**.
2. Locate the media server that you want to stop and click the Actions menu on the right side of the row. Click **Stop**.
3. In the **Stop Media Server** window, type YES to confirm that you want to stop and select **Stop**.

   > **Note:** Stopping the media server will immediately terminate all running backup and restore jobs on this server. Active jobs may fail and require manual retry. New jobs will not be processed until the server is started again. Connected clients may experience interruptions.

   The selected media server is stopped.

Restart Media Server

To restart a media server:

1. Navigate to **Summary > Services**.
2. Locate the media server that you want to restart and click the Actions menu on the right side of the row. Click **Restart**.

3. In the **Restart Media Server** window, type YES to confirm that you want to restart and select **Restart**.

   > **Note:** Restarting the media server will temporarily stop services and terminate running jobs. Active jobs may fail and require manual retry. New job processing will resume once the server is fully operational.

   The selected media server is restarted.

Start Media Server

To start a media server:

1. Navigate to **Summary > Services**.
2. Locate the media server that you want to start and click the Actions menu on the right side of the row. Click **Start**.This option will be available only when the media server is in stopped state.

   The selected media server is started.

Activate Media Server

To activate a media server:

1. Navigate to **Summary > Services**.
2. Locate the media server that you want to activate and click the Actions menu on the right side of the row. Click **Activate**. This option will be available only when the media server is in paused state.

   The selected media server is activated.

Manage Storage Domains

From the Manage Storage Domains page, you can view, add and manage the storage domains that have been created on the cluster.

A storage domain is a named storage location on a cluster. A storage domain defines the policy and frequency for deduplication among other configurations such as compression, erasure coding, and replication factor. A storage domain contains views.

You can perform the following operations:

- Create Storage Domain
- Delete Storage Domain
- Edit Storage Domain

Create Storage Domain

To create a storage domain:

1. Navigate to **Settings > Storage Management** and select the **Storage Domains** tab.
2. Select **Create Storage Domain** located at the right side of the page. Fill all the required details in the Create Storage Domain window.
3. **Storage Domain Name**: Enter a name for the storage domain.
4. **Authentication Provider**: You can map the storage domain to an AD or LDAP domain. Mapping to a specific domain means that users from only that domain have permission to view the storage domain's contents when accessing through SMB. By default, a storage domain is accessible to users from all joined domains.
5. **Deduplication**: To enable deduplication, turn it on. If deduplication is enabled, the cluster eliminates duplicate blocks of repeating data stored on the cluster thus reducing the amount of storage space needed to store data.

   Deduplication can be:

   - Inline (Recommended): Deduplication occurs as the cluster saves blocks to the partition.
   - Post Process: Deduplication occurs after the cluster writes data to the partition.
6. **Compression**: To enable compression, turn it on. If compression is enabled,the cluster compresses all the data stored in the storage domain thus reducing the amount of storage space needed to store the data. If compression is on and then turned off later, the current data on the storage domain is not decompressed but any new data stored on the storage domain is not compressed.

   Compression can be:

   - Inline (Recommended): Compression occurs as the cluster saves blocks to the partition.
   - Post Process: Compression occurs after the cluster writes data to the partition.

   > **Note:** If both compression and inline deduplication are on, Cohesity recommends you turn Inline Compression on.
7. **Encryption**: To enable encryption for the storage domain, turn it on. This option is enabled only if cluster level encryption is not enabled.

> **Note:** Encryption can be enabled only when you are creating the storage domain. Once the storage domain is created, the options to enable encryption or modify its encryption settings are disabled in the cluster.

> **Note:** If encryption is enabled at the cluster level, then data within the storage domains are also encrypted and this option is not enabled for each storage domain.

By default, Cohesity uses Internal KMS for encryption.

You can register an external KMS and select one of the following from the drop-down list:

- AWS
- KMIP Compliant

For more information, see External Key Management Service.

8. Expand **More Options** to see the list of additional settings.
9. **Physical Quota & Alerts**: Turn it on to set the total quota for the storage domain. Includes space on the cluster as well as cloud tier. The capacity must be a whole number.

   Alert Threshold: Turn it on to trigger an alert when the storage domain reaches the specified threshold.
10. **Default View Quota & Alerts**: Turn it on to set a default logical capacity limit for the storage domain and set the capacity. Logical data is fully hydrated and expanded.

    Alert Threshold: Turn it on to trigger an alert when the view reaches the specified threshold.
11. **Fault Tolerance & Redundancy**: Select the number of hard drive, node, chassis, or rack failures that can occur before a cluster health alert is triggered. The available choices are determined by the number of nodes, chassis, racks in the cluster. The cluster's global custom failure domain setting determines the failure domain (node, chassis or rack) and the number of failures tolerated at the storage domain level.

1D:1N	Tolerate failure of 1 drive or 1 node. The available options to select are RF2 or EC x:1.
2D:1N	Tolerate failure of 2 drives or at least 1 node. The available options to select include RF3 and EC x:2. This option is available only if the Cluster has three or more nodes. Cohesity will always make the best effort to place each EC stripe unit to be on separate nodes when sufficient nodes are available. Only

when sufficient nodes are not present the EC stripe units will be placed within the same node but on separate disks.
Tolerate failure of 2 drives or 2 nodes.

This setting is only available if there are at least four nodes.

With this setting, the available options are RF3 or EC x:2. Cohesity will always try to place each EC stripe unit on separate nodes when sufficient nodes are available.

> **Note:** You cannot change the fault tolerance setting to 2D:2N if the metadata usage is >=50% of the available storage.

2D:2N

> **Note:** This Erasure Coding (EC) setting is not recommended for clusters with fewer than five nodes because the data cannot be healed after one node failure. Your data will be at risk until the failed node is replaced.

Tolerate failure of 3 drives or 1 node.

3D:1N

This setting is only available if there are at least three nodes. With this setting, the available options are RF4, EC x:3, or EC x:4.

Tolerate failure of 3 drives or 2 nodes.

3D:2N

This setting is only available if there are at least seven nodes. With this setting, the available options are RF4, EC x:3, or EC x:4.

12. **Erasure Coding**: Turn it on for better fault tolerance and capacity utilization. You can also select an Erasure Coding setting from the drop-down list. The available choices are determined by the Fault Tolerance setting and the number of nodes, chassis, or racks in the cluster. The settings indicate the ratio of the data stripes to the code stripes, expressed as x:y. For example, 2:1, 2:2, 3:1, 3.2, 4:1, 4:2, 4:3, 5:1, 5:2, or 5.3.
13. Select **Create**.

Delete Storage Domain

You can delete a storage domain to free up storage space if required.

To delete a storage domain:

1. Navigate to **Settings > Storage Management** and select the **Storage Domains** tab.
2. Select the Actions menu of the storage domain you want to delete and select **Delete**.

Edit Storage Domain

You can edit a storage domain to free up storage space if required.

To edit a storage domain:

1. Navigate to **Settings > Storage Management** and select the **Storage Domains** tab.
2. Select the Actions menu of the storage domain you want to edit and select **Edit**.
3. Make the required changes in the **Edit Storage Domain** window.
4. Select **Save**.

Manage Storage Servers

A storage server is the set of media servers in the cluster that are registered as a storage server with NetBackup. The **Storage Servers** tab gives the details of the storage servers that are registered along with the media servers assigned to each storage server.

> **Note:** This tab appears only if you have opted to build the cluster with media servers.

You can perform the following actions from this page:

- Register a Storage Server
- Assign Media Servers to a Storage Server
- Unregister a Storage Server

Register a Storage Server

To register a new storage server:

1. Navigate to **Settings > Storage Management** and select the **Storage Servers** tab.
2. Select **Register Storage Server**.
3. In the **Register Storage Server** dialog box, you can:
    - Choose Use FQDN and select the FQDN from the drop-down to connect to the NetBackup storage server.
    - Choose Use VIP and enter the address to connect to the NetBackup storage server.
4. Enter the cluster credentials to register the cluster as a storage server with the NetBackup primary server.
5. Select **Register**.

> **Note:** In a multi-VLAN environment, for each VLAN you have to register the storage server individually using its cluster FQDN or virtual IP.

Assign Media Servers to a Storage Server

All media servers of the cluster are automatically assigned to the storage server when you register the storage server.

In a multi-VLAN environment, you have to manually add newly added media servers to storage server.

To manually assign media servers to a storage server:

1. Navigate to **Settings > Storage Management** and select the **Storage Servers** tab.
2. Select the Actions menu of the storage server to which you want to assign media servers and select **Assign Media Servers**.
3. In the **Assign Media Servers** dialog box, enter the cluster's credentials to establish the connection. The list of media servers that will be assigned to the storage servers is displayed.
4. Select **Assign Media Servers**.

The media servers are assigned to the storage server and get displayed on the **Storage Servers** page.

Unregister a Storage Server

You can edit a storage domain to free up storage space if required.

To unregister a storage server:

1. Navigate to **Settings > Storage Management** and select the **Storage Servers** tab.
2. Select the Actions menu of the storage server to which you want to unregister and select **Unregister**.
3. In the **Unregister Storage Server** dialog box, confirm that you want to unregister the storage server and select **Delete**.

The storage server is no longer registered as a storage server with NetBackup.

Manage Logical Storage Units

The Logical Storage Unit tab gives details of the Logical Storage Units (LSUs) created on the cluster. You can also create new LSUs and delete an existing LSU from this page.

You can perform the following actions from this page:

- Create a Logical Storage Unit
- Delete a Logical Storage Unit

Create a Logical Storage Unit

To create a new LSU:

1. Navigate to **Settings > Storage Management** and select the **Logical Storage Unit** tab.
2. Select **Create LSU** to create one or more LSUs to allocate space from the cluster for NetBackup disk pools.
3. In the **Create Logical Storage Unit (LSU)** dialog box, enter the LSU name.
4. Select the **Storage Domain** from the dropdown.

   If you want to add a new Storage domain, select **Add Storage Domain**. For more information, see Manage Storage Domains
5. Select the LSU Immutability mode. The following options are available:
   - Non-WORM
     This mode does not support WORM storage and data retention capabilities.
   - Enterprise mode
     In this mode, you can create WORM storage and define the duration for which you want to retain the data, but you cannot override the WORM policy. A user with Appliance administrator role can remove the retention lock and expire data but cannot reduce the lock duration.
   - Compliance
     In this mode, you can create WORM storage and define the duration for which you want to retain the data. However, you cannot expire data before the defined lock duration.You must specify the minimum and maximum lock duration if you choose Compliance or Enterprise mode.
6. Select **Create LSU**.

Delete a Logical Storage Unit

To delete an LSU:

1. Navigate to **Settings > Storage Management** and select the **Logical Storage Unit** tab.
2. Select the **Delete** icon in the row corresponding to the LSU you want to delete and select **Delete LSU**.

Add Nodes to the Cluster

Before you begin:

- Ensure connectivity to all the existing sources on the newly added nodes.
- VLAN virtual IPs need to be provisioned first before adding the nodes. Typically, one virtual IP per node should be provisioned in each VLAN. VLAN virtual IPs need not be added for the primary interface VLAN.
- If cluster usage is above 85%, Cohesity recommends adding additional nodes to the cluster.

To add nodes to the cluster:

1. Navigate to **Settings > Summary** and select the **Nodes** tab.
2. Select **Add Node**. The cluster looks for free nodes in the same subnet which are discoverable from the node from where add node option is being done to add to the cluster. The node list is displayed on the Dashboard.
3. Select the node(s) you want to add to the cluster. The cluster can accommodate a maximum of 8 nodes when adding them at the same time. Select **Continue**.
4. Configure the node IP addresses. Specify IP addresses for each of the new nodes. You can use first node's IP to auto-fill remaining IP and IPMI addresses. Select **Continue**.
5. (If you have opted to build the cluster with media servers) In clusters on which multiple VLANs are configured, add virtual IPs for the new node if the number of virtual IPs previously assigned to the VLAN is less than the total number of nodes currently in the cluster. This ensures better workload balancing. Add a media server VLAN IP and FQDN for the new node.
6. (If you have opted to build the cluster with media servers) Configure IP addresses and FQDN for each of the media servers. You can use the first node's IP to auto-fill the remaining media server IP addresses.
7. (If you have opted to build the cluster with media servers) Under **Storage Server Registration**, provide this cluster's credentials to add new media servers to the storage server already registered with the NetBackup Primary server.
8. Select **Add**. You can see the progress of the add node progress using alerts.

> **Note:** If a node's software version is different from the cluster's software version, the node's software is automatically updated to the cluster's version when you add the node. The IPMI configuration is optional. You can configure IPMI after adding nodes to the cluster using the CLI or Dashboard.

> **Note:** During the **Add Node** operation, progress details may temporarily disappear from the UI. This is a known behavior. The node addition continues in the background even if progress is no longer visible. If progress details disappear, do not retry the operation or attempt additional configuration actions. Wait for the operation to complete in the background.

Chapter 7: Organizations (Multitenancy)

Cohesity Platform cluster supports the multitenancy feature. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. You can use a single Cohesity cluster to provide service to multiple tenants. See the following topics for details:

- Service Providers (SPs) 866
- Large Enterprises 866
- Organization User Types 867
- Organizations Subfeatures 868
- Prerequisites and Considerations for Organizations 870
- Privileges for Service Provider Admin and Organization Admin 871
- Best Practices for Organizations 872
- Enable Organizations (Multitenancy) on a Cluster 873
- Add an Organization (Tenant) 874
- Infrastructure for Organizations 876
- Access Management for Organizations 908
- Users and Groups for Organizations 916
- Data Protection for Organizations 917
- Impersonate an Organization 920
- Deactivate an Organization 920
- Activate an Organization 921
- Delete an Organization 921
- CloudArchive and CloudTier for Organizations 922
- Replication for Organizations 923
- Reports and Chargeback for Organizations 923

Supported Multitenancy Workflows 927
Tenant-Level View Allowlist 1213
Key Management Service (KMS) in Multitenancy 1214

> **Note:**
> - In the context of Cohesity, the multitenancy feature is termed as Organizations. Each tenant on the Cohesity cluster is called an Organization.
> - Once the Organizations feature is enabled on a cluster, you cannot disable it.

You can configure the Organizations feature on the Cohesity cluster for the following environments:

- Service Providers (SPs)
- Large Enterprises

Service Providers (SPs)

As a service provider, you can use the Cohesity cluster to provide storage-related, data backup, data management services, etc. to your customers.

Cohesity Platform allows tenants to manage their organizations based on the privileges provided by you as a service provider. This allows organizations (tenants) to perform routine tasks without depending on you as the service provider.

You can provide the following services as a service provider:

- Backup as a Service (BaaS)
- Disaster Recovery as a Service (DRaaS)
- Filer as a Service (FaaS)

For more information on solutions, see the following documents:

- Backup as a Service Solution Guide
- Disaster Recovery as a Service (DRaaS) Solution Guide

Large Enterprises

Multitenancy is not just for service providers; it's also useful for large enterprises. Large enterprises can use multitenancy to isolate departments or business units as an individual tenant from each other for security and compliance reasons.

Organization User Types

In the Organization environment, we define the following user types for clarity of communication and understanding only. These user types may be different from jargon used elsewhere in the Cohesity documentation.

> **Note:** The following user types are used in Cohesity's Organizations (multitenancy) environment. These user types are not defined in the Cohesity Dashboard (UI).

Service Provider Administrator

A service provider is a company that provides IT services to its customers using the Cohesity cluster. A service provider administrator is typically a user with the default Administrator role.

> **Note:** Technically, the service provider administrator is the same as the **Cluster administrator**. We use the term **service provider administrator** in the context of the **Organization** feature, where the service provider manages multiple Organizations.

Organization Users

An Organization on the Cohesity cluster can have users and groups.

An Organization user with the Admin privilege is called the **Organization administrator** and users without the Admin privilege are called **Organization users**.

Organization Administrator

An Organization administrator can manage its own Organization in a multitenant environment. The Organization administrator role and its privileges are defined by the service provider administrator. The Organization administrator can create an Organization user or group. However, the Organization administrator cannot create or modify roles.

> **Note:** To add Active Directory as an Organization Admin, the AD domain's IP must be part of your assigned VLAN. Otherwise, please contact the Service Provider admin to add the AD domain to the cluster and assign it to your Organization. For Service Provider use cases, see Privileges for Service Provider Admin and Organization Admin.

Organization Users

An Organization user or group is created by the Organization administrator or it can be assigned to the Organization by a service provider administrator. The Organization administrator can only assign the inherited roles to a new Organization user or group.

The roles can be Viewer, Operator, or other custom roles.

Organizations Subfeatures

The following are some key subfeatures of the Organizations feature.

- A Service Provider Administrator can add Organizations and assign them the following Cohesity resources:
  - Storage Domain
  - View
  - AD, LDAP, SSO, or Keystone
  - Add Organization administrator, users, or groups
  - Objects
  - Policies and protection groups
  - VLAN(s)
- Organization users can log in to the Cohesity GUI for self-service access to their assets and services, in a manner that is securely isolated from other tenants. Cluster-level administrative capabilities are not available to Organization users.
- Organization administrators can configure Role-Based Access Control to control what is visible to tenant users.
- For data protection, each Organization owns private sources, protection jobs.

Storage efficiency

- As a service provider, you can also isolate data in Organizations from each other logically by assigning each Organization a storage domain, which creates separate, independent deduplication domains and encryption keys.
- As a service provider, you can alternatively configure Organizations to use the same storage domain for storage efficiency, which acts as a single deduplication domain. However, access is still secure and limited to each Organization. For more information, see Infrastructure for Organizations.

Network Isolation

- Service provider admins can isolate Organizations from each other by assigning VLAN.
- On each VLAN, service providers can assign cluster virtual IP addresses to each tenant.
- Alternatively, you can use Cohesity's Hybrid Extender to overcome IP address conflicts across the tenants. For more information, see Configure Network Settings for an Organization and Hybrid Extender

Impersonate Organizations

- Service providers can impersonate Organizations to preview and verify what their Organization administrator sees and helps Organizations diagnose and debug issues.
    - The Service provider administrator acquires tenant administrator privileges during impersonation.
    - The Service provider operator acquires tenant operator privileges during impersonation.
    - Service provider viewers will get tenant viewer privileges during impersonation.
- All actions performed during impersonation are audit-logged for security reasons. The audit logs are available to both service providers and Organizations. For more information. For more information, see Impersonate an Organization.

Backup and Restore

Cohesity provides efficient backup and restore workflows for service providers and Organizations and helps in managing large environments with multiple workloads using a single pane of glass. For more information, see Data Protection for Organizations.

Encryption

- Cohesity has a built-in key manager that generates keys and stores them internally on the SSDs in encrypted form.
- Data on the cluster is encrypted using a Data Encryption Key (DEK) and the DEK is in turn encrypted using a Key.
- Key Encryption Key (KEK). The encryption and decryption process is done within the Cohesity cluster and is transparent to all inbound and outbound protocols and applications, such as backup, archiving, and file services. Also, the keys are stored in a distributed fashion to be resilient in the rare event of a hardware failure.
- The default key rotation policy for KEK is 90 days, but the administrator can configure a different value.

Third-party Integration

Cohesity is based on an API-first architecture, valuing automation, consumable APIs, and self-service capabilities from the get-go. Users can enjoy error-free consistency and self-service manageability that extends the capabilities of individual applications and enables IT to spend less time with management, and more time with ambitious, value-oriented projects.

Supported automation frameworks include:

- VMware vCloud Director
- VMware vRealize Automation (vRA) and Orchestration (vRO)
- ServiceNow
- Python/Powershell SDK

Reports

- All Organizations can view their storage consumption report.
- As a service provider, you can view storage consumption reports by each tenant and generate custom reports. For more information, see Reports and Chargeback for Organizations.

Prerequisites and Considerations for Organizations

As a Cohesity cluster Administrator, you can enable multi-tenancy permanently for a cluster, add Organizations to the cluster, and assign resources and local users to the organization.

Prerequisites

- You must enable the Organizations feature to support multi-tenancy.
- Use VLANs to segregate traffic across different tenants. A VLAN should be assigned to each tenant. If it's not assigned, then the views in the tenant may become visible but not accessible to other tenants.

Considerations

- Once the Organizations feature is enabled on a cluster, you cannot disable it.
- Once an organization is created with a particular name/OrgID, the name/OrgID cannot be reused even after the org is deleted.

- If a single-tenant cluster is configured with remote access to a multi-tenant-enabled cluster, the Organizations page will not be available when accessing the multi-tenant cluster. The workaround is to enable multi-tenancy on the single tenancy cluster (it is not necessary to add any organizations.)
- Storage Consumed by Storage Domains: The Storage Consumed by Storage Domains report is available to the organization only if Storage Domain Sharing is disabled.
- View names must be unique. Two Organizations or service providers admin, or tenant cannot have two same view names.
- Users are assigned to the organization with specific roles by the service administrator. Cohesity does not allow an organization user to create a new user with a different role other than the organization's inherited roles.
    - For example, the service provider administrator creates two users with the "Viewer" and "Admin" roles and assigns them to the Organization. The organization administrator can now create a new user with either a "Viewer" or "Admin" role but cannot create a new user with a different role.
- No DNS per tenant. Global DNS must have stub zones for all tenant DNS.
- No multi-tenancy support for Marketplace Apps.
- If the **Enable for all Organizations** option is selected for a VLAN and the VLAN is assigned to an Organization in the OpenStack environment, the Swift clients cannot access Cohesity Swift Service through this shared VLAN.
- Organizations (tenants) do not have access to the Dashboard in the Cohesity UI.
- Replication to a source or remote cluster fails if the source or remote cluster has multi-tenancy enabled and VLAN assigned to the tenant. For more information, see the Replication stops sending data after upgrade to Cohesity release 6.4.1c and later KB article.
- To create a SmartFiles S3 View in a multi-tenant environment, log in to the Cohesity cluster as an Organization user. If you create the S3 View while impersonating an organization, the Service Provider administrator becomes the owner of the S3 View.

Privileges for Service Provider Admin and Organization Admin

The following table lists the privileges of the service provider admin that are distinct from those of the organization administrator.

Resource	Service Provider Administrator/Cluster Admin					Organization (Tenant) Administrator	
	Create	Assign	Unassign	Delete	Can be shared across tenants?	Add/Create	Delete
Storage Domain	Yes	Yes	No	Yes	Yes	No	No
View	Yes	Yes	No	Yes*	No	Yes	Yes
Full Source Tree	Yes	Yes	No	No	No	Yes	Yes
Partial Source Tree	Yes	Yes	Yes	NA	Yes	No	No
Protection Policy	Yes	Yes	Yes	Yes	Yes	No	No
Protection Job	Yes	Yes	No	No	No	Yes	Yes
VLAN	Yes	Yes	Yes	Yes	Yes	No	No
User/Group Management	Yes	Yes	Yes	Yes	No**	Yes	Yes

Note:

- * - Service provider administrator cannot delete a view assigned to an organization. But it can delete any unassigned views on the cluster.
- ** - Users cannot be shared across tenants. However, in the Single Sign-on (SSO) environment, groups can be shared.

Best Practices for Organizations

The following are best practices you can implement in Cohesity's Organizations (multitenancy) environment.

- To get better performance for VMware backups, use VMkernel ports (usually 10G) for backup instead of ESXi management (usually 1G) ports. For more information, see How to determine which vmkernel port is being used in a Cohesity VMware VM Backup in the Cohesity Knowledge Base.
- For a Hosted Backup solution with a shared, multitenant hypervisor, as a service provider, you should make sure the VLAN carrying the backup traffic has the Enable For All Organizations option enabled. This ensures the same backup network can be used for backing up different tenants' data.
- Use VLANs to segregate traffic across different tenants. Each tenant can be assigned a VLAN.
- Use different VPN tunnels for different tenants if the tenants are remote and accessing the Cohesity cluster over the Internet. This is important for secure self-service restores.
- Naming conventions in a multitenancy environment:
  - Create your naming convention for the Organization ID to work with 10 characters.
  - The Organization ID cannot be edited later. If the organization is deleted later, the Organization ID cannot be reused.
  - Use the same Organization ID across clusters.
- Follow the same practice when using CloudRetrieve; the same Organization ID and Storage Domains must exist on both clusters prior to running a CloudRetrieve download task.
- For service providers who are implementing a BaaS solution using Cohesity Data Cloud (Self-managed), Cohesity recommends that you segregate replication and backup traffic on the source cluster to avoid bandwidth issues. For more information, see Disaster Recovery as a Service (DRaaS) Solution Guide.
- Cohesity also recommends having a VPN connection between the source and DR clusters to set up replication across the clusters if the clusters are not using public IPs for replication communication.
- Alternatively, you can replicate over WAN without the need for a VPN connection.

Enable Organizations (Multitenancy) on a Cluster

To use the Organizations feature and create Organizations, you must enable the feature. From the **Edit Cluster Settings** page, you can enable the feature.

> **Note:** You cannot disable the Organizations feature once it is enabled on a cluster.

After enabling the Organizations feature, you can add new Organizations from the **Access Management** page.

To enable the Organizations feature, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Summary** and click **Configure**.
2. Toggle **Enable Organizations**.

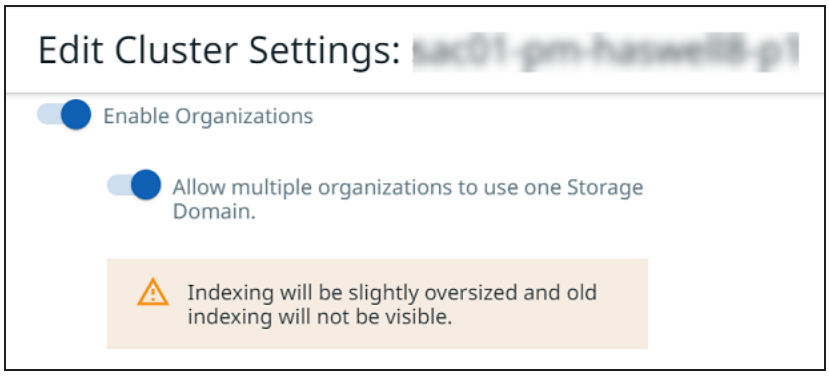

> **Note:** Ignore the above warning on indexing as it is not actionable and does not indicate any issue when enabling Organizations on the cluster.

3. (Optional) Toggle on **Allow multiple Organization to use one Storage Domain**.

> **Note:** When you select this option, the storage consumption on the cluster reduces by providing Cohesity deduplication across all the Organizations using the storage domain. Organizations will be isolated from other Organizations whether this option is on or off.

4. Click **Save**.

Add an Organization (Tenant)

An Organization has a name and an organization ID. After you create an organization, you can edit and configure it according to your requirements. You can also deactivate or delete the Organization.

Prerequisite

The Organization feature is enabled. For information, see Enable Organizations (Multitenancy) on a Cluster.

To add an organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Organizations** and click the **Add Organization** button on the top-right of the screen.

2. In the **Add Organization** page, type a name for an organization in the **Organization Name** field.

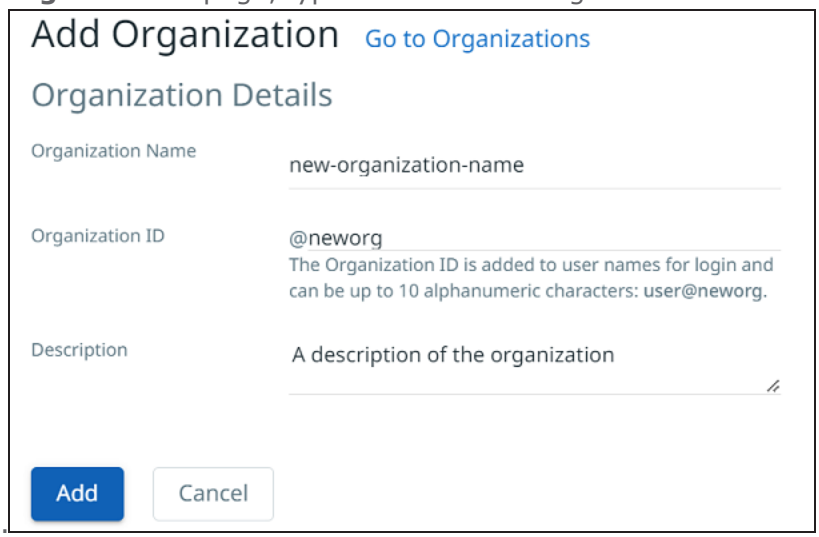

> **Note:** You can edit the Organization name after adding the Organization.

3. Type an ID for the Organization.
   - An identifier, up to ten alphanumeric characters, which is used as the suffix for all user names in the organization. For example, if you enter the ID abc123, user joe would log in to the cluster with the user name joe@abc123.
   - The Organization ID cannot be edited later. If the organization is deleted, the Organization ID cannot be reused.
   - Use the same Organization ID across clusters.

> **Note:** If you need to increase the number of characters in an Organization name, contact the Cohesity Support team.

4. (Optional)Type description for the organization.
5. Click **Add**.

   The **Edit Organization** page appears. You can configure your organization immediately after adding the Organization or editing the Organization later and configure it.

Infrastructure for Organizations

Assign a Storage Domain and View

You must assign a Storage Domain to the Organization for users in the Organization to use the Cohesity cluster.

> **Note:** Once you assign a storage domain to the Organization, you cannot unassign it. However, you can add multiple storage domains to the Organization.

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization, and select **Edit**.
3. In the **Edit Organization** page, click the **Storage Domains** drop-down and select desired storage domains.
4. Click **Views** drop-down list and select the required View.

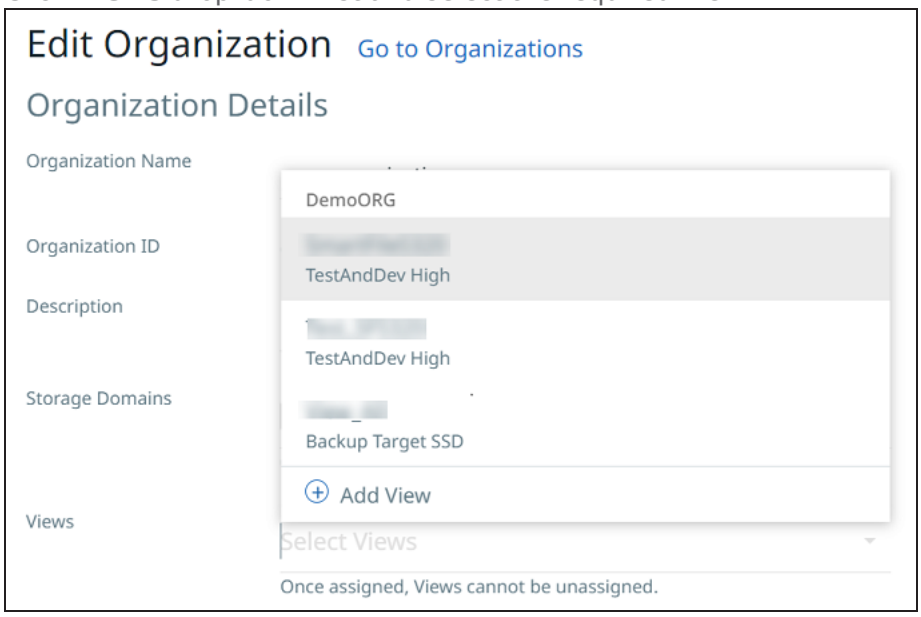

> **Note:** A View assigned to an Organization cannot be accessed or unassigned by the cluster admin. However, the **View** can be deleted by the Organization admin or by the cluster admin by impersonating the Organization.

5. Click **Save**.

Configure Network Settings for an Organization

The network settings for an Organization involves assigning VLANs (optional) and set service provider side NAT.

VLANs

To further isolate network traffic for the Organization, assign VLAN IDs. VLANs can be added and removed later. Use VLANs to segregate traffic across different tenants. If VLANs are used they must be assigned to tenants. Each tenant can be assigned a VLAN.

> **Note:** For a Hosted Backup solution with a shared hypervisor, as a service provider, you should make sure the VLAN carrying the backup traffic has the **Enable For All Organizations** option enabled. This ensures the same backup network can be used for backing up different tenants' data.

For more information, see Manage VLANs

IPs on the NAT

Specify the hostname or IP Address. The Hostname or Address(es) provided here corresponds to the public IPs on the NAT (in front of the Cohesity nodes) so that the Organizations can use these entries to reach out to the Cohesity nodes.

To assign VLAN to an Organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization, and select **Edit**.
3. In the **Network Segment** section of the **Edit Organization** page, click **Assign VLANs**.
4. In the **Assign VLANs** page, select a **VLAN** and click **Assign**.
5. In the **Cohesity Node IP Range** field of the **Edit Organization** page, type hostname, or IP address.
6. Click **Save**.

Hybrid Extender

As a service provider, if you are serving multiple tenants (customers) in isolated networks with the same IP address ranges and subnets on the same Cohesity cluster, you may need to enable the Hybrid Extender.

Once the service provider administrator or cluster administrator enables the Hybrid Extender, the Organization (tenant) administrator can deploy and configure the Hybrid Extender.

> **Note:** If you are upgrading the Cohesity cluster from one maintenance release to another maintenance release, you need not upgrade the Hybrid Extender. For example, if you are upgrading the Cohesity cluster from 7.2.1 to 7.2.2, you need not upgrade the Hybrid Extender. However, Cohesity recommends that the version of Cohesity cluster and the Hybrid Extender should be the same.

In a multitenant environment, if organizations are controlling the IP addressing for their virtualization hosts, virtual machines, physical servers, NAS servers, Active Directories, and other servers (collectively "tenant entities"), there is a possibility that the tenant entities' IP addresses overlap with each other. As a result, the Cohesity cluster will be unable to uniquely identify a given tenant entity using its IP address (Example: tenant A and tenant B might both have SQL Server VMs with identically the same IP address on their respective tenant private networks). To resolve this issue, you can deploy a Hybrid Extender VM on each organization to uniquely identify the virtual entities using the tuple of organization ID and virtual entity's IP address. The Hybrid Extender acts as a proxy VM between the organization and the Cohesity cluster and also provides source-side deduplication.

Hybrid Extender Support

Communication between the Hybrid Extender VM and the Cohesity cluster is over-secure gRPC and encrypted using mutual TLS. The Hybrid Extender VM handles all the outbound traffic from the Cohesity cluster to supported platforms (VMware vCenter, VMware ESXi, and to the "physical server", and the inbound traffic of VMware backup ingest traffic from VMware ESXi (or physical server) to Cohesity cluster.

For more information about the supported workflows, see Supported Multitenancy Workflows for Hybrid Extender.

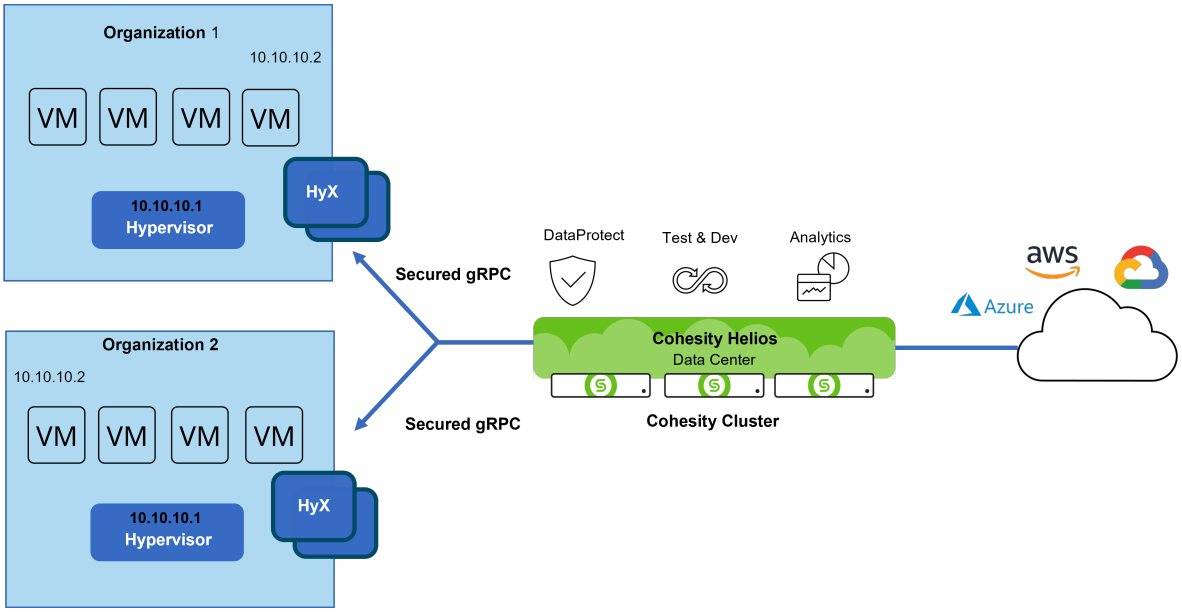

Prerequisites

- Ensure that the Hybrid Extender option is enabled while creating the organization.
- (Default) Ensure port 29994 is open on the Hybrid Extender VM, and ports 11117 and 29991 are open on the Cohesity cluster.
- DNS configuration—The DNS configuration on your Hybrid Extender is used for general name resolution and to look up the domain controller to join the Cohesity cluster to the tenant AD. This is used during source registration for name resolution and mount-based restores/clones for data protection of MS-SQL server, Oracle, and so on.

  The following are the considerations for DNS configuration on Hybrid Extender:

  - When you assign an IP address using DHCP, the Hybrid Extender automatically gets the DNS configuration.
  - When you assign a static IP address, you must add the DNS in the Hybrid Extender VM by accessing the VM using SSH. You can access the Hybrid Extender VM and add your DNS in the `/etc/resolv.conf` file. For more information on the steps to edit the `/etc/resolv.conf` file, see Editing for Hybrid Extender VM. For any assistance, you may contact the Cohesity Technical Support team.
- Register the AD domain of the tenant and on the DNS server of the tenant, create DNS entries for the registered AD machine accounts with Hybrid Extender IPs. This ensures that the clients connected to the tenant network are able to access the views.
- Ensure that Network Time Protocol (NTP) is set up on the Hybrid Extender. For more information, see Set Up NTP Servers.

  If NTP is not setup:

  - RPC requests between the Hybrid Extender and the cluster can timeout.
  - If the time skew between the AD server and the Hybrid Extender is greater than five minutes, Kerberos authentication may fail.
  - Hybrid Extender may not be able to connect if the time lag between the Hybrid Extender and the cluster is extensive enough to invalidate TLS certificates.

Authentication

NFS—Hybrid Extender supports only LDAP group membership lookups.

SMB—Hybrid Extender supports the following authentication methods:

- NTLM authentication—Tenants can mount their shares using the server IP address or FQDN. The tenant must be registered with only one AD domain.
- Local user account—The Service Provider can create a local user account and assign this user to a tenant.

- Kerberos authentication.
- Anonymous users.

Considerations

- Hybrid extender supports source registration and backup only for Windows and Linux physical sources. AIX, HPUX, Solaris physical sources are not supported with hybrid extender.
- Currently, Cohesity does not support the auto-upgrade of the Hybrid Extender. Therefore, you must upgrade the Hybrid Extender after upgrading the Cohesity cluster from one major release to another major release. For example, if you are upgrading the Cohesity cluster from 6.5.1 to 6.6, use the Hybrid Extender version provided with 6.6.
- When you're upgrading to maintenance releases such as 6.5.1e, you need not upgrade the Hybrid Extender. However, Cohesity recommends that the version of Cohesity cluster and the Hybrid Extender to be same.
- If a tenant deploys multiple Hybrid Extender VMs, SMB and NFS sessions do not failover to the next available Hybrid Extender VM. Cohesity depends on the hypervisor that is hosting the Hybrid Extender VM to ensure high availability. If the hypervisor does not support high availability, I/O requests fail.
- Hybrid Extender does not support the following features:
  - S3
  - SMB Multichannel
  - Keystone
  - Kerberos client for NFS
  - SSO
  - NFS authentication

Hardware Requirements for Hybrid Extender VM

The default configuration is as mentioned below:

- **Memory**: 4GB RAM
- **Processor**: 4 vCPU

Enable Hybrid Extender

To enable the Hybrid Extender, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab, click ⋮ for required organization, and select **Edit**.

3. In the **Network Segments** section of the **Edit Organization** page, toggle the **Hybrid Extender**.
4. Click **Save**.

Disable Hybrid Extender

Before you disable the Hybrid Extender, ensure that no protection or recovery job is running. After you disable the Hybrid Extender, the tenant network falls back to VLAN-based multitenancy. For information about assigning VLANs, see Configure Network Settings for an Organization.

To disable the Hybrid Extender, complete the following steps:

1. In the Cohesity Dashboard, select **Settings > Access Management**.
2. In the **Organizations** tab, click ⋮ for required organization, and select **Edit**.
3. In the **Network Segments** section of the **Edit Organization** page, turn off the **Hybrid Extender** option.
4. Click **Save**.

Support Two Interfaces on Hybrid Extender

Hybrid Extender can be configured with two interfaces so that it can simultaneously talk to a tenant's private network and an entity (Cohesity) outside of the private network. This feature eliminates the need for making network configuration changes (such as NAT mappings or changing firewall rules) for the Hybrid Extender interface to reach the Cohesity. Following are the two network interfaces:

- Primary interface (ens160)
- Secondary interface (ens192)

Configuring Two Interfaces

You can configure the Hybrid Extender using one or two network interfaces. Each network interface can be configured with DHCP or static IPs. DHCP is the default. Static IPs will be used only if you configure a static IP.

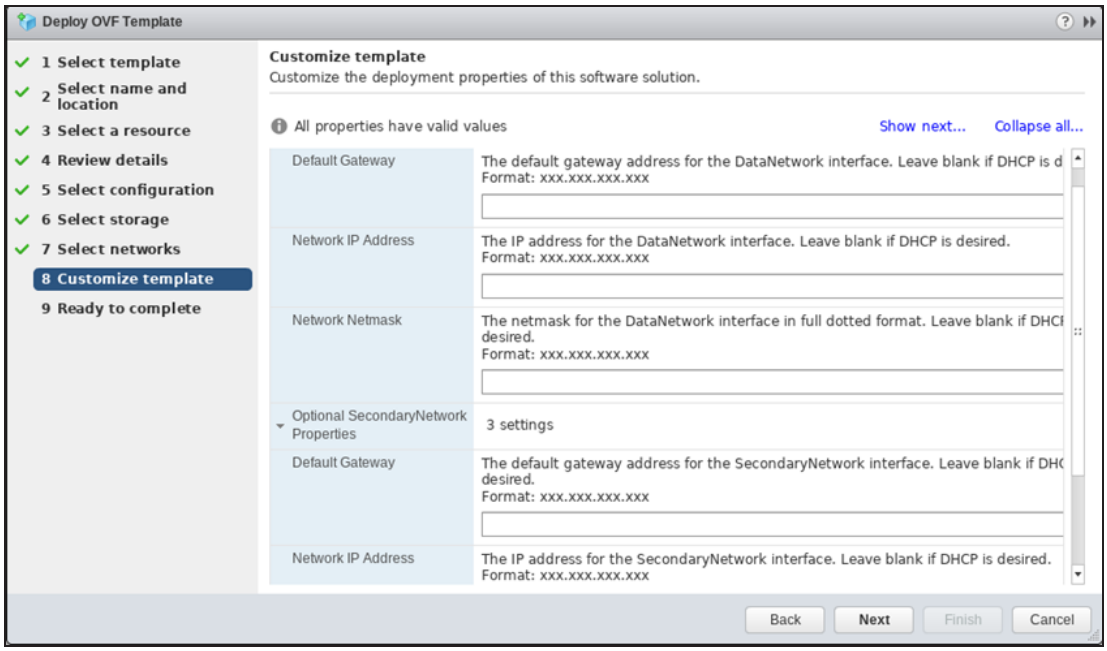

To use Hybrid Extender with one network interface:

1. Configure the first network interface with DHCP or static IPs.
2. Do not specify any configuration for the second network interface. This unused secondary interface will not affect user functionality.

To use Hybrid Extender with two network interfaces:

- Configure both network interfaces as required (using either DHCP or static IPs)

> **Note:**
> - By default, the primary interface is used to communicate with the Cohesity cluster and the secondary interface is used to communicate with entities in the tenant's internal network.
> - If the **Network IP Address** field value for an interface is left empty, Cohesity uses a DHCP server to configure the interface.
> - If network field values are provided as input, Cohesity uses those static values to configure the interface and fall back to DHCP if there is an issue with the values.
> - After deployment, you can remove the second interface manually from the VM if required.

Network Interfaces Combinations

- If you specify the network fields for both interfaces, the interfaces will come up with the specified static IP addresses.

- If you do not specify the network fields for both the interfaces, the interfaces will be automatically configured as DHCP.
- If you specify the network field only for the primary interface, the primary interface will come up with the static IP and the secondary interface is automatically configured as DHCP.
- If you specify the network field only for the secondary interface, the secondary interface will come up with the static IP and the primary interface is automatically configured as DHCP.

Use Cases

Use only Primary Interface

1. In the primary interface, provide an IP address if you want to use a static IP address or leave it blank to use DHCP.
2. Leave the secondary interface blank.

> **Note:**
> The secondary interface comes up with DHCP as you left the field blank. Remove the interface from the VM configuration or use the following command to bring down the secondary interface:
>
> ip link set dev ens192 down

Use the Secondary Interface to communicate with Cohesity Cluster

1. Configure both interfaces as required. For more details, see Network Interfaces Combinations.
2. Once the VM is up, create a **customer_settings.json** file in the data/bifrost directory. This file is used to override default interface behavior.
3. Explicitly mention the interface to be used for communication with the cluster and the interface to be used for internal tenant side communication (Note that explicitly mentioning both of these is important. If only one is mentioned, we will still use the default interfaces).

Sample **customer_settings.json** file:

```
{
"tenant_source_side_interface_name": "ens160",
"cohesity_side_interface_name": "ens192"
}
```

Deploy Hybrid Extender VM

1. In the Cohesity Dashboard, select **Settings > Access Management**.
2. In the **Organizations** tab, click ⋮ for required organization, and select **Edit**.
3. In **the Network Segments** section of the **Edit Organization** page, click **Add Hyx**.
4. On the **How to configure Hybrid Extender?** page, click the download button under **step 1(Download and provision the HyX)** to download the .ova file. You can also download the .ova file from MyCohesity portal.
5. Deploy this OVA file in your ESXi hypervisor as a virtual machine by following the instructions given on the ESXi console.
6. Go back to the **How to configure Hybrid Extender?** page and click the download button under **step 2 (Download the configuration file)**.
7. Go to http://<ip_address>:29994/upload and upload the configuration file you downloaded.

   > **Note:** Where <ip_address> is the IP address of the Hybrid Extender VM.

8. Refresh the Hybrid Extender page. The Hybrid extender VM will connect.

You can deploy more than one Hybrid Extender for each organization. If the backup adapter supports multi-streaming, then having multiple Hybrid Extenders per organization improves high availability, load balancing, and yields faster backups and restores.

- High availability—If one Hybrid Extender instance fails, the Cohesity cluster will retry backup, recovery, and clone workflows using a surviving Hybrid Extender.
- Load-balancing and throughput—When multi-streaming a backup or restore, different network flows may traverse different Hybrid Extenders to parallelize the task and speed the overall operation.

The configuration file you upload (steps 6 and 7) on each Hybrid Extender VM should be unique for each organization.

It is recommended to use a different config file for every Hybrid Extender but you may choose to reuse the same config file for deploying Hybrid Extenders of the same organization. However, for different organizations, use a different configuration file on different Hybrid Extender VMs.

Set Up NTP Servers

The tenant administrator must provide the NTP servers. To update the machine clock, execute the following command:

```
ntpdate -u <ntpserver>
```

Example: `ntpdate -u ntp.for.mydomain`

Using Non-Secure gRPC

Run these commands in Hybrid Extender to use non-secure gRPC:

1. `bifrost.sh stop>`
2. `bifrost.sh start --magneto_grpc_timeout_msecs=300000 --tcp_proxy_connector_ports=445:445,2049:2049,80:80,443:443,11117:11117,11113:11113`
3. `sudo firewall-cmd --add-port=11113/tcp --zone=public --permanent && sudo firewall-cmd --reload`

Recommendations

- Deploy multiple Hybrid Extender VMs to distribute client traffic and achieve high availability and increased throughput.
- Download a new configuration file from Cohesity Cluster every time you deploy a new Hybrid Extender VM.
- Cohesity recommends that you use the Hybrid Extender version provided with the Cohesity version to get access to the supported sources. For example, use the Hybrid Extender version provided with Cohesity 6.6.0a.
- Once the Hybrid Extender is connected, the backup and recovery communication between cluster and source will be proxied through the Hybrid Extender VM.
- If you are upgrading the Cohesity cluster, you must upgrade the Hybrid Extender. To upgrade the Hybrid Extender, you should first redeploy the new version and then decommission the existing version of the Hybrid Extender. After you upgrade a cluster, then the version of Hybrid Extender should be either the same as Cohesity cluster or it can remain in the previous version. However, Cohesity recommends that the Cohesity cluster version and the Hybrid Extender version should be the same. As an example, you can refer to the table below:

Cohesity Cluster Version	Hybrid Extender Version
6.6.0a	• 6.6.0 • 6.6.0a

Ports Used by Hybrid Extender

For more information about ports used by Hybrid Extender, see Hybrid Extender.

Network Realms for Multi-tenant Environments

Cohesity introduces Network Realms for multi-tenant environments. Network Realm ensures source-level selection in a multi-tenant environment, and allows you to create and manage multiple isolated networks within a tenant. On a Cohesity cluster, a source and a hybrid extender can be associated with a Network Realm to create an isolated network.

How it works?

When a backup task is initiated on a Cohesity cluster, it reaches out to the Hybrid Extender associated with the Network Realm for the source IP address. You can associate the source with the Network Realm during its registration and each source can be associated with only one Network Realm. However, you can associate multiple Hybrid Extenders with a Network Realm. The Network Realm maintains the list of registered sources on a Cohesity cluster to create isolated networks. An isolated network ensures that a task on a Cohesity cluster doesn't fail due to a mismatch of IP addresses from multiple sources.

> **Note:** You must have an **Admin** or **Super Admin** role or similar privileges in a custom role to create and manage Network Realms.

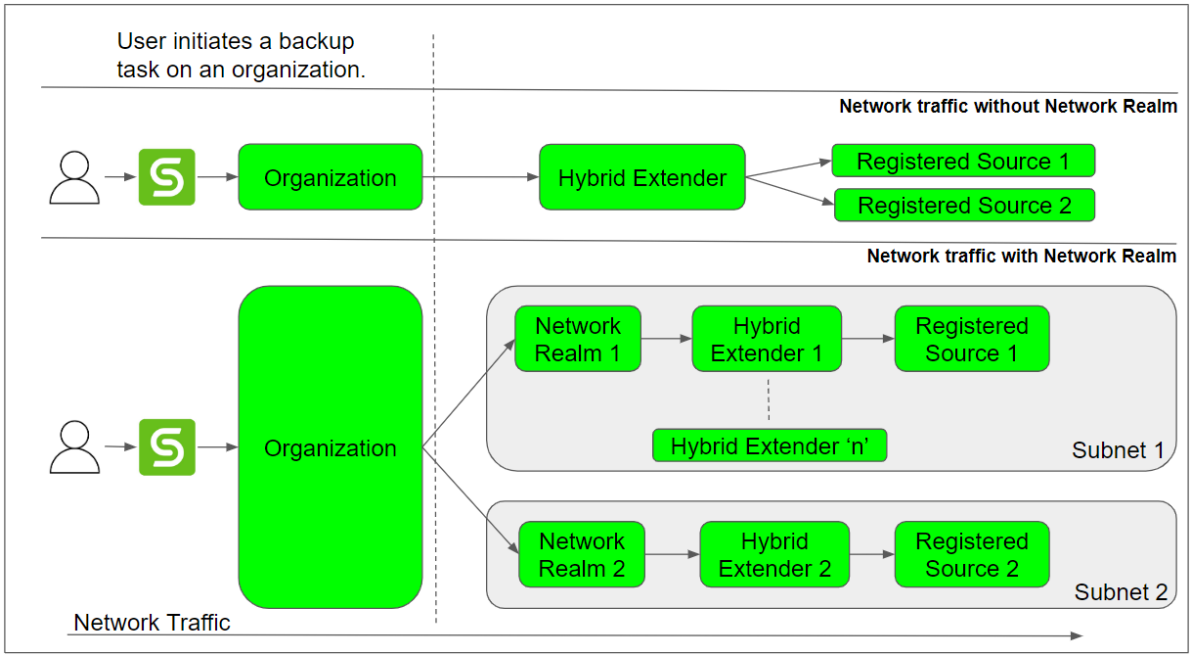

Technical Support for Hybrid Extender VM

Hybrid Extender now has a support channel that can be enabled to let Cohesity Technical Support engineers access it for troubleshooting. The login credentials for the Hybrid Extender will be needed for accessing it. The support channel can be enabled or disabled as an Organization admin or service provider admin impersonating the tenant. Contact the ***Cohesity Technical Support*** team for assistance.

Data Source Connection

In a Service Provider (SP) multi-tenant environment, tenants often operate in isolated networks and may use overlapping IP address ranges and subnets. When multiple

organizations use identical IP ranges on the same Cohesity cluster, the cluster cannot uniquely identify tenant workloads based on IP address alone.

To address this, Cohesity provides Data Source Connections.

A Data Source Connection consists of one or more Data Source Connectors (virtual machines) that act as secure data movers between a tenant organization and the Cohesity cluster. These connectors enable backup and recovery of workloads that are not directly reachable from the cluster or that exist in overlapping IP environments.

In earlier releases, this functionality was provided by the Hybrid Extender (HyX). Starting with release 7.4, you can also use Data Source Connections.

Why Use Data Source Connection

Data Source Connections provide significant improvements in manageability, scalability, security, and user experience for both Service Providers and their tenants.

Following are the key benefits of Data Source Connection:

- **Improved Lifecycle Management**: Data Source Connectors can be securely claimed using a token-based workflow. They support in-place upgrades and controlled patching, eliminating the need for redeployment during version changes and reducing operational overhead.
- **Connection-Based Traffic Segmentation**: Data Source Connections let you group connectors logically and configure VIPs per connection. This gives you better control over how backup traffic flows, making it easier to manage routing, balance load across nodes, and isolate network traffic when needed.
- **Improved Monitoring and Visibility**: Data Source Connections provides visibility into connector health, CPU/memory, along with proactive alerts.

How Data Source Connection Works

In a multi-tenant deployment, tenant organizations may independently control IP addressing for their virtualization hosts, virtual machines, physical servers, and other infrastructure components (collectively referred to as “tenant entities”).

Because tenants operate in isolated private networks, overlapping IP addresses are common. For example:

- Tenant A may have a SQL Server VM with IP address 10.0.0.15
- Tenant B may also have a SQL Server VM with IP address 10.0.0.15

Without network isolation, the Cohesity cluster cannot uniquely identify these workloads using IP address alone.

Data Source Connections resolve this by logically isolating traffic per organization. Workloads are uniquely identified using a combination of:

- Organization ID
- Workload IP address

Each Data Source Connector acts as a proxy VM between the tenant environment and the Cohesity cluster. The connector:

- Handles outbound traffic from the cluster to protected sources
- Handles inbound backup ingest traffic (for example, VMware backups) from the tenant environment to the cluster
- Supports source-side deduplication
- Ensures tenant-level network isolation

Secure Communication Model

Communication between the Data Source Connector and the Cohesity cluster uses:

- Secure gRPC channels
- Mutual TLS (mTLS) encryption
- Certificate-based authentication

All data and control traffic between the connector and cluster is encrypted in transit. The connector maintains secure, outbound connections to the cluster, minimizing firewall and inbound exposure requirements within tenant environments.

Data Source Connector Requirements

Before deploying the Data Source Connector, review and understand the following requirements.

Supported Sources

You can protect the following sources in an organization using the Data Source Connection:

- VMware vCenter
- Physical Servers (Windows – File and Block Based)
- Physical Servers (Linux – File and Block Based)
- MS SQL VM, FCI, AG (Volume-based, and File-based, Native backup)
- MS SQL Native
- Oracle Standalone
- Oracle RAC

Required User Role

Service Provider Administrator

The Service Provider (SP) Administrator role is required to:

- Enable Data Source Connections for Organizations in a multitenant cluster
- Upgrade Data Source Connections for Organizations in a multitenant cluster

This role ensures centralized control over connector enablement and lifecycle management across tenant organizations.

Organization Administrator or Organization User

An Organization Administrator or a user with the following privileges is responsible for managing Data Source Connections within their own Organization.

- `Modify data-source connections`
- `View data-source connections`
- `Modify data-source connectors`
- `View data-source connectors`

This role is required to:

- Deploy Data Source Connectors for their Organization
- Manage Data Source Connections

Supported Data Source Connection

Cohesity supports the following Data Source Connection:

- VMware Data Source Connection
- Hyper-V Data Source Connection

System Prerequisites

Ensure that the Data Source Connector virtual machine (VM) meets the following system requirements before deployment:

- **Compute Requirements**
  - 4 vCPUs
  - 10 GB RAM
- **Disk Requirements**
  - The required disk size depends on the hypervisor platform:
  - 171 GB for VMware deployments
  - 166.6 GB for Hyper-V deployments
- **Network Requirements**
  - Outbound Internet connectivity is required for secure communication, updates, and patch downloads.
  - Data Source Connectors require IPv4 network connectivity to communicate with the Cohesity cluster. IPv6-only network environments are not supported for Data Source Connector deployments.

Data Source Connector Sizing Recommendations

We recommend that you have one Data Source Connector for each 160 VMs or 16 TB of source data. If you have more data, we recommend that you stagger their first full backups.

> **Note:** These requirements are subject to change.

Check Firewall Ports

A Data Source Connector communicates with both:

- The cluster
- The protected data sources (for example, VMware vCenter, physical servers, and other workloads)

To ensure proper connectivity, the required firewall ports must be open between the Data Source Connector, the cluster, and the relevant data sources.

These firewall rules allow outgoing traffic from a Data Source Connector to the Cohesity cluster endpoint. The Data Source Connector opens a secure encrypted gRPC tunnel to the endpoint and uses it for both backup and recovery traffic.

For more information, see Data Source Connector.

Configure Data Source Connections

To configure Data Source Connections on a source you want to protect:

1. Enable Data Source Connection and download the installer for the organization.
2. Deploy Data Source Connection on the Organization cluster.
3. Register source on the Organization cluster with Data Source associated.

Enable Data Source Connections and Download the Installer

> **Important:** This action can be performed only by the Service Provider administrator.

To deploy Data Source Connections, the service provider administrator or cluster administrator must:

- Enable the Data Source Connection for the Organization
- Download the Data Source Connector Installer and make it available for the organization admin .

To enable the Data Source Connections, complete the following steps:

1. Login to the cluster as a Service Provider administrator.
2. Navigate to **Settings > Access Management**.
3. In the **Organizations** tab, click for required organization, and select **Edit**.
4. In the **Network Segments** section of the Edit Organization page, enable **Data Source Connections**.
5. Click the **Download Data Source Connection** drop-down to download the Data Source Connector Installer:

   > **Note:** If you are redirected to the Cohesity downloads portal, download the installer (OVA/VHD) and share it with the organization administrator using your standard software distribution method.

   - **VMware**
   - **Hyper-V**
6. Click **Save**.

Deploy Data Source Connection

> **Important:** This action can be performed only by the Organization administrator.

After the Service Provider administrator enables Data Source Connections and make it available for the for the Organization to download the installer, the Organization Administrator can deploy the appropriate Data Source Connector based on their data center environment.

> **Tip:** The organization administrator needs access to the downloaded installer file or URL. Make sure the installer is reachable from the tenant environment before you proceed.

Depending on the hypervisor platform, deploy one of the following:

- Deploy VMware Data Source Connectors
- Deploy Hyper-V Data Source Connectors

Register Source

After deploying the Data Source Connection, you can register the source. During registration, associate the source with the appropriate Data Source Connection.

Deploy VMware Data Source Connectors

The organization administrator can deploy the VMware Data Source Connection using the installer OVA. The installer OVA can be obtained either from the service provider administrator or from the Cohesity Downloads Portal.

Deploy the OVA in your VMware environment on:

- A vCenter Server, or
- A Standalone ESXi host

The selected host must:

- Have network access to your protected data sources.
- Meet the system requirements.
- Meet the required firewall and port configurations.

After deployment, each Data Source Connector runs as a virtual machine within your data center.

Best Practice Recommendation

For improved performance and redundancy, Cohesity recommends deploying at least two Data Source Connectors for each Data Source Connection in your data center.

Cohesity can concurrently back up individual VM disks using multiple Data Source Connectors, leading to faster backups. Cohesity determines the concurrency based on various factors, such as disk size and the workload of the Data Source Connectors.

All the data that a Data Source Connection handles, from your sources to the cluster, is encrypted in flight and at rest.

Create VMware Data Source Connection

To create a VMware Data Source Connection, you must perform the following steps:

1. Deploy Data Source Connector OVA.
2. Obtain the Connection Token from the Cohesity cluster UI.
3. Authenticate the Data Source Connector with Cohesity using the Connection Token.
4. Verify the Connection.

Deploy Data Source Connector OVA

To deploy the Data Source Connector OVA in your data center:

1. Log in to your vCenter host.
2. From the **Hosts and Clusters** tab in the vSphere Web Client, right-click on any cluster that can host your VM and select **Deploy OVF Template**. The **Deploy OVF**

Template wizard opens.

3. On the **Select an OVF template** page, select **Local file**, click **UPLOAD FILES**, and browse to the location of the Data Source Connector OVA file you downloaded .
4. Clikc **Next**.
5. On the **Select a name and folder** page, enter the following and click **Next**:
    1. In the **Virtual machine name** field, enter a unique name for your Data Source Connector.
    2. In the **Select a location for the Virtual Machine** field, select where your VM should reside from the displayed list of inventory locations.
6. On the **Compute Resources** page, select a compute resource for the Data Source Connector VM and click **Next**.
7. On the **Review details** page, verify the Data Source Connector information and click **Next**.
8. On the **Configuration** page, verify **SAAS-CONNECTOR** is selected and click **Next**.
9. On the **Select storage** page, select a datastore with at least 181 GB of free disk space and click **Next**.
10. On the **Select networks** page, select a destination network and click **Next**. You can select VLANs from both the **DataNetwork** and the **SecondaryNetwork** fields. The Data Network is used for communication with Cohesity Data Source Connector, and the Secondary Network is used for communication with your data sources. Based on your requirements:
    - To deploy the Data Source Connector on a single network, select the same VLAN in both **DataNetwork** and **SecondaryNetwork**.
    - To deploy the Data Source Connector on a dual network, select different VLANs in **DataNetwork** and **SecondaryNetwork**, respectively.

    > **Note:**
    > - The Data Source Connector must have dual IP addresses if your data sources are in a private non-routable VLAN.
    > - Once you have deployed the Data Source Connector on a single network, you cannot modify the Data Source Connector to use a dual network or vice versa.
11. On the **Customize template** page, enter the network settings: **Network IP Address**, **Network Netmask**, and **Default Gateway**. If you have selected a different VLAN for the secondary network, enter the **Network IP Address**, **Network Netmask**, and **Default Gateway** for the secondary network, as well. Click **Next**.

> **Note:**
> - To set the network settings using static IP addresses, manually enter the details in the respective fields for both **DataNetwork** and **SecondaryNetwork**.
> - To set the network settings using DHCP, leave the fields blank in both the **DataNetwork** and **SecondaryNetwork** sections.
> - **Data Network** and **Secondary Network** must be configured using the same network configuration method. That is static IP addresses or DHCP.

12. Review the summary on the **Ready to complete** page and click **Finish**.
13. Once the VM is created, power it on.

After it boots, the services in the Data Source Connector VM (including the UI) can take 4-5 minutes to start.

Obtain the Connection Token

To obtain the connection token required for authenticating the Data Source Connection with the Cohesity cluster:

1. Log in to the Cohesity cluster UI as an Organization administrator, navigate to **Systems > Data Source Connections**.
2. Click **New Connection**.

   Th Create Data Source Connection dialog appears.
3. From the **Destination Cloud Region for this Connection** drop-down, select a region for your data backups.
4. From the **Deployment Platform** drop-down, select **VMware**.
5. Copy the **Connection token**.
6. Click **Skip Verification** or **Cancel**. You will need to verify the connection later once Data Source Connector authenticate to the Cohesity cluster.

Authenticate the Data Source Connector with Cohesity

1. Log in to the Data Source Connector by entering the IP address of the Data Source Connector VM in the address bar of your browser and click **Enter**.
2. On the Data Source Connector's User Interface, enter "admin" in the **Username** and **Password** fields to log into the Data Source Connector.

3. On the next screen, you are prompted to change your password. Change your default password and log in again with your new password.
4. Verify the network configuration settings, make necessary changes, and click **Continue**.
5. On the **Data Source Connector** page, paste the Connection token in the **Connection Claim Token** field and click **Save**.

It can take another few minutes for the Data Source Connector to authenticate to the Cohesity cluster.

Verify the Connection

Once the Data Source Connector authenticates successfully, return to the **Create New Connection** dialog in the Cohesity cluster UI, and then click **Verify Connection**.

Deploy Hyper-V Data Source Connectors

The organization administrator can deploy the Hyper-V Data Source Connection using the template VHD. The template VHD can be obtained either from the service provider administrator or from the Cohesity Downloads Portal.

Deploy the VHD in your Hyper-V data center environment that:

- Has network access to the data sources you want to protect
- Meets the required system specifications
- Meets the required firewall and port configurations

Once deployed, each Data Source Connector is a virtual machine that runs on a Hyper-V host in your data center.

Best Practice Recommendation

For improved performance and redundancy, Cohesity recommends deploying at least two Data Source Connectors for each Data Source Connection in your data center.

Create Hyper-V Data Source Connection

To create a Hyper-V Data Source Connection, you must perfrom the following steps:

1. Deploy Data Source Connector VHD.
2. Obtain the Connection Token from the Cohesity cluster UI.
3. Authenticate the Data Source Connector with Cohesity using the Connection Token.
4. Verify the Connection.

Deploy Data Source Connector VHD

To deploy the Data Source Connector VHD in your data center:

1. Log in to your SCVMM server or Standalone Hyper-V host.
2. Download the VHD file to the SCVMM server or Standalone Hyper-V host using the VHD URL.
3. From the Hyper-V Manager, open the **New (Create) Virtual Machine** wizard. (For detailed instructions, see Create a virtual machine in Hyper-V in the Microsoft documentation.)
    1. **Configure** the name, location, generation, hardware for the VM.
    2. **Select Virtual Hard Disk**. Select 'Use existing virtual hard disk' and choose the downloaded VHD file.
    3. **Configure Networking**. Select an operational Virtual Switch to connect the VM to.
    4. **Review**. Review the configuration from the Summary section and click Finish.

    After it boots, the services in the Data Source Connector VM (including the UI) can take 4-5 minutes to start.

Obtain the Connection Token

To obtain the connection token required for authenticating the Data Source Connection with the Cohesity cluster:

1. Log in to the Cohesity cluster UI as an Organization administrator, navigate to **Systems > Data Source Connections**.
2. Click **New Connection**.

    Th Create Data Source Connection dialog appears.
3. From the **Destination Cloud Region for this Connection** drop-down, select a region for your data backups.
4. From the **Deployment Platform** drop-down, select **Hyper-V**.
5. Copy the **Connection token**.
6. Click **Skip Verification** or **Cancel**. You will need to verify the connection later once Data Source Connector authenticate to the Cohesity cluster.

Authenticate the Data Source Connector with Cohesity

1. Log in to the Data Source Connector by entering the IP address of the Data Source Connector VM in the address bar of your browser and click **Enter**.
2. On the Data Source Connector's User Interface, enter "admin" in the **Username** and **Password** fields to log into the Data Source Connector.

3. On the next screen, you are prompted to change your password. Change your default password and log in again with your new password.
4. Verify the network configuration settings, make necessary changes, and click **Continue**.
5. On the **Data Source Connector** page, paste the Connection token in the **Connection Claim Token** field and click **Save**.

It can take another few minutes for the Data Source Connector to authenticate to the Cohesity cluster.

Verify the Connection

Once the Data Source Connector authenticates successfully, return to the **Create New Connection** dialog in the Cohesity cluster UI, and then click **Verify Connection**.

View Data Source Connections

> **Important:** This action can be performed only by the Organization administrator.

You can navigate to the Data Source Connections page on the cluster UI to:

- View the list of Data Source Connections deployed and their details.
- Manually deploy a new Data Source Connection.

To view the Data Source Connections page, in the Cohesity UI, navigate **System > Data Source Connections.** On the Data Source Connections page, you can view:

- Glance Bar
- Data Source Connections Details Table

Glance Bar

The glance bar on the Data Source Connections page provides a summary of the following information:

- The total number of Healthy Data Source Connections.
- The total number of Unhealthy Data Source Connections.
- The total number of Unused Data Source Connections.
- The total percentage of CPU utilized by the Data Source Connections.
- The total percentage of memory utilized by the Data Source Connections.

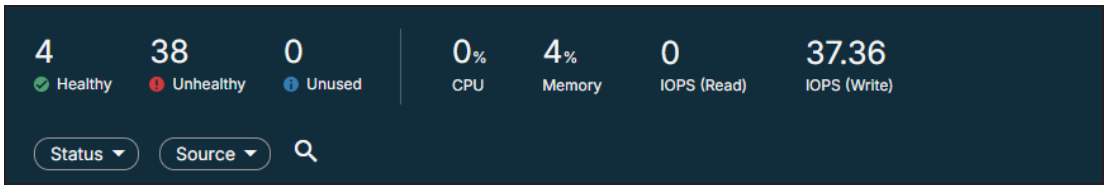

Data Source Connection Details Table

The table below describes the data displayed in the Data Source Connection table:

You can use the:

- Search bar to filter Data Source Connections by name.
- **Status** filter option to filter the Data Source Connection based on:
  - **Healthy**: Displays only the healthy Data Source Connections.
  - **Unhealthy**: Displays only the unhealthy Data Source Connections.
  - **Unused**: Displays only the unused Data Source Connections.
- Source filter option to filter the Data Source Connection based on the Source it is deployed on.

> **Note:** The hybrid extenders that were deployed in the previous versions are displayed by the name, `<default_realm_xxx>`.

Column Name	Description
Name	The name of the Data Source Connection.
Health Status	The status of the Data Source Connection: • Healthy • Unhealthy • Unused
Sources	The number of sources that use the Data Source Connection.
Connectors	Number of Data Source connectors in the Data Source Connection.
Last Connection Time	The timestamp when the Data Source Connection was last found to be healthy.
Connected Since Time	The timestamp from when the Data Source Connection was continuously connected.

View Data Source Connection Details

To view the list of Data Source Connectors of a Data Source Connection, their status, resource matrices, and the sources utilizing the Data Source Connection, select the Data Source Connection from the Data Source Connection details table.

The Data Source Connection details page contains the following tabs:

- Connection Details Tab
- Linked Sources Tab

Connection Details Tab

The **Connection Details** tab of a Data Source Connection displays the following:

- A graphical representation of how the Data Source Connection acts as the data mover between the data sources and the cluster.
- The number of Data Source Connectors in the Data Source Connection.
- The percentage of CPU and memory utilized.
- Connector Details Table.

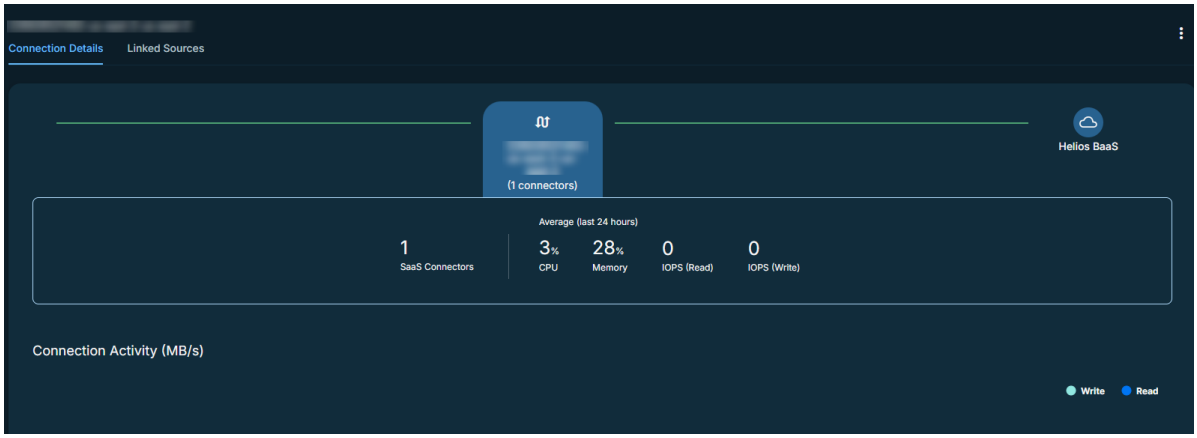

Connector Details Table

The table below describes the data displayed in the Data Source Connector details table:

Column Name	Description
Connector	The list of Connectors in the Data Source Connection.

Column Name	Description
Status	The status of the Data Source Connector: • **Healthy**: The cluster is connected with the Data Source Connector.. • **Unhealthy**: The cluster is not connected with the Data Source Connector.
Release	Data Source Connector version.
CPU	The percentage of CPU utilized by the Data Source Connector.
Memory	The percentage of memory utilized by the Data Source Connector.
IOPS (Read)	The average number of read operation per second on the Data Source Connector.
IOPS (Write)	The average number of write operation per second on the Data Source Connector.
Last Connection Time	The timestamp when the Connector was last found to be healthy.
Connected Since Time	The timestamp from when the connection was continuously healthy.

View Data Source Connector Details

To view the resource utilization of a Data Source Connector, as well as the status, ports, and connection timestamps of its connected endpoints, select the desired Data Source Connector from the Data Source Connector Details table.

In the Data Source Connector Details page, you can view:

- A glance bar displaying the percentage of CPU and memory utilized, and the read and write I/O operations.
- Graphical representation of the percentage of CPU and memory utilized, and the read and write I/O operations.

Linked Sources Tab

The **Linked Sources** tab displays the sources that use the selected Data Source Connection to protect the source data.

Manage Data Source Connectors

> **Important:** This action can be performed only by the Organization administrator.

To optimize performance, we recommend that you use at least two Data Source Connectors in each Data Source Connection you create, and that you have one Data Source Connector for each 160 VMs or 16 TB of source data. (If you have more VMs, we recommend that you stagger their first full backups.)

Cohesity clusters can concurrently back up individual VM disks using multiple Data Source Connectors, leading to faster backups. Cohesity clusters determine the concurrency based on various factors, such as disk size and the workload of the Data Source Connectors. Cohesity clusters cannot recover individual VM disks using multiple Data Source Connectors because VMware and HyperV do not support it.

Add Data Source Connector

To add a Data Source Connector to an existing user-deployed Data Source Connection:

1. In the product UI, navigate to **Systems > Data Source Connections**.
2. Click the Actions menu (⋮) next to the Data Source Connection and select **Download Installer** to save the OVA to your data center.
3. To deploy the OVA or VHD. For more information, see Configure Data Source Connections.
4. Back in the **Data Source Connections** page, click the Actions menu (⋮) next to the Data Source Connection again and select **Connection Token**. In the dialogue, click the **Copy to Clipboard** button.
5. Browse to the Data Source Connector IP and log in as admin/admin. On initial login, change the default password and log in again with your new password. Enter the **Connection Token** and common configuration settings and click **Save**.
6. Once the Data Source Connector authenticates successfully to the Cohesity cluster, click the **Expand (v)** button next to the Data Source Connection to confirm that the new Data Source Connector is listed.

To add more Data Source Connectors to the same Data Source Connection, repeat the steps above.

Remove Data Source Connector

To remove a Data Source Connector from one of your Data Source Connections:

1. In the product UI, navigate to **Systems > Data Source Connections**.
2. Select the Data Source Connection.
3. Click the Actions menu (⋮) next to the Data Source Connector and select **Remove from Connection**.

The Data Source Connector is removed from the Data Source Connection. If other healthy Connectors remain in the Data Source Connection, it will continue to function over those Connectors.

Modify Data Source Connector Network Settings

After you have deployed your Data Source Connector, you can use the Data Source Connector UI to:

- Modify the Data Source Connector's existing network configuration.
- Change the network settings of the Data Source Connector from manual (static IP addresses) to DHCP, or from DHCP to manual configuration.

To modify the Data Source Connector network settings:

1. Using a browser, type the IP address of the Data Source Connector VM and log in to the Data Source Connector UI.
2. On the Data Source Connector UI:
    1. Select **DHCP** to change the network configuration from static IP addresses to DHCP, then click **Continue**.
    2. Select **Configure Manually** to change the network configuration from DHCP to static IP addresses, then click **Continue**.
3. The Data Source Connector verifies the network connection to the Cohesity cluster using the primary network. Click **Close** if the verification is successful or if errors are found during the verification process. You can also click **Continue with Errors** to ignore the errors and continue using the Data Source Connector.

Enable Data Source Connection Support Channel Access

You can enable the Support Channel for Data Source Connector when you need assistance from Cohesity. When enabled, Support Channel access applies to all Data Source connectors under that connection.

> **Important:** This action can be performed only by the Organization administrator.

To enable the Data Source Connection support channel:

1. Login to the Data Source Connector UI (*https://<saas-connector-ip>/*) using the credentials you have set while creating the Data Source Connector.
2. Click the user icon on the top-right corner of the Data Source Connector UI, and then click **Support Access**.
3. On the support Access page, enable the **Support Channel** option to enable the support channel access.

4. From the channel active until field, you can select the date till you want the support channel to be active.
5. Copy the **Support Access Token** and **Cluster ID** from the UI and send these details to Cohesity Support.

   Cohesity will enable support access for the Data Source Connector.

Migrate Hybrid Extender to Data Source Connection

If you have already installed Hybrid Extenders in the previous versions to backup the data of a source from the organization environment to the Cohesity cluster, then you can migrate the hybrid extender to use the new Data Source Connectors.

Edit the Source Registered

As an Organization administrator, perform the following task on the cluster:

> **Note:**
> Before proceeding, ensure:
>
> - Data Source Connection is enabled by the Service Provider administrator for the cluster.
> - Data Source Connection is configured on the cluster by the organization administrator.

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu next to the source and select **Edit**.
3. In the **Edit <source_name> Source** page, and from the SaaS Connection drop-down, select the Data Source Connection you want to use. You can also create a new Data Source Connection by clicking **Create Data Source Connection** from the SaaS Connection drop-down.

Disable Hybrid Extender

Before you disable the Hybrid Extender, ensure that no protection or recovery job is running. After you disable the Hybrid Extender, the tenant network falls back to VLAN-based multitenancy. For information about assigning VLANs, see Configure Network Settings for an Organization.

To disable the Hybrid Extender, complete the following steps:

1. In the Cohesity cluster UI, select **Settings > Access Management**.
2. In the **Organizations** tab, click for required organization, and select **Edit**.
3. In the **Network Segments** section of the **Edit Organization** page, disable **Data Source Connections**.
4. Click **Save**.

Upgrade and Patch Workflows for Data Source Connectors

> **Important:** This action can be performed only by the Service Provider administrator.

As a service provider administrator, you can upgrade, apply patch, or upgrade and apply patch to the Data Source Connectors across all organizations.

Download Package

Before performing an upgrade or applying patch to the Data Source Connectors, you must download the package for upgrade, patch, or both. To download an upgrade or patch package, you can select the desired package type from either a URL or your local machine. Once the package is chosen, you can proceed with downloading, or opt to apply the upgrade or patch immediately.

Perform the following steps to download the package:

1. Log in to the Cohesity cluster UI as a service provider administrator.
2. Go to **Settings** > **Software Update**.
3. In the **Available Connector Packages** section, click **Get New Connector Package**.

   A dialog appears with the following package options:

   1. **Upgrade Package**
   2. **Patch Package**
   3. **Upgrade and Patch Packages**
4. Select the required package and click **Continue**.
   A **Get New Upgrade Package** dialog appears.
5. Choose one of the following options to download the package.

Option	Field	Description
From URL	**Enter the download URL**	Provide the package's URL. You can obtain the package from the Cohesity downloads portal.
	Repository credentials for package (Optional)	Repository credentials can be used to download packages from secure repositories that require authentication to access the files. Choose one of the following options from the drop-down: • **Basic** - Select this if your repository requires HTTP basic authentication. Enter the **Username** and **Password**. • **Custom** - Select this if your repository requires custom authentication schemes and requires you to pass credentials as HTTP headers. Enter **Key** and **Value**. Key is the HTTP header name and Value is the HTTP header value.
From Local Machine	**Select File**	Browse and select the downloaded package from your local machine.

> **Note:**
> When you select the **Upgrade and Patch Package** option, you can either specify the download URLs of both upgrade and patch packages or you can upload the downloaded packages from your local machine to Cohesity cluster.

6. Depending on your package selection, you can click:
    1. **Get Only** to download the package.
    2. **Get and Upgrade** to download and immediately upgrade the package.
7. Once the package is downloaded, you can perform the following actions by clicking on the package:
    1. **Upload CRL for Signed Package** - To upload Certificate Revocation List (CRL) for signed packages. CRLs are used to validate whether the certificate used for signing the package has been revoked after the package was downloaded to the Cohesity cluster.
    2. **View Checksum** - To view and verify the checksum of the package.
    3. **Delete** - To delete the package.

Run Pre-check

The pre-check option is available after the package download. It is important to run the pre-checks once the packages have been downloaded. Doing so in advance allows you to identify and resolve any issues, ensuring a smooth upgrade and patching process without any interruptions.

To run the pre-check on the downloaded package, click **Run**.

The pre-check will start and you can see the pre-check details by clicking **View Details**.

Upgrade Data Source Connectors

To upgrade the Data Source Connectors, you can select the upgrade package and initiate the upgrade process after performing the pre-check. You can track the upgrade progress for all connectors until the operation is completed.

To upgrade the Data Source Connector:

1. On the **Software Update** page, click **Upgrade**.

   The **Upgrade Review** dialog appears.
2. After reviewing the upgrade information, click **Start Upgrade**.

   The **Ok to Upgrade?** dialog appears.
3. If the pre-check fails, you can either:
    1. **Stop immediately** to stop the upgrade process (recommended). This allows you to review and address the pre-check failures before retrying the upgrade.
    2. **Proceed with the upgrade by ignoring pre-check failure** to continue the upgrade process despite the pre-check failure.
4. Enter **Yes** in the text field and click **Upgrade** to confirm and start upgrading.

The upgrading process begins. You can observe the upgrading progress for all the available connectors.

Apply Patch

To apply a patch, select the patch package and begin the process after performing the pre-check. The patch will be applied across all connectors, with progress tracked until completion.

Perform the following steps to apply a patch:

1. On the **Software Update** page, click **Apply**.

   The **OK to apply?** dialog appears.
2. If the pre-check fails, you can either:
    1. **Stop immediately** to stop the patching process (recommended). This allows you to review and address the pre-check failures before retrying the patch.

    2. **Proceed with the patch by ignoring pre-check failure** to continue applying the patch despite the pre-check failure.
3. Enter **Yes** in the text field and click **Apply** to confirm and start applying the patch.

You can observe the patching progress for all the available Data Source Connectors. The patching process begins after the pre-checks are completed on all connectors.

To cancel the patching, click **Cancel**.

The **Cancel Patching?** dialog appears.

Enter **Cancel Patching** in the text field and click **Yes** to cancel the patching or **No** to continue the patching.

Upgrade with Patch

Upgrading with a patch allows you to apply both an upgrade and patch in one action. After downloading the upgrade and patch packages, you can proceed with the upgrade, including the patch, following a pre-check process.

Perform the following steps to upgrade with the available patch:

1. On the **Software Update** page, click **Upgrade**.

   The **Upgrade with the Available Patch?** dialog appears.
2. Click **Yes** to upgrade the Data Source Connectors with the available patch.

   The **Upgrade Review** dialog appears
3. After reviewing the upgrade information, click **Start Upgrade**.

   The **Ok to Upgrade?** dialog appears.
4. If the pre-check fails, you can either:
    1. **Stop immediately** to stop the upgrade process (recommended). This allows you to review and address the pre-check failures before retrying the upgrade.
    2. **Proceed with the upgrade by ignoring pre-check failure** to continue the upgrade process despite the pre-check failure.
5. Enter **Yes** in the text field and click **Upgrade** to confirm and start upgrading.

The upgrading process begins. You can observe the upgrading progress for all the available Data Source Connectors.

To monitor upgrading progress, click **View Details**.

Access Management for Organizations

Assign an Authentication Server

Assigning an authentication server enables Organization users to access the Organization securely.

You can assign one of the following authentication server types:

- Active Directory
- LDAP
- Keystone

Assign Active Directory

An entire AD domain can be dedicated to each Organization. Or An AD domain can be registered on the cluster by the Cohesity administrator user (in which case no AD domain is added to the Organization) and individual users/groups from that AD domain may be added to each Organization.

To assign an Active Directory, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization, and select **Edit**.
3. In the **Edit Organization** page, click the **Assign AD/LDAP/Keystone** link.
4. In the Assign **AD/LDAP** page, click **Active Directory**.
5. Select one or more Active Directories from the list.
6. Click **Save**.

Assign LDAP

To assign an LDAP server to an Organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Organizations**.
2. In the **Organizations** tab, click ⋮ for the required Organization, and select **Edit**.
3. In the **Edit Organization** page, click the **Assign AD/LDAP/Keystone** link.
4. In the **Assign AD/LDAP** page, click **LDAP**.
5. Select one or more **LDAP** providers from the list or create a new one.

> **Note:**
> - You can map the LDAP providers to the Active Directory domain so users can access the same files on the Cohesity cluster from both Windows and Linux systems.
> - An AD/LDAP that is assigned to an Organization cannot be assigned to another Organization.

6. Click **Save**.

Assign a Keystone Config

After you create an Organization on the Cohesity cluster, you must map a Keystone config to use the Cohesity Swift service in the Organization. For more information, see OpenStack Swift Storage Protocol.

To assign a Keystone config to an Organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Organizations**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required organization, and select **Edit**.
3. In the **Edit Organization** page, click the **Assign AD/LDAP/Keystone** link.
4. In the **Assigning AD/LDAP** page, the **Keystone** tab displays the registered Keystone servers. Select one Keystone server and click **Assign**.
5. In the **Swift Roles** section of the **Edit Organization** page, click the pencil icon.
6. In **the Edit Default Swift** options, add or modify the required Swift roles and click **Save**

> **Note:** You can only add the roles defined on the OpenStack server.

7. (Optional)To further isolate network traffic for the organization, assign VLAN IDs.

> **Note:**
> - You can add or remove VLANs later.
> - All the swift requests whose tenancy cannot be determined will fail.

8. Click **Save**.

Single Sign-On for Organizations

An Organization in a multitenancy cluster can register an IDP and use it for single sign-on authorization within the Organization.

For information on SSO, see Single Sign-On (SSO).

You can configure Single Sign-On at two different levels on a Cohesity cluster with the Organizations feature is enabled:

- Single Sign-On at a Service Provider level.
- Single Sign-On at an Organization level.

To successfully sign in to the Cohesity cluster, after successful validation at IDP, at least one of the following must be true:

1. The user is registered as an SSO user on the Cohesity cluster.
2. The user is part of an SSO group that is registered on the Cohesity cluster.
3. A default role is available in the SSO configuration.
    1. If a default role is added to the SSO, then the user needn't be added to the cluster and all users will have default access to the cluster using the default role.
    2. If no default role is added to the cluster, then the user has to be added to the cluster with a particular role, and then only the user login to the cluster.

Single Sign-On at a Service Provider Level

When a Service Provider Administrator registers an SSO (IDP), the SSO is configured at the cluster level and it is not assigned to a specific Organization.

The following are a few points you should consider for Service Provider-level SSO:

- The Service Provider Administrator can create new SSO users and groups and assign roles in multiple Organizations.
- The Service Provider Administrator can configure multiple SSOs (IDPs) on the cluster. However, one SSO per tenant and one per service provider admin.
- The users and groups that are part of the SSO (IDP) configured on the cluster are authenticated through the same SSO (IDP).
- If a default role is assigned to the SSO, the role is applicable to all the SSO users and groups that are part of the SSO.
- If no default role is assigned to the SSO, the role that is assigned to the individual group or user is applicable.
- An SSO group can be assigned to multiple Organizations. If an SSO user is part of a group and that group is part of multiple Organizations, the user can access other Organizations by using the Switch Organization option in the Cohesity Dashboard and then manage more than one Organization based on the roles assigned to the group.

- A user who is a part of a group and the group is assigned to multiple Organizations can have different roles inside different Organizations.

For example, In the following configuration:

- **IDP**: Okta is configured at the Cohesity cluster level.
- **SSO Groups**: Group 1 with Admin role and Group 2 with Operator role are created on the Cohesity cluster and the same groups are present on the Okta IDP.
- **SSO User on Okta**: User1 is part of both Group 1 and Group 2 on the Okta IDP.
- **Organizations**: Org 1, Org 2, and Org 3 are created on the Cohesity cluster.
- Group 1 with the Admin role is assigned to Org 1 and Org 2.
- Group 2 with the Operator role is assigned to Org 3.

Now, when the User1 accesses the Cohesity cluster, the user can switch among Org1, Org2, and Org3. User1 inherits the role assigned to the group.

In this scenario, the User1 plays the role of the Admin for Org1 and Org2 and for Org 3 as an Operator.

- All the existing SSO users in the SSO of an Organization can be quickly onboard to Cohesity by configuring the SSO and registering groups on the Cohesity cluster.
- An Organization Admin can switch among multiple Organizations. The Service provider/tenant admin can only see the entities that belong to a logged-in Organization.

Single Sign-On at an Organization Level

When an Organization Administrator registers an SSO (IDP), the SSO is configured at the Organization level and is applicable only to that Organization.

The following are a few points you should consider for the Organization-level SSO:

- To manage the SSO configured at the Organization level, you must sign in to the Organization.
- When you sign in to the Organization with SSO configured at the Organization level, you are authenticated through the SSO installed on the Organization and you cannot switch to any other Organization.
- As an Organization user or administrator, depending on the role assigned, you can create new SSO users and groups, and assign roles to them within the Organization.
- All the existing SSO users in the SSO of an Organization can be quickly onboard to Cohesity by configuring the SSO and registering groups on the Cohesity cluster.

Configure SSO at an Organization Level

As an Organization Admin, you can sign in to the Organization and register an SSO specific to your Organization.

Prerequisites

- As the Organization admin, you must have the Admin role assigned or you should be part of an SSO group with the Admin role assigned.
- Add the Cohesity cluster as an application on your IDP.
- Get the SSO URL and certificate file.

To configure SSO at an Organization level, complete the following steps:

1. Access the Cohesity Dashboard.
2. On the **Sign in** page, click **Sign in with SSO** and click **Sign In**.
3. Slide the **Sign in to Organization** toggle button.
4. Type your Organization ID and click **Sign In**.
5. On the IDP login page, provide your username and password.
6. Click **Sign in**.
7. After signing in, click **Settings > Access Management**.
8. Click the **SSO** tab and click the **Configure SSO** button on the top-right of the screen.
9. Complete the SSO configuration steps.

Switch Organizations as an Organization Admin or User

If your role requires you to manage multiple Organizations, you can switch among multiple Organizations to perform Organization-specific operations.

Prerequisites

- Access the cluster as a Cohesity cluster admin and configure SSO (IDP) on the cluster.
- Create a group and add the required users to the group on your IDP. This allows users to become part of multiple Organizations.
- Add SSO group(s) on the cluster the same as the IDP group.
- Assign groups to multiple Organizations. For more information, see the Assign an SSO Group to Organizations.
- Add the Cohesity cluster as an application on your IDP

To switch to another Organization, complete the following steps:

1. Access the Cohesity Dashboard.
2. On the **Sign in** page, click **Sign in with SSO** and click **Sign In**.
3. On the IDP login page, provide your IDP username and password.
4. Click **Sign in**.
5. In the Cohesity Dashboard, click the user icon in the top-right corner and select **Switch**.
6. In the **Switch** window, select a required Organization from the list.

   Click **Switch**.

Assign an SSO Group to Organizations

Currently, a Cohesity user can be part of only one Organization. However, you can assign an SSO group to multiple Organizations. When you assign the group to multiple Organizations, a user that's part of the group can switch among the Organizations. However, the user inherits the role assigned to the group.

> **Note:** After assigning a group to multiple Organizations, users within that group can sign in to the Cohesity Dashboard and switch among the Organizations. All users have the same privileges that are assigned to the group.

To assign an SSO user or group to Organizations, complete the following steps:

1. Click **Settings > Organizations**.
2. On the **Organizations** page, click the button and select **Edit**.
3. On the **Edit Organization** page, click **Assign Users/Groups** of the **Users/Groups** section.
4. Select a required IDP domain from the drop-down list at the top of the page.
5. In the **Assign Users/Groups** window, select the groups from the list and click **Assign**.
6. Click **Save**.
7. Repeat steps 1 to 6 to assign the same group to a different Organization.

Sign in Using an Organization ID

Users of an Organization with SSO assigned can sign in using the Organization ID. When a user signs in using an Organization ID, the sign-in goes through the SSO assigned to the Organization. The user can access only the objects assigned to the Organization and the role assigned to the user/group. The user cannot switch to any other Organization.

Prerequisites

- Users must be part of the IDP.
- Add the Cohesity cluster as an application on your IDP.
- A default role must be assigned for the user/group that is signing in to the Cohesity cluster.
- You must know the Organization ID. Cohesity Admin or Service provider can provide the Organization ID.

To sign in using the Organization ID, complete the following steps:

1. Access the Cohesity Dashboard.
2. On the **Sign in** page, click **Sign in with SSO**.
3. On the **Single Sign-On** page, click **Sign in to Organization**.

4. Provide the organization ID and click Sign in.
5. In the IDP login page, provide your username and password.
6. Click **Sign in**.

Sign in as a Group User

A user who is part of a group can sign in to the Cohesity cluster and manage the Organizations.

Prerequisites

- A user must be part of the IDP and added to the required group.
- Add the Cohesity cluster as an application on your IDP.
- Create a group and add the desired user on the cluster and assign that group to all the desired Organizations.

After signing in, the group user can see all the Organizations assigned to the groups that the user is a member and perform admin role functions within those Organizations.

To sign in as a group user, complete the following steps:

1. Access the Cohesity Dashboard.
2. On the **Sign in** page, click **Sign in with SSO** and click **Sign In**.
3. On the IDP login page, provide your username and password.
4. Click **Sign in**.

Edit or Deactivate an SSO for an Organization

You can edit or deactivate an SSO for any Organization only after accessing the Organization. If you are the Service Provider, you can impersonate and edit the SSO and if you are an Organization user or admin, you can log in to the Organization and edit.

> **Note:** If you deactivate SSO, only the Organization admin or a user with Access Management privileges can reactivate it.

To edit or deactivate an SSO for any Organization, complete the following steps:

1. Access the Cohesity Dashboard.
2. On the **Sign in** page, click **Sign in with SSO** and click **Sign In**.
3. Slide the **Sign in to Organization** toggle button.
4. Type your Organization ID and click **Sign In**.
5. On the IDP login page, provide your username and password.
6. Click **Sign in**.

7. After signing in, click **Settings > Access Management**.
8. Click the **SSO** tab, click the action menu (⋮) and select either **Edit** or **Deactivate**.

Multi-factor Authentication for Organizations

Multi-factor Authentication (MFA) is an additional layer of security used to verify the identity of a user. In a multitenant environment, the Cohesity supports MFA for local users. For more information about organization users, see Organization User Types.

A service provider administrator can enable or disable MFA at the cluster level. An Organization administrator can enable or disable MFA for user(s) in their organization.

To enable and set up MFA, follow the instructions in the Multi-factor Authentication for Local and Active Directory Users section.

To disable MFA for an Organization user, the service provider administrator or Organization administrator must complete the following steps:

1. In the Cohesity GUI, click the user icon ( ) in the upper-right corner and select **Impersonate**.
2. In the **Impersonate** dialog, select the required Organization from the list and click **Switch**.
3. Navigate to **Settings > Access Management** and click the **MFA** tab.
4. In the **MFA** column, turn off the toggle corresponding to the user:

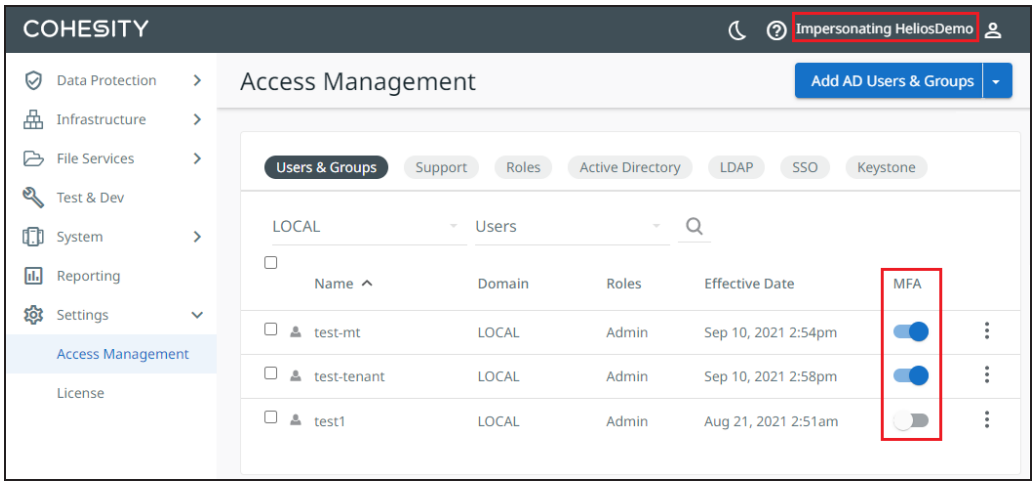

5. Click **Disable**.
6. Click **Save**.

Related Topics

- Supported TOTP Apps
- Multi-factor Authentication for Local and Active Directory Users
- Multi-factor Authentication for the Support User Account

Users and Groups for Organizations

Assign an Organization Administrator

After creating an Organization, you need to assign an Organization Administrator. The organization administrator can be a local user. You must provide the Administrator privileges for the user.

To assign an administrator to an organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization and select **Edit**.
3. In the **Edit Organization** page, click the **Assign Users/Groups**.
4. In the **Assign Users/Groups** page, select the user that you want to assign as an Administrator.

   > **Note:** If you have created a user with Administrator privileges, you can create one by clicking the **Add User** link on the **Assign Users/Groups** screen. After creating the user, assign the same user.

5. Click **Assign**.
6. Click **Save**.

Assign Local or AD Users and Groups

You can assign local users, AD users and groups to an Organization. Select one or more users from the list or create new ones. You can also add the users later.

> **Note:** When logging in to the cluster, use the syntax as given below: Username: <user>@<Organization ID>

Points to Remember

- When you assign a user or group to the Organization, the assigned user or group can access the Cohesity cluster using the user name and tenant ID.
- Users in the Organization can see all the users and roles in the Organization, but they cannot create or delete users or roles.
- Restricted users currently cannot be assigned to Organizations. If a user with the Administrator role is assigned to an Organization, they will have reduced cluster-level privileges.
    - AD users/groups can be assigned to the Organization. To login to the cluster or to access Views via SMB, use the AD credentials.
    - Who can perform this task: Service Provider Administrator/Cluster Administrator.

  To assign AD or local users to an Organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization and select **Edit**.
3. In the **Edit Organization** page, click **Assign Users/Groups**.
4. In the **Assign Users/Groups** page, select the required local or AD users or groups.
5. Click **Assign**.
6. Click **Save**.

Data Protection for Organizations

Assign Objects

Select the Sources to assign exclusively to the Organization.

Consider the following points for assigning objects:

- If you select a top-level node in the Object hierarchy, the entire Source is assigned to the Organization. The Organization becomes the owner and can manage the Source (unregister, edit, refresh, and protect it after the source is assigned to an Organization, the cluster administrator cannot edit or remove the source(read-only mode).
- If you select a sub-tree(VMware sources) in the Object hierarchy, only those Objects are assigned to the Organization. The Organization can protect and recover those Objects, but it cannot manage the Source. Conversely, the cluster administrator

cannot protect and recover those Objects, but can see and manage the Source. For more information, see Assigning Protection Sources to an Organization.

- Organizations can also register their own Sources after they log in to the cluster.

To assign an object to an Organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization and select **Edit**
3. In the **Edit Organization** page, click the **Assign Objects** link.

   > **Note:** You can filter the objects or search the objects using the filter and search options.

4. In the **Assign Objects** page, select the required objects and click **Assign**.
5. Click **Save**.

Assigning VCD objects

Assigning Sub-Trees

In vCloud Director, you can assign the following vCD sub-trees to a Cohesity organization.

- vCD Organisation
- Organisation VDC
- vApp
- VMs

When an organization(tenant) a user accesses the organization, the user can see only the sources assigned to that organization ensuring isolation among the organizations.

> **Note:** If a sub-tree (like a VM folder, ESXi host, or ResourcePool) is assigned to an organization, there must be at least one VM in the sub-tree to properly determine permissions during VM recovery.

The organization user can perform all supported operations for the assigned VCD sources using tenant login and the Cohesity vCD plugin.

> **Note:** The sub-trees can be assigned to Cohesity organizations in all network configurations (i.e. default-subnet-based, VLAN-based, and Hybrid Extender-based).

Assigning Top-level VCD

Although rare, it is possible to assign a top-level node in the VCD hierarchy (in other words, an entire VCD instance) to an organization.

Note: This is not supported for Cohesity organizations using a Hybrid Extender, but is supported for Cohesity organizations using other network configurations (i.e. default-subnet-based, VLAN-based).

Assigning Existing Protection Groups to an Organization

- Assign the Protection Group's storage domain to the organization.
- Assign the protection source to the organization.
- Assign the Protection Group's protection policy to the organization.
- Assign the Protection Group to the organization.

> **Note:** When you assign an existing protection group that is created at the cluster-level to an organization(tenant) and if there is any backup run after assigning the protection group to the organization(tenant), it will result in reindexing of files and folders if there are any changes in them. After the reindex, only the organization (tenant) users can search for files and folders.

Assign Policies and Protection Groups

Protection Groups once assigned to Organization cannot be unassigned. When you assign a protection group, you can assign protection policy, storage domain, and source to a particular job.

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization and select **Edit**
3. On the **Edit Organization** page, click **Assign Policies And Protection Groups**.
4. In the **Assign Policies And Protection Groups**, select a policy name, and click

Assign.

5. Click **Save**.

Impersonate an Organization

As a Service Provider or Cluster administrator, you can impersonate an Organization to verify what the Organization administrator will see in the cluster.

Actions performed when impersonating an Organization are logged and both the cluster administrator and Organization can see the audit entries by selecting **System > Audit Log.**

To impersonate an Organization administrator, complete the following steps:

1. In the Cohesity Dashboard, click the user icon ( ) in the top-right corner and select **Impersonate**

   > **Note:** To return to your own dashboard, click the user icon, and select **Switch Back**.

2. In the **Impersonate** screen, select the required Organization from the list and click **Switch**.

   > **Note:** Once you switch, the dashboard is refreshed to display what the Organization administrators will see (the top right corner of the dashboard indicates that you are impersonating the Organization administrator). You can impersonate multiple Organization administrators at the same time by opening the dashboard in different browser windows or tabs.

Deactivate an Organization

You can deactivate an Organization to stop access to the Organization.

> **Note:** Deactivating an Organization does not delete it. You can activate it back.

When an Organization is deactivated, its users cannot log in to the Cohesity Dashboard, future Protection Group runs are paused, and any replication runs will fail. Protection Groups that are currently in progress will complete.

To deactivate an Organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization, and select **Deactivate**.
3. In **the Deactivate Tenant** screen, click the **Deactivate** button.

> **Note:** The Status column in the Organizations page changes to Inactive.

Activate an Organization

You can activate an Organization that is in the deactivated status.

To activate an Organization, complete the following steps:

1. In the Cohesity Dashboard, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization, and select **Activate**.
3. In the **Activate Tenant** screen, click the **Activate** button.

> **Note:** The Status column in the Organizations page changes to Active.

Delete an Organization

When you delete an Organization, all the resources and users for the Organization are permanently deleted from the Cohesity cluster.

Deleting an Organization deletes all the entities of the Organization, such as protection groups, protection sources, LDAP, remote targets, users, groups, and recovery targets. Once deleted the above entities will be completely lost. However, service providers can still impersonate a deleted Organization. This allows the service providers to remove any entity (that gets stuck while deletion) manually.

> **Note:**
> - Policies and Storage Domains of the Organization will retain their association with the tenant even if the tenant is deleted.
> - The Organization ID and name cannot be reused after the Organization is deleted.

To delete an Organization, complete the following steps:

1. In the Cohesity **Dashboard**, navigate to **Settings > Access Management**.
2. In the **Organizations** tab of the **Access Management** page, click ⋮ for the required Organization, and select **Delete**.
3. In the **Tenant Delete** screen, click the **Delete** button.

The status of the Organization in the Organizations page changes to **Marked for Deletion**. The deletion process is an asynchronous process. Once the Organization is completely deleted, the status changes to **Deleted**.

CloudArchive and CloudTier for Organizations

CloudArchive

Using the Cohesity, you can quickly archive data to multiple external targets such as:

- Public clouds
- Private clouds - S3-compatible device
- Any NAS-NFSv3 from storage vendors
- QStar managed tape libraries

As a service provider, you can register an external target and add the target to a policy that you can assign to an Organization for CloudArchive.

If you want your Organization Administrator to perform the Cloud Archive, you should create a protection policy with CloudArchive and assign the protection policy to the required Organization.

CloudTier

Cohesity CloudTier enables an additional storage tier where cold data can be stored. After you register supported existing cloud storage as a CloudTier target with the Cohesity cluster, you can enable CloudTier to move rarely used and inactive data to the cloud when HDD used capacity exceeds a set threshold (typically 80%).

As a service provider, you can enable the CloudTier setting in a storage domain.

> **Note:** If you plan to allow any CloudTier functionality, do not share a storage domain among Organizations. If you enable CloudTier, Cohesity recommends mapping each storage domain to a single Organization.

Replication for Organizations

The Replication feature in Cohesity enables the service providers to provide Disaster Recovery as a Service (DRaaS) solution to their tenants.

For more information about replication, see: About Replication.

Prerequisites

- The Organizations feature must be enabled on the source cluster or target cluster to use the replication feature in the Organizations (multitenancy) environment.
- The source and target cluster must be running Cohesity Data Cloud (Self-managed) version 6.1 or later for Organizations to use the replication feature.
- When replicating data between multitenancy-enabled Clusters, the same Organization ID and matching storage domain(s) must exist on both the transmitting and receiving Clusters prior to running replication tasks.
- When using replication, the same org ID should exist on the primary and DR clusters for the Organization.

Disaster Recovery for Organizations (Multitenancy)

As a service provider, you can implement replication in the Organizations (multitenancy) environment using your Cohesity cluster for your Organizations (tenants) and provide Disaster Recovery as a Service (DRaaS). You can implement DRaaS depending on the backup network topology implemented. DRaaS service enables Organizations (tenants) to recover data in the event of data loss.

Replication is the fundamental operation that makes any Disaster Recovery as a Service (DRaaS) solution possible. Cohesity's DRaaS refers to the approach of maintaining a regularly updated replica of Organizational backup data in another Cohesity cluster in the same or different locations for data retention and offsite storage. Cohesity leverages service providers' multi-tenant infrastructure wherein they replicate data from multiple customers on shared but logically isolated infrastructure.

For more information on implementing Disaster Recovery as a Service (DRaaS), see the solution document: Use Cohesity Multi-tenancy to Offer Disaster Recovery as a Service (DRaaS).

Reports and Chargeback for Organizations

Cohesity provides many built-in reports that help you with planning, charge-back, compliance, and more.

Some of the other benefits of Cohesity Reports include:

- Schedule and receive reports through emails.
- Filter reports based on type, operation, time frame, or tenant Organization.
- Integrate with third-party business intelligence tools.

Cohesity provides required reports for Service Providers and Organizations:

- Storage Consumed by Organizations
- Storage Consumed by Storage Domains Report(Only for tenants)
- Backup Summary
- Clone and Recover VM Tasks Summary
- Failed Objects
- Protected Objects Heatmap (since 6.5.1 only)
- Protection Runs Summary
- Protection Summary by Object Type
- Storage Consumed by Backups
- Storage Consumed by Organizations
- Unprotected/Protected VMs

Important Points to Consider When Using Reports for Organizations

The following are the important points you must consider when using Reports for organizations:

- You must use the **Physical Usage** tab from the **Storage Consumed by Organizations** report that is available in the **Service Provider login** for billing purposes. This report is computed every two days.
- Do not use the **Storage Consumed by Organizations** report in the **tenant login** for billing purposes.
- The **Storage Consumed by Domains Report** displays the most current space usage data. However, you notice a mismatch in data between the **Storage Consumed by Domains Report** and **Storage Consumed by Organizations** as the latter is computed every 2 days.
- Some of the UI labels in different reports have different UI label names for the same attribute. So, consider them as synonyms. For example, the **Raw used**, **Physical used**, **Storage Consumed** labels in different reports are synonyms. These labels are respectively used in the following reports:
    - Storage Consumed by Storage Domains report (Service Provider login)

- Storage Consumed by Organizations (Tenant login)
- Storage Consumed by Organizations (Service Provider login)

For more information on Reports, see Reports.

Chargeback

As a service provider, you can charge your tenants for storage consumption. You can charge based on the pre-deduplication or post-deduplication of storage consumption by a tenant. Cohesity provides you with required raw data in the form of reports for your billing purposes.

You can get the required data for chargeback from the following:

- Storage Consumed by Organizations report
- Custom report that you can generate using REST APIs

Built-in Reports

View Storage Consumed by Organizations Report

This report provides information about the Organizations consuming the most storage.

By default, the report displays all the Organizations in the cluster, but you can specify one or more in the Organizations field. For more information, Storage Consumed by Organizations

To view the report, complete the following steps:

1. Access the Cohesity **Dashboard**.
2. In the **Dashboard**, click **Reporting**.
3. In the **Reports** tab, click the **Select Report** drop-down and select **Storage Consumed by Organization**.

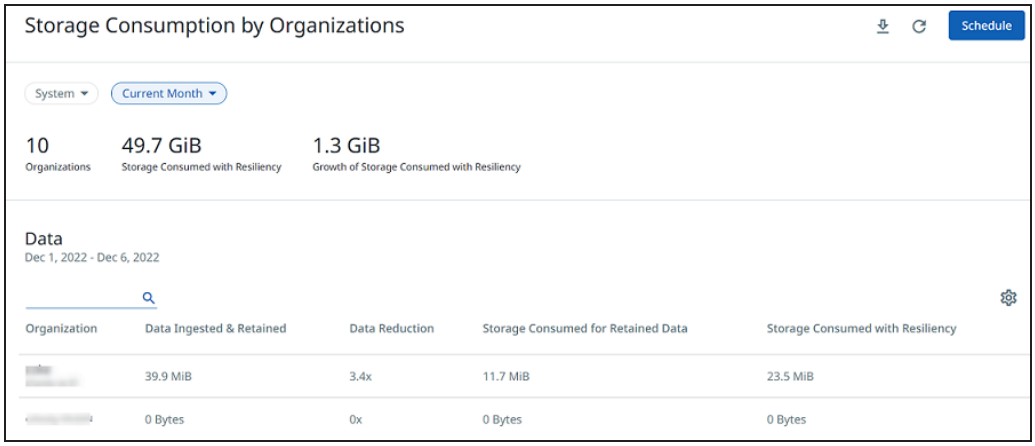

> **Note:** After selecting the **Storage Consumed by Organization** report, a report for all Organizations is displayed. However, you can apply filters and view reports that are specific to the storage domain and Organizations.

4. (Optional)Select the required **Storage Domain** from the dropdown.
5. (Optional)Select the required **Organizations** from the dropdown list.
6. (Optional) Click **Apply Filter**.
7. Click the **Physical Storage** or **Work Load** tabs to view the required details.

> **Note:** The Physical Storage tab displays the Top 3 Organizations Storage Consumed and the Work Load tab displays the Top 3 Organizations by Backup Data.

View Storage Consumed by Storage Domains Report

The Storage Consumed by Storage Domains report provides information about the Storage Domains that store the most logical data and raw data on the Cohesity cluster.

This report provides the following charts:

- Top <Num> Storage Domains by Storage Used—This chart lists up to ten Storage Domains that have stored the largest amount of data.
- Top <Num> Jobs by Logical Data Managed—This chart lists up to ten Storage Domains that have stored the largest amount of logical data from a Source.

To view the report, complete the following steps:

1. On the Cohesity Dashboard, click **Reporting**.
2. In the **Reports** tab, click the **Select Report** drop-down and select **Storage Consumed by Storage Domains**.

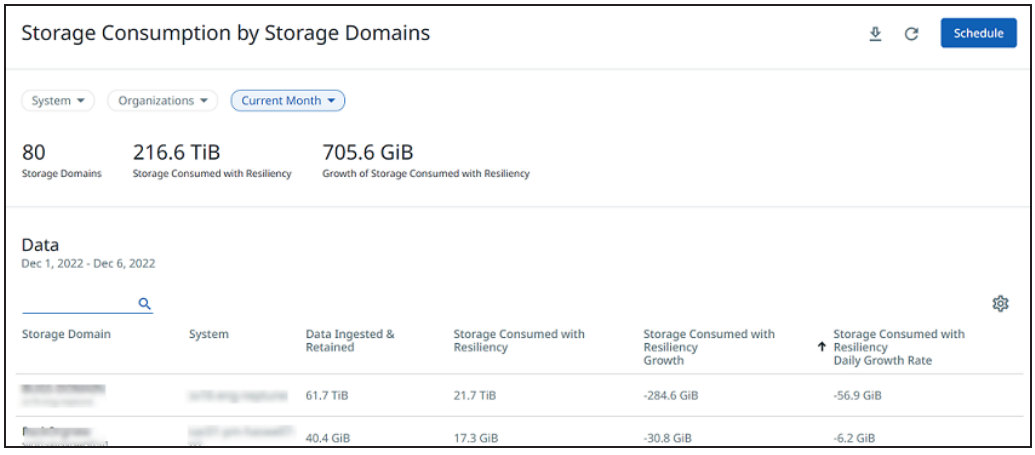

3. (Optional)Provide number of days to Calculate Growth Over Specified Days and click **Apply Filter**.

Custom Reports

The Custom Reporting Database is a new feature built into the platform using a PostgreSQL database. An internal service called Groot collects statistics and metrics from various internal platform services and populates data into Postgres. To generate reports, service providers write SQL statements with the schema Cohesity has published. For more information, Custom Reporting

> **Note:** The Custom Reports feature is not available for tenants.

Supported Multitenancy Workflows

This section lists the supported multitenancy workflows.

Protection Group and External Targets Workflows

- Virtual Server (VMs) and Physical Server Workflows
  - Protection Based Workflow
  - Archival Based Workflow
  - Replication Based Workflow
- Databases Workflows
  - Protection Based Workflow for MS SQL VM, FCI, and AG
  - Archival Based Workflow for MS SQL VM, FCI, and AG
  - Replication Based Workflow for MS SQL VM, FCI, and AG
  - Protection Based Workflow for MS SQL File-based, Volume-based, and Native Backup
  - Archival Based Workflow for MS SQL File-based, Volume-based, and Native Backup
  - Replication Based Workflow for MS SQL File-based, Volume-based, and Native Backup
  - Protection Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE
  - Archival Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE
  - Replication Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE

- NAS Workflows
  - Protection Based Workflow
  - Archival Based Workflow
  - Replication Based Workflow
- Workflows for Views, Remote Adapters
  - Protection Based Workflows
  - Archival Based Workflows
  - Replication Based Workflows
- Workflows for On-prem Applications
  - Protection Based Workflows
  - Archival Based Workflows
  - Replication Based Workflows
- Workflows for SaaS Applications
  - Protection Based Workflows
  - Archival Based Workflows
  - Replication Based Workflows
- AWS Cloud Workflows
  - Protection Based Workflows
  - Archival Based Workflows
  - Replication Based Workflows
- External Targets for CloudArchive
  - AWS External Targets for CloudArchive - Standard
  - AWS External Targets for CloudArchive - C2S
  - AWS External Targets for CloudArchive - Gov
  - Microsoft Azure External Targets for CloudArchive
  - GCP External Targets for CloudArchive
  - Oracle External Targets for CloudArchive
  - S3 - Compatible, NAS, Tape External Targets for CloudArchive
- External Target for Incremental Forever Archival/CloudArchive Direct with Dedup
- External Target for Incremental Forever Archival/CloudArchive Direct without Dedup
  - AWS External Targets for CloudArchiveDirect - Standard
  - AWS External Targets for CloudArchiveDirect - C2S
  - AWS External Targets for CloudArchiveDirect - Gov
  - GCP External Targets for CloudArchiveDirect
  - Oracle External Targets for CloudArchiveDirect

    - Microsoft Azure External Targets for CloudArchiveDirect
    - S3 - Compatible, NAS, Tape External Targets for CloudArchiveDirect
- Supported Workflow for NAS Tiering

Legend for Charts Below:

Yes = Supported

No = Not Supported

N/A = Not Applicable

CA = Cluster Administrator

TA = Tenant Administrator

Create = Create for an Org

Delete = Delete from Cluster

Assign = Assign to an Org

Un-assign = Un-assign from an Org

Virtual Server (VMs) and Physical Server Workflows

Protection Based Workflow

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
REGISTRATION	Register a source	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Auto-protect new objects	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
BACKUP	Incremental backup with CBT	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No
	Full backup with CBT	Yes	Yes	N/A	N/A	N/A	N/A	Supported (CA,TA)	Supported (CA,TA)	No	No
	Incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Full backup without CBT	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Backup, archive and replicate now	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	CloudSpin to AWS and Azure	Supported for Azure	Supported for Azure	No	Supported (CA,TA)	No	No	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recovering an entity in the Job Run	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Recover an entire backup job run	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Recover a SQL DB to the original server and instance	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover DB to Alternate Instance with overwite	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Instant volume mount	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Crash-consistent recover for apps	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLONE	Clone an object of a job run	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Clone all the objects of a job run	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Clone db to FCI instance	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
TEARDOWN	Tear down objects	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Tear down instant volume mount	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CANCEL	Cancel a running job run	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for a backed-up object	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download files/folders	Yes	Yes	Yes	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file	N/A	N/A	N/A	Supported (TA)	No	Supported (TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Microsoft Hyper-V 2016 / 2019 / 2022 VM Servers	Microsoft Hyper-V 2012 R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	*** vCloud Director	Cisco HyperFlex Storage Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A

Archival Based Workflow

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
ARCHIVAL	Archive an incremental backup	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Archive a full backup	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recover a single object of a job run from an archive to a source	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Recover an entire backup job run from an archive to a source	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Instant volume mount from an archive	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Crash-consistent recover for apps from an archive	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLONE	Clone objects of a job run from an archive	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Clone all objects of a job run from an archive	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CANCEL	Cancel a running archival task	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders from an archive	Yes	Yes	No	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for files or folders from an archive snapshot	No	No	No	No	No	No	No	No	No	No
	Download files and folders from an archive	Yes	Yes	Yes	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover files from an archive to a source (to the same and different sources)	No	No	No	Supported (CA,TA)	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download n VMX file from an archive	No	No	No	Supported (CA,TA)	No	No	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Yes	No	Yes	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download metadata for all runs of a job	Yes	No	Yes	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download a snapshot (archived data) for a job run	Yes	No	Yes	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Clone all objects of a job run from an archive	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	clone DB to FCI	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Recover all objects of job run from an archive to a source	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Instant volume mount from an archive	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Search files or folders from the metadata of a retrieved job run	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse files or folders from a retrieved snapshot	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Browse files or folders from the metadata of a downloaded job run	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Download files and folders from an archive	Yes	Yes	Yes	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover files from an archive to a source (to the same and different sources)	No	No	No	Supported (CA,TA)	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from an archive	N/A	N/A	N/A	No	No	No	N/A	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A

The Cohesity admin can either assign the entire VMware source or the components of the VMware source to an organization.

The following table lists different granularity at which VMware components can be assigned as a source to an organization (especially applicable for service providers that provide hosted IaaS).

Source	Components
VMware vCenter	• Entire VC • VMware DC • ESXi cluster • ESXi host • Folders • Resource pools • VMs
VMware vCloud Director	• Entire vCD • vCD Org • Org vDC • vApp • VMs

Replication Based Workflow

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Replicate a full backup with CBT to a remote cluster	Yes	Yes	N/A	N/A	N/A	N/A	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate a full backup without CBT to a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recover a single object of a job run from a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Recover an entire backup job run from a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A	N/A	N/A	N/A	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A	N/A	N/A	N/A	Supported (CA,TA)	N/A	N/A	N/A	N/A
	Instant volume mount of replicated data	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
DB Migration	Crash-consistent recover for apps from a remote cluster	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLONE	Clone an object of a job run from a remote cluster	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Clone all objects of a job run from a remote cluster	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	N/A	N/A	N/A			

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Crash-consistent clone for apps from a remote cluster	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
TEARDOWN	Tear down object on a remote cluster	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	No	No	No
	Tear down instant volume mount on a remote cluster	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CANCEL	Cancel a running replication task	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders on a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for replicated object on a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Yes	Yes	Yes	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover replicated files from a remote cluster	Yes	Yes	Yes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from a remote cluster	N/A	N/A	N/A	Supported (TA)	No	Supported (TA)	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	Yes	Yes	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	No	No	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources									
		Hyper-V 2016 / 2019 / 2022 VM Servers	Hyper-V 2012R2 VM Servers	Nutanix Acropolis (AHV)	ESXi VM Servers	vCloud Director	Cisco HyperFlex Storage Snapshot Integration	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

*** = For more information on cloning vCD VMs, refer to the Considerations section.

Databases Workflows

Protection Based Workflow for MS SQL VM, FCI, and AG

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	Yes using VM edit workflow	Supported (CA,TA)	Supported using Edit workflow for VM(CA,TA)
	Auto-protect new objects	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
BACKUP	Incremental backup with CBT	Yes as log-based incremental	Supported (CA,TA)	Supported (CA,TA)
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Incremental backup with SQL server differential	N/A	N/A	N/A
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A
	Backup, archive and replicate now	Supported at DB and server level both (CA,TA)	Supported at DB and server level both(CA,TA)	Supported at DB and server level both (CA,TA)
	CloudSpin to AWS and Azure	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	Only DB and VM supported(CA,TA)	Only DB and VM supported (CA,TA)
	Recover an entire backup job run	Yes as crash-consistent VM not as SQL restore	No	Supported as Crash Consistent VM not as SQL Restore(CA,TA)
	Recover a SQL DB to the original server and instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to the original server but a different instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover DB to Alternate Instance with overwite	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Yes
	Crash-consistent recover for apps	Supported (CA,TA)	Supported (CA,TA)	Yes using IVM (physical)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLONE	Clone an object of a job run	Supported (CA,TA)	Cloning DB from FCI to standalone server / node Supported. cloning db to fci instance not supported	Supported (CA,TA)
	Clone all the objects of a job run	Yes as crash-consistent VM, not as SQL clone	No	Supported as Crash Consistent VM not as SQL Clone(CA,TA)
	Clone db to FCI instance	N/A	N/A	N/A
	Crash-consistent recover for apps	Supported (CA,TA)	No	Supported using Instant volume mount (physical) (CA,TA)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
TEARDOWN	Tear down objects	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Tear down instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for a backed-up object	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download files/folders	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file	Supported (TA)	N/A	Supported (TA)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A

Archival Based Workflow for MS SQL VM, FCI, and AG

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
ARCHIVAL	Archive an incremental backup	Supported (Not for log backup) (CA,TA)	Supported (Not for log backup) (CA,TA)	Supported (Not for log backup) (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A
	Archive a full backup	Supported (Not for log backup) (CA,TA)	Supported (Not for log backup) (CA,TA)	Supported (Not for log backup) (CA,TA)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (No Databases, Only Files) (CA,TA)	Supported (No Databases, Only Files) (CA,TA)	Supported (No Databases, Only Files) (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical)(CA,TA)	Not Supported for SQL (use VM or physical)(CA,TA)	Not Supported for SQL (use VM or physical)(CA,TA)
	Crash-consistent recover for apps from an archive	Supported (No Databases, Only Files and VM) (CA,TA)	Supported (No Databases, Only Files)(CA,TA)	Supported (No Databases, Only Files and VM) (CA,TA)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLONE	Clone objects of a job run from an archive	No	No	No
	Clone all objects of a job run from an archive	No	No	No
	clone db to FCI	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	No	No	No
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	No	No	No
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
FILE RECOVERY	Search for files or folders from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for files or folders from an archive snapshot	No	No	No
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download n VMX file from an archive	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported for volume based job (CA)	Supported for volume based job (CA)	Supported for volume based job (CA)
	Download metadata for all runs of a job	Supported for volume based job (CA)	Supported for volume based job (CA)	Supported for volume based job (CA)
	Download a snapshot (archived data) for a job run	Supported for volume based job (CA)	Supported for volume based job (CA)	Supported for volume based job (CA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)
	Clone all objects of a job run from an archive	No	No	No
	clone DB to FCI	N/A	N/A	N/A
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)
	Recover all objects of job run from an archive to a source	No	No	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical, volume based)(CA,TA)	Not Supported for SQL (use VM or physical, volume based)(CA,TA)	Not Supported for SQL (use VM or physical, volume based)(CA,TA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)
	Search files or folders from the metadata of a retrieved job run	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)
	Browse files or folders from a retrieved snapshot	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)
	Browse files or folders from the metadata of a downloaded job run	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)
	Download files and folders from an archive	Supported for volume based job (TA)	Supported for volume based job (TA)	Supported for volume based job (TA)
	Recover files from an archive to a source (to the same and different sources)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)	Supported for volume based job (CA,TA)
	Download a VMX file from an archive	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Replication Based Workflow for MS SQL VM, FCI, and AG

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported as Log based incremental (CA,TA)	Supported as Log based incremental (CA,TA)	Supported as Log based incremental (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	Yes only for DBs & not full server for physical	Yes only for DBs & not full server for physical
	Recover an entire backup job run from a remote cluster	Yes as crash-consistent, not as SQL restore	No	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount of replicated data	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Crash-consistent recover for apps from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Only DB and VM supported (CA,TA)
	Clone all objects of a job run from a remote cluster	Yes as crash-consistent VM, not as SQL clone	No	No
	clone db to FCI	N/A	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	Supported (CA,TA)	No	Yes for DB and VM

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for replicated object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Supported (TA)	Supported (TA)	Supported (TA)
	Recover replicated files from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from a remote cluster	Supported (TA)	N/A	Supported but only VM(TA)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A

Protection Based Workflow for MS SQL File-based, Volume-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Auto-protect new objects	Supported (CA,TA)	Supported exclusion is not supported (CA,TA)	Supported exclusion is not supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
BACKUP	Incremental backup with CBT	Yes as log-based incremental	Supported (CA,TA)	N/A
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)	N/A
	Incremental backup with SQL server differential	N/A	N/A	Supported (CA,TA)
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	Supported (CA,TA)	Supported (CA,TA)
	Backup, archive and replicate now	Supported at DB and server level both(CA,TA)	Supported at DB and server level both (CA,TA)	Supported at DB and server level both (CA,TA)
	CloudSpin to AWS and Azure	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
RECOVER	Recovering an entity in the Job Run	Supported only for DBs and not Full Server (CA,TA)	Supported at DB level (CA,TA)	Supported at DB level (CA,TA)
	Recover an entire backup job run	No	No	No
	Recover a SQL DB to the original server and instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to the original server but a different instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover DB to Alternate Instance with overwite	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount	Supported (CA,TA)	N/A	N/A
	Crash-consistent recover for apps	Yes using IVM (physical)	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
CLONE	Clone an object of a job run	Supported (CA,TA)	Supported (CA,TA)	No
	Clone all the objects of a job run	No	N/A	N/A
	Clone db to FCI instance	N/A	Supported (CA,TA)	N/A
	Crash-consistent recover for apps	No	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	Supported (CA,TA)	N/A
	SQL DB Migration with auto sync	N/A	Supported (CA,TA)	N/A
TEARDOWN	Tear down objects	Supported (CA,TA)	Supported (CA,TA)	N/A
	Tear down instant volume mount	Supported (CA,TA)	N/A	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Yes at server level	Yes at server level

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	N/A	N/A
	Browse for a backed-up object	Supported (CA,TA)	N/A	N/A
	Download files/folders	Supported (TA)	N/A	N/A
	Recover files	Supported (CA,TA)	N/A	N/A
	Download a VMX file	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A

Archival Based Workflow for MS SQL File-based, Volume-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
ARCHIVAL	Archive an incremental backup	Supported (Not for log backup) (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	Supported (CA,TA)
	Archive a full backup	Supported (Not for log backup) (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (No Databases, Only Files) (CA,TA)	Supported at DB level (CA,TA)	Supported at DB level (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical) (CA,TA)	N/A	N/A
	Crash-consistent recover for apps from an archive	Supported (No Databases, Only Files)(CA,TA)	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
CLONE	Clone objects of a job run from an archive	No	Supported (CA,TA)	No
	Clone all objects of a job run from an archive	No	No	No
	clone db to FCI	N/A	Supported (CA,TA)	N/A
	Crash-consistent clone for apps from an archive	No	N/A	N/A
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	No	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
FILE RECOVERY	Search for files or folders from an archive	Supported (CA,TA)	N/A	N/A
	Browse for files or folders from an archive snapshot	No	N/A	N/A
	Download files and folders from an archive	Supported (TA)	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	Supported (CA,TA)	N/A	N/A
	Download n VMX file from an archive	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported for volume based job (CA)	No	No
	Download metadata for all runs of a job	Supported for volume based job (CA)	No	No
	Download a snapshot (archived data) for a job run	Supported for volume based job (CA)	No	No
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported for volume based job (CA,TA)	No	No
	Clone all objects of a job run from an archive	No	No	No
	clone DB to FCI	N/A	Supported (CA,TA)	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported for volume based job (CA,TA)	No	No
	Recover all objects of job run from an archive to a source	No	No	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical, volume based)(CA,TA)	No	No
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported for volume based job (CA,TA)	N/A	N/A
	Search files or folders from the metadata of a retrieved job run	Supported for volume based job (CA,TA)	N/A	N/A
	Browse files or folders from a retrieved snapshot	Supported for volume based job (CA,TA)	N/A	N/A
	Browse files or folders from the metadata of a downloaded job run	Supported for volume based job (CA,TA)	N/A	N/A
	Download files and folders from an archive	Supported for volume based job (TA)	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	Supported for volume based job (CA,TA)	N/A	N/A
	Download a VMX file from an archive	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Replication Based Workflow for MS SQL File-based, Volume-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported as Log based incremental (CA,TA)	Supported (CA,TA)	N/A
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	Supported (CA,TA)
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
RECOVER	Recover a single object of a job run from a remote cluster	Yes only for DBs & not full server for physical	Yes	Yes
	Recover an entire backup job run from a remote cluster	No	No	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount of replicated data	Supported (CA,TA)	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	No	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Only DB is supported (CA,TA)	No
	Clone all objects of a job run from a remote cluster	No	No	No
	clone db to FCI	N/A	Supported (CA,TA)	N/A
	Crash-consistent clone for apps from a remote cluster	No	No	No

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Only DB is supported (CA,TA)	Only DB is supported (CA,TA)
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)	N/A	N/A
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	N/A	N/A
	Browse for replicated object on a remote cluster	Supported (CA,TA)	N/A	N/A
	Download replicated files and folders from a remote cluster	Supported (TA)	N/A	N/A
	Recover replicated files from a remote cluster	Supported (CA,TA)	N/A	N/A
	Download a VMX file from a remote cluster	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A

Protection Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	N/A	Supported
	Refresh Application registration	Supported (CA,TA)	Supported (CA,TA)	N/A	Supported
	Auto-protect new objects	No	No	N/A	Supported

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Incremental backup with SQL server differential	N/A	N/A	N/A	Not Applicable
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Not Applicable
	Ability to run a log backup in parallel to a full backup.	No	No	Supported (CA,TA)	Not Applicable
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)	Yes except archive	Not Applicable
	CloudSpin to AWS and Azure	N/A	N/A	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Recover an entire backup job run	No	No	Supported (CA,TA)	Supported
	Recover a SQL DB to the original server and instance	Supported (CA,TA)	Supported (CA,TA)	N/A	Supported (Cloning is not supported)
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A	Supported
	Recover a SQL DB to a different server that has the same SQL version	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Recover DB to Alternate Instance with overwite	Supported (CA,TA)	N/A	N/A	Not Applicable
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	No	No	N/A	Not Applicable
	Instant volume mount	N/A	N/A	N/A	Not Applicable
	Crash-consistent recover for apps	N/A	N/A	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLONE	Clone an object of a job run	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Clone all the objects of a job run	No	No	N/A	Not Applicable
	Clone db to FCI instance	N/A	N/A	N/A	Not Supported
	Crash-consistent recover for apps	N/A	N/A	N/A	Not Applicable
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	Not Applicable
	SQL DB Migration with auto sync	N/A	N/A	N/A	Not Supported
TEARDOWN	Tear down objects	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Supported
	Tear down instant volume mount	N/A	N/A	N/A	Not Supported
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (Cloning is not supported) (CA,TA)	Supported
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
FILE RECOVERY	Search for files or folders	N/A	N/A	N/A	Not Applicable
	Browse for a backed-up object	N/A	N/A	N/A	Not Applicable
	Download files/folders	N/A	N/A	N/A	Not Applicable
	Recover files	N/A	N/A	N/A	Not Applicable
	Download a VMX file	N/A	N/A	N/A	Not Applicable
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	Not Applicable
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	Not Applicable

Archival Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A	Not Applicable
	Archive a full backup	Supported (CA,TA)	Supported (CA,TA)	No	Supported
RECOVER	Recover a single object of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Recover an entire backup job run from an archive to a source	No	No	No	Supported
	Instant volume mount from an archive	N/A	N/A	No	Not Applicable
	Crash-consistent recover for apps from an archive	N/A	N/A	N/A	Not Applicable
DB Migration	SQL DB Migration without auto sync	N/A	N/A	No	Not Applicable
	SQL DB Migration with auto sync	N/A	N/A	No	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Clone all objects of a job run from an archive	Not supported	Not supported	No	Not Applicable
	clone db to FCI	N/A	N/A	N/A	Not Applicable
	Crash-consistent clone for apps from an archive	N/A	N/A	N/A	Not Applicable
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Cancel a running clone or recover task that is retrieving from an archive	Not supported	Not supported	No	Supported (Cloning is not supported)
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	No	Supported

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
FILE RECOVERY	Search for files or folders from an archive	N/A	N/A	No	Not Applicable
	Browse for files or folders from an archive snapshot	N/A	N/A	No	Not Applicable
	Download files and folders from an archive	N/A	N/A	No	Not Applicable
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A	No	Not Applicable
	Download n VMX file from an archive	N/A	N/A	No	Not Applicable
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	No	Not supported	No	Not Supported
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)	No	Not Applicable
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)	No	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	No	Not Supported
	Clone all objects of a job run from an archive	No	Not supported	No	Not Supported
	clone DB to FCI	N/A	N/A	N/A	Not Supported
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Recover all objects of job run from an archive to a source	No	No	No	Supported
	Instant volume mount from an archive	N/A	N/A	No	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	N/A	N/A	No	Not Applicable
	Search files or folders from the metadata of a retrieved job run	N/A	N/A	No	Not Applicable
	Browse files or folders from a retrieved snapshot	N/A	N/A	No	Not Applicable
	Browse files or folders from the metadata of a downloaded job run	N/A	N/A	No	Not Applicable
	Download files and folders from an archive	N/A	N/A	No	Not Applicable
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A	No	Not Applicable
	Download a VMX file from an archive	N/A	N/A	N/A	Not Applicable
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable

Replication Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CBT is not applicable)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A	Not Applicable
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Not Applicable
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CBT is not applicable)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	N/A	Not Supported

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A	Supported
	Recover an entire backup job run from a remote cluster	Not supported	No	N/A	Not Supported
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	No	No	N/A	Not Applicable
	Instant volume mount of replicated data	N/A	N/A	N/A	Not Applicable
	Crash-consistent recover for apps from a remote cluster	N/A	N/A	N/A	Not Supported
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	Not Applicable
	SQL DB Migration with auto sync	N/A	N/A	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Clone all objects of a job run from a remote cluster	No	No	N/A	Not Applicable
	clone db to FCI	N/A	N/A	N/A	Not Applicable
	Crash-consistent clone for apps from a remote cluster	N/A	N/A	N/A	Not Applicable
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Supported
	Tear down instant volume mount on a remote cluster	N/A	N/A	N/A	Not Applicable
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Cancel a clone or recover task on a remote cluster	No	No	No	Supported (Clone not supported)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Supported

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
FILE RECOVERY	Search for files or folders on a remote cluster	N/A	N/A	No	Not Applicable
	Browse for replicated object on a remote cluster	N/A	N/A	No	Not Applicable
	Download replicated files and folders from a remote cluster	N/A	N/A	No	Not Applicable
	Recover replicated files from a remote cluster	N/A	N/A	No	Not Applicable
	Download a VMX file from a remote cluster	N/A	N/A	N/A	Not Applicable
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	Not Applicable
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	Not Applicable
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A	Not Applicable
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A	Not Applicable

NAS Workflows

Protection Based Workflow

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	N/A	N/A	N/A	N/A	N/A	N/A
	Auto-protect new objects	No	No	N/A	N/A	No	No

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Full backup with CBT	N/A	N/A	N/A	N/A	N/A	N/A
	Incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A	N/A
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A	N/A	N/A	N/A
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	CloudSpin to AWS and Azure	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run	No	No	No	No	No	No
	Recover a SQL DB to the original server and instance	N/A	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A	N/A	N/A	N/A	N/A
	Recover DB to Alternate Instance with overwite	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A	N/A	N/A	N/A	N/A
	Instant volume mount	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A	N/A	N/A	N/A	N/A
CLONE	Clone an object of a job run	No	No	No	No	No	No
	Clone all the objects of a job run	No	No	No	No	No	No
	Clone db to FCI instance	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A	N/A
TEARDOWN	Tear down objects	N/A	N/A	N/A	N/A	N/A	N/A
	Tear down instant volume mount	N/A	N/A	N/A	N/A	N/A	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
FILE RECOVERY	Search for files or folders	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Browse for a backed-up object	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download files/folders	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file	N/A	N/A	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	N/A	N/A	N/A

Archival Based Workflow

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A	N/A
	Archive a full backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
RECOVER	Recover a single object of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No	No	No	No	No
	Instant volume mount from an archive	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps from an archive	N/A	N/A	N/A	N/A	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from an archive	No	No	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	N/A	N/A	N/A	N/A	N/A	N/A
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
FILE RECOVERY	Search for files or folders from an archive	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Browse for files or folders from an archive snapshot	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download n VMX file from an archive	N/A	N/A	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from an archive	No	No	No	No	No	No
	clone DB to FCI	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover all objects of job run from an archive to a source	No	No	No	No	No	No
	Instant volume mount from an archive	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Search files or folders from the metadata of a retrieved job run	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Browse files or folders from a retrieved snapshot	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse files or folders from the metadata of a downloaded job run	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Download files and folders from an archive	Supported (Protection Run must be indexed) (TA)	Supported (Protection Run must be indexed) (TA)	Supported (Protection Run must be indexed) (TA)	Supported (Protection Run must be indexed) (TA)	Supported (Protection Run must be indexed) (TA)	Supported (Protection Run must be indexed) (TA)

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
	Recover files from an archive to a source (to the same and different sources)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from an archive	N/A	N/A	N/A	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A	N/A	N/A

Replication Based Workflow

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A
	Replicate a full backup without CBT to a remote cluster	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No	No	No	No

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
RECOVER	Recover a single object of a job run from a remote cluster	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery
	Recover an entire backup job run from a remote cluster	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery	Yes for single vol recovery
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
	Instant volume mount of replicated data	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
CLONE	Clone an object of a job run from a remote cluster	No	No	No	No	No	No
	Clone all objects of a job run from a remote cluster	No	No	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A
TEARDOWN	Tear down object on a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A
	Tear down instant volume mount on a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed)
	Browse for replicated object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover replicated files from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from a remote cluster	N/A	N/A	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources					
		NetApp - NFS	NetApp - SMB	Isilon - NFS	Isilon - SMB	Generic NAS - NFS	Generic NAS - SMB
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A	N/A	N/A	N/A

Workflows for Views, Remote Adapters

Protection Based Workflows

Categories	Workflows	Sources	
		Views	Remote Adapters
REGISTRATION	Register a source	N/A	N/A
	Refresh object hierarchy	N/A	N/A
	Refresh Application registration	N/A	N/A
	Auto-protect new objects	N/A	N/A
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)
	Full backup with CBT	N/A	N/A
	Incremental backup with SQL server differential	N/A	N/A
	Full backup without CBT	N/A	N/A
	Ability to run a log backup in parallel to a full backup.	N/A	N/A
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)
	CloudSpin to AWS and Azure	N/A	N/A

Categories	Workflows	Sources	
		Views	Remote Adapters
RECOVER	Recovering an entity in the Job Run	N/A	N/A
	Recover an entire backup job run	N/A	N/A
	Recover a SQL DB to the original server and instance	N/A	N/A
	Recover a SQL DB to the original server but a different instance	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A
	Recover DB to Alternate Instance with overwite	N/A	N/A
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A
	Instant volume mount	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A
CLONE	Clone an object of a job run	Supported (CA,TA)	Supported (CA,TA)
	Clone all the objects of a job run	Yes	Yes
	Clone db to FCI instance	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A
TEARDOWN	Tear down objects	N/A	N/A
	Tear down instant volume mount	N/A	N/A

Categories	Workflows	Sources	
		Views	Remote Adapters
CANCEL	Cancel a running job run	N/A	Supported (CA,TA)
	Cancel a running clone or recover tasks	No	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	No
	Browse for a backed-up object	No	No
	Download files/folders	Supported (TA)	No
	Recover files	No	No
	Download a VMX file	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A

Archival Based Workflows

Categories	Workflows	Sources	
		Views	Remote Adapters
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A
	Archive a full backup	N/A	N/A
RECOVER	Recover a single object of a job run from an archive to a source	N/A	N/A
	Recover an entire backup job run from an archive to a source	N/A	N/A
	Instant volume mount from an archive	N/A	N/A
	Crash-consistent recover for apps from an archive	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A
CLONE	Clone objects of a job run from an archive	N/A	N/A
	Clone all objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)
	clone db to FCI	N/A	N/A
	Crash-consistent clone for apps from an archive	N/A	N/A
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	No	Supported (CA,TA)

Categories	Workflows	Sources	
		Views	Remote Adapters
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders from an archive	Supported (CA,TA)	No
	Browse for files or folders from an archive snapshot	No	No
	Download files and folders from an archive	Supported (TA)	No
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A
	Download n VMX file from an archive	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported (CA)	Supported (CA)
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)
	clone DB to FCI	N/A	N/A

Categories	Workflows	Sources	
		Views	Remote Adapters
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	N/A	N/A
	Recover all objects of job run from an archive to a source	N/A	N/A
	Instant volume mount from an archive	N/A	N/A
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	No	No
	Search files or folders from the metadata of a retrieved job run	No	No
	Browse files or folders from a retrieved snapshot	No	No
	Browse files or folders from the metadata of a downloaded job run	No	No
	Download files and folders from an archive	No	No
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A
	Download a VMX file from an archive	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A

Replication Based Workflows

Categories	Workflows	Sources	
		Views	Remote Adapters
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Yes	Yes
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	N/A	N/A
	Replicate a full backup without CBT to a remote cluster	N/A	N/A
FAILOVER	Failover an inactive Job on a remote cluster	No	No
RECOVER	Recover a single object of a job run from a remote cluster	N/A	N/A
	Recover an entire backup job run from a remote cluster	N/A	N/A
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A
	Instant volume mount of replicated data	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A

Categories	Workflows	Sources	
		Views	Remote Adapters
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)
	clone db to FCI	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	N/A	N/A
TEARDOWN	Tear down object on a remote cluster	N/A	N/A
	Tear down instant volume mount on a remote cluster	N/A	N/A
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	No	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders on a remote cluster	No	No
	Browse for replicated object on a remote cluster	No	No
	Download replicated files and folders from a remote cluster	No	No
	Recover replicated files from a remote cluster	No	No
	Download a VMX file from a remote cluster	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A

Categories	Workflows	Sources	
		Views	Remote Adapters
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A

Workflows for On-prem Applications

Protection Based Workflows

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
REGISTRATION	Register a source	Yes	Yes
	Refresh object hierarchy	Yes	Yes
	Refresh Application registration	Yes	Yes
	Auto-protect new objects	N/A	Yes

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
BACKUP	Incremental backup with CBT	Yes	Yes
	Full backup with CBT	Yes	Yes
	Incremental backup with SQL server differential	N/A	N/A
	Full backup without CBT	Yes	Yes
	Ability to run a log backup in parallel to a full backup.	N/A	N/A
	Backup, archive and replicate now	Yes	Yes
	CloudSpin to AWS and Azure	N/A	N/A
RECOVER	Recovering an entity in the Job Run	No	Yes
	Recover an entire backup job run	Yes	No
	Recover a SQL DB to the original server and instance	N/A	N/A
	Recover a SQL DB to the original server but a different instance	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A
	Instant volume mount	Yes	No
	Crash-consistent recover for apps	No	No

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
CLONE	Clone an object of a job run	N/A	N/A
	Clone all the objects of a job run	N/A	N/A
	Clone db to FCI instance	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A
TEARDOWN	Tear down objects	Yes	Yes
	Tear down instant volume mount	Yes	No
CANCEL	Cancel a running job run	Yes	Yes
	Cancel a running clone or recover tasks	Yes except for cloning	Yes (except for cloning)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Yes	Yes
FILE RECOVERY	Search for files or folders	Yes	No
	Browse for a backed-up object	Yes	No
	Download files/folders	Yes	No
	Recover files	Yes	No
	Download a VMX file	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A

Archival Based Workflows

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
ARCHIVAL	Archive an incremental backup	Yes	Yes
	Archiving an incremental backup with SQL server differential	N/A	N/A
	Archive a full backup	Yes	Yes
RECOVER	Recover a single object of a job run from an archive to a source	No	Yes
	Recover an entire backup job run from an archive to a source	Yes	No
	Instant volume mount from an archive	Yes	No
	Crash-consistent recover for apps from an archive	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
CLONE	Clone objects of a job run from an archive	N/A	N/A
	Clone all objects of a job run from an archive	N/A	N/A
	clone db to FCI	N/A	N/A
	Crash-consistent clone for apps from an archive	N/A	N/A
CANCEL	Cancel a running archival task	Yes	Yes
	Cancel a running clone or recover task that is retrieving from an archive	Yes (Except cloning)	Yes (Except cloning)
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Yes	Yes
FILE RECOVERY	Search for files or folders from an archive	No	No
	Browse for files or folders from an archive snapshot	No	No
	Download files and folders from an archive	No	No
	Recover files from an archive to a source (to the same and different sources)	No	No
	Download n VMX file from an archive	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	No	No
	Download metadata for all runs of a job	Yes	Yes
	Download a snapshot (archived data) for a job run	Yes	Yes
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	No	No
	Clone all objects of a job run from an archive	No	No
	clone DB to FCI	N/A	N/A
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	No	Yes
	Recover all objects of job run from an archive to a source	Yes	No
	Instant volume mount from an archive	Yes	No

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	No	No
	Search files or folders from the metadata of a retrieved job run	No	No
	Browse files or folders from a retrieved snapshot	Yes	No
	Browse files or folders from the metadata of a downloaded job run	No	No
	Download files and folders from an archive	Yes	No
	Recover files from an archive to a source (to the same and different sources)	Yes	No
	Download a VMX file from an archive	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A

Replication Based Workflows

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Yes	Yes
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	Yes	Yes
	Replicate a full backup without CBT to a remote cluster	Yes	Yes
FAILOVER	Failover an inactive Job on a remote cluster	No	No
RECOVER	Recover a single object of a job run from a remote cluster	No	Yes
	Recover an entire backup job run from a remote cluster	Yes	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A
	Instant volume mount of replicated data	Yes	No
	Crash-consistent recover for apps from a remote cluster	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
CLONE	Clone an object of a job run from a remote cluster	N/A	N/A
	Clone all objects of a job run from a remote cluster	N/A	N/A
	clone db to FCI	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	N/A	N/A
TEARDOWN	Tear down object on a remote cluster	Yes	Yes
	Tear down instant volume mount on a remote cluster	Yes	No
CANCEL	Cancel a running replication task	Yes	Yes
	Cancel a clone or recover task on a remote cluster	Yes (Except cloning)	Yes (Except cloning)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Yes	Yes
FILE RECOVERY	Search for files or folders on a remote cluster	No	No
	Browse for replicated object on a remote cluster	Yes	No
	Download replicated files and folders from a remote cluster	Yes	No
	Recover replicated files from a remote cluster	Yes	No
	Download a VMX file from a remote cluster	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A

Categories	Workflows	Sources	
		Microsoft Active Directory	Microsoft Exchange OnPrem
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A

Workflows for SaaS Applications

Protection Based Workflows

Categories	Workflows	Sources
		Office 365
REGISTRATION	Register a source	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)
	Refresh Application registration	N/A
	Auto-protect new objects	Supported (CA,TA)
BACKUP	Incremental backup with CBT	Supported (CA,TA)
	Full backup with CBT	N/A
	Incremental backup with SQL server differential	N/A
	Full backup without CBT	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A
	Backup, archive and replicate now	Supported (CA,TA)
	CloudSpin to AWS and Azure	N/A

Categories	Workflows	Sources
		Office 365
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)
	Recover an entire backup job run	Supported (CA,TA)
	Recover a SQL DB to the original server and instance	N/A
	Recover a SQL DB to the original server but a different instance	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A
	Instant volume mount	N/A
	Crash-consistent recover for apps	N/A
CLONE	Clone an object of a job run	N/A
	Clone all the objects of a job run	N/A
	Clone db to FCI instance	N/A
	Crash-consistent recover for apps	N/A
DB Migration	SQL DB Migration without auto sync	N/A
	SQL DB Migration with auto sync	N/A
TEARDOWN	Tear down objects	N/A
	Tear down instant volume mount	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (Cloning is not supported) (CA,TA)

Categories	Workflows	Sources
		Office 365
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)
FILE RECOVERY	Search for files or folders	Searching for Files/Folders, mails, and mailboxes Supported (CA,TA)
	Browse for a backed-up object	Supported (CA,TA) for Onedrive and SharePoint Not Supported for Outlook
	Download files/folders	Supported for SharePoint Not Supported for Outlook Only Files Suported for Onedrive; not Folders (TA)
	Recover files	Recovering Files/Folders, mails, and mailboxes supported (CA,TA)
	Download a VMX file	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A
	Register AWS source and clone VMs to AWS	N/A

Archival Based Workflows

Categories	Workflows	Sources
		Office 365
ARCHIVAL	Archive an incremental backup	No
	Archiving an incremental backup with SQL server differential	N/A
	Archive a full backup	No
RECOVER	Recover a single object of a job run from an archive to a source	No
	Recover an entire backup job run from an archive to a source	No
	Instant volume mount from an archive	N/A
	Crash-consistent recover for apps from an archive	N/A
DB Migration	SQL DB Migration without auto sync	N/A
	SQL DB Migration with auto sync	N/A
CLONE	Clone objects of a job run from an archive	N/A
	Clone all objects of a job run from an archive	N/A
	clone db to FCI	N/A
	Crash-consistent clone for apps from an archive	N/A
CANCEL	Cancel a running archival task	No
	Cancel a running clone or recover task that is retrieving from an archive	No
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	No

Categories	Workflows	Sources
		Office 365
FILE RECOVERY	Search for files or folders from an archive	No
	Browse for files or folders from an archive snapshot	No
	Download files and folders from an archive	No
	Recover files from an archive to a source (to the same and different sources)	No
	Download n VMX file from an archive	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	No
	Download metadata for all runs of a job	No
	Download a snapshot (archived data) for a job run	No
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	No
	Clone all objects of a job run from an archive	No
	clone DB to FCI	N/A
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	No
	Recover all objects of job run from an archive to a source	No
	Instant volume mount from an archive	N/A

Categories	Workflows	Sources
		Office 365
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	No
	Search files or folders from the metadata of a retrieved job run	No
	Browse files or folders from a retrieved snapshot	No
	Browse files or folders from the metadata of a downloaded job run	No
	Download files and folders from an archive	No
	Recover files from an archive to a source (to the same and different sources)	No
	Download a VMX file from an archive	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A

Replication Based Workflows

Categories	Workflows	Sources
		Office 365
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A
	Replicate a full backup with CBT to a remote cluster	N/A
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No

Categories	Workflows	Sources
		Office 365
RECOVER	Recover a single object of a job run from a remote cluster	No
	Recover an entire backup job run from a remote cluster	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A
	Instant volume mount of replicated data	N/A
	Crash-consistent recover for apps from a remote cluster	N/A
DB Migration	SQL DB Migration without auto sync	N/A
	SQL DB Migration with auto sync	N/A
CLONE	Clone an object of a job run from a remote cluster	N/A
	Clone all objects of a job run from a remote cluster	N/A
	clone db to FCI	N/A
	Crash-consistent clone for apps from a remote cluster	N/A
TEARDOWN	Tear down object on a remote cluster	N/A
	Tear down instant volume mount on a remote cluster	N/A
CANCEL	Cancel a running replication task	No
	Cancel a clone or recover task on a remote cluster	No
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	No

Categories	Workflows	Sources
		Office 365
FILE RECOVERY	Search for files or folders on a remote cluster	No
	Browse for replicated object on a remote cluster	No
	Download replicated files and folders from a remote cluster	No
	Recover replicated files from a remote cluster	No
	Download a VMX file from a remote cluster	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A
	Register AWS source and clone VMs to AWS	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A
	Create a snapshot in Azure and replicate to another region	N/A

AWS Cloud Workflows

Protection Based Workflows

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	N/A	N/A	N/A
	Auto-protect new objects	Supported (CA,TA)	No	No
BACKUP	Incremental backup with CBT	Supported (CA,TA)	No	Supported(AWS does incremental backups)(CA,TA)
	Full backup with CBT	N/A	N/A	N/A
	Incremental backup with SQL server differential	N/A	N/A	N/A
	Full backup without CBT	N/A	N/A	N/A
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A
	Backup, archive and replicate now	Supported (CA,TA)	No	N/A
	CloudSpin to AWS and Azure	N/A	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run	Supported (CA,TA)	Supported (CA,TA)	No
	Recover a SQL DB to the original server and instance	N/A	N/A	N/A
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A	N/A
	Recover DB to Alternate Instance with overwite	N/A	N/A	N/A
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A	N/A
	Instant volume mount	No	N/A	N/A
	Crash-consistent recover for apps	No	N/A	N/A
CLONE	Clone an object of a job run	No	N/A	N/A
	Clone all the objects of a job run	No	N/A	N/A
	Clone db to FCI instance	N/A	N/A	N/A
	Crash-consistent recover for apps	No	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
TEARDOWN	Tear down objects	No	N/A	N/A
	Tear down instant volume mount	No	N/A	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	N/A	N/A
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	No	No
	Browse for a backed-up object	Supported (CA,TA)	No	No
	Download files/folders	Supported (TA)	No	No
	Recover files	Supported (CA,TA)	No	No
	Download a VMX file	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	No	N/A	N/A

Archival Based Workflows

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
ARCHIVAL	Archive an incremental backup	No	N/A	N/A
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A
	Archive a full backup	No	N/A	N/A
RECOVER	Recover a single object of a job run from an archive to a source	No	N/A	N/A
	Recover an entire backup job run from an archive to a source	No	N/A	N/A
	Instant volume mount from an archive	N/A	N/A	N/A
	Crash-consistent recover for apps from an archive	N/A	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
CLONE	Clone objects of a job run from an archive	N/A	N/A	N/A
	Clone all objects of a job run from an archive	N/A	N/A	N/A
	clone db to FCI	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	N/A	N/A	N/A
CANCEL	Cancel a running archival task	N/A	N/A	N/A
	Cancel a running clone or recover task that is retrieving from an archive	N/A	N/A	N/A
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	N/A	N/A	N/A
FILE RECOVERY	Search for files or folders from an archive	N/A	N/A	N/A
	Browse for files or folders from an archive snapshot	N/A	N/A	N/A
	Download files and folders from an archive	N/A	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A	N/A
	Download n VMX file from an archive	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	N/A	N/A	N/A
	Download metadata for all runs of a job	N/A	N/A	N/A
	Download a snapshot (archived data) for a job run	N/A	N/A	N/A
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	N/A	N/A	N/A
	Clone all objects of a job run from an archive	N/A	N/A	N/A
	clone DB to FCI	N/A	N/A	N/A
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	N/A	N/A	N/A
	Recover all objects of job run from an archive to a source	N/A	N/A	N/A
	Instant volume mount from an archive	N/A	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	N/A	N/A	N/A
	Search files or folders from the metadata of a retrieved job run	N/A	N/A	N/A
	Browse files or folders from a retrieved snapshot	N/A	N/A	N/A
	Browse files or folders from the metadata of a downloaded job run	N/A	N/A	N/A
	Download files and folders from an archive	N/A	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A	N/A
	Download a VMX file from an archive	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Replication Based Workflows

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CA,TA)	N/A	N/A
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	N/A	N/A	N/A
	Replicate a full backup without CBT to a remote cluster	N/A	N/A	N/A
FAILOVER	Failover an inactive Job on a remote cluster	Supported(Only during Failback) (CA,TA)	N/A	N/A
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	N/A	N/A
	Recover an entire backup job run from a remote cluster	Supported (CA,TA)	N/A	N/A
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A	N/A
	Instant volume mount of replicated data	No	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	N/A	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
CLONE	Clone an object of a job run from a remote cluster	No	N/A	N/A
	Clone all objects of a job run from a remote cluster	No	N/A	N/A
	clone db to FCI	N/A	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	No	N/A	N/A
TEARDOWN	Tear down object on a remote cluster	No	N/A	N/A
	Tear down instant volume mount on a remote cluster	No	N/A	N/A
CANCEL	Cancel a running replication task	Supported (CA,TA)	N/A	N/A
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	N/A	N/A
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	N/A	N/A

Categories	Workflows	Sources		
		AWS Instances (Native Backup)	AWS Instances (CSM)	AWS RDS
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	N/A	N/A
	Browse for replicated object on a remote cluster	Supported (CA,TA)	N/A	N/A
	Download replicated files and folders from a remote cluster	Supported (TA)	N/A	N/A
	Recover replicated files from a remote cluster	Supported (CA,TA)	N/A	N/A
	Download a VMX file from a remote cluster	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	No	Supported (CA,TA)
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A

External Targets for CloudArchive

AWS External Targets for CloudArchive - Standard

Workflows	External Targets					
	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone	AWS - Glacier (Legacy)	AWS - S3-Glacier	AWS - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	No	No	No	No	No	No
Enable/disable incremental archive at external target level	No	No	No	No	No	No
Enable encryption at external target level	No	No	No	No	No	No
Once enabled, disable encryption at external target level	No	No	No	No	No	No
Enable/disable compression at external target level	No	No	No	No	No	No
Delete external target	No	No	No	No	No	No
Enable secure archive at external target level	No	No	No	No	No	No
Once enabled, disable Secure Archive at external target level	No	No	No	No	No	No
Perform search for CloudRetrieve operations	No	No	No	No	No	No
Download jobs (CloudRetrieve)	No	No	No	No	No	No

Workflows	External Targets					
	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone	AWS - Glacier (Legacy)	AWS - S3-Glacier	AWS - S3 Glacier Deep Archive
Download files/folders from archive	No	No	No	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No	No	No	No
Recover files/folders to source	No	No	No	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No	No	No	No

AWS External Targets for CloudArchive - C2S

Workflows	External Targets	
	AWS C2S - S3	AWS C2S - Glacier
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No

Workflows	External Targets	
	AWS C2S - S3	AWS C2S - Glacier
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

AWS External Targets for CloudArchive - Gov

Workflows	External Targets	
	AWS Gov - S3	AWS Gov - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No

Workflows	External Targets	
	AWS Gov - S3	AWS Gov - S3 Glacier Deep Archive
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

Microsoft Azure External Targets for CloudArchive

Workflows	External targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Enable/disable deduplication at external target level	No	No	No	No
Enable/disable incremental archive at external target level	No	No	No	No
Enable encryption at external target level	No	No	No	No
Once enabled, disable encryption at external target level	No	No	No	No
Enable/disable compression at external target level	No	No	No	No
Delete external target	No	No	No	No
Enable secure archive at external target level	No	No	No	No
Once enabled, disable Secure Archive at external target level	No	No	No	No
Perform search for CloudRetrieve operations	No	No	No	No

Workflows	External targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Download jobs (CloudRetrieve)	No	No	No	No
Download files/folders from archive	No	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No	No
Recover files/folders to source	No	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No	No

GCP External Targets for CloudArchive

Workflows	External Targets		
	Google - Standard	Google - Nearline	Google - Coldline
Enable/disable deduplication at external target level	No	No	No
Enable/disable incremental archive at external target level	No	No	No
Enable encryption at external target level	No	No	No
Once enabled, disable encryption at external target level	No	No	No
Enable/disable compression at external target level	No	No	No
Delete external target	No	No	No
Enable secure archive at external target level	No	No	No
Once enabled, disable Secure Archive at external target level	No	No	No
Perform search for CloudRetrieve operations	No	No	No
Download jobs (CloudRetrieve)	No	No	No

Workflows	External Targets		
	Google - Standard	Google - Nearline	Google - Coldline
Download files/folders from archive	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No
Recover files/folders to source	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No

Oracle External Targets for CloudArchive

Workflows	External Targets	
	Oracle Standard	Oracle Archive
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No

Workflows	External Targets	
	Oracle Standard	Oracle Archive
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

S3 - Compatible, NAS, Tape External Targets for CloudArchive

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Enable/disable deduplication at external target level	No	No	No
Enable/disable incremental archive at external target level	No	No	No
Enable encryption at external target level	No	No	No
Once enabled, disable encryption at external target level	No	No	No
Enable/disable compression at external target level	No	No	No
Delete external target	No	No	No
Enable secure archive at external target level	No	No	No
Once enabled, disable Secure Archive at external target level	No	No	No
Perform search for CloudRetrieve operations	No	No	No
Download jobs (CloudRetrieve)	No	No	No
Download files/folders from archive	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No
Recover files/folders to source	No	No	No

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No

External Target for Incremental Forever Archival/CloudArchive Direct with Dedup

External Targets for CloudArchiveDirect

Workflows	External Targets						
	NAS	S3 - Compatible (E.g., Cohesity)	Azure Hot Blob	Azure Cool Blob	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone
Enable encryption at external target level	No	No	No	No	No	No	No
Once enabled, disable encryption at external target level	No	No	No	No	No	No	No
Enable/disable compression at external target level	No	No	No	No	No	No	No
Delete external target	No	No	No	No	No	No	No
Enable secure archive at external target level	No	No	No	No	No	No	No

Workflows	External Targets						
	NAS	S3 - Compatible (E.g., Cohesity)	Azure Hot Blob	Azure Cool Blob	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone
Once enabled, disable Secure Archive at external target level	No	No	No	No	No	No	No
Once enabled, going back to legacy archival	No	No	No	No	No	No	No
Perform search for CloudRetrieve operations	No	No	No	No	No	No	No
Download jobs (CloudRetrieve)	No	No	No	No	No	No	No
Download files/folders from archive	No	No	No	No	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No	No	No	No	No
Recover files/folders to source	No	No	No	No	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No	No	No	No	No

External Target for Incremental Forever Archival/CloudArchive Direct without Dedup

AWS External Targets for CloudArchiveDirect - Standard

Workflows	External Targets			
	AWS - S3-Intelligent	AWS - S3-One-Zone	AWS - Glacier	AWS - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	No	No	No	No
Enable/disable incremental archive at external target level	No	No	No	No
Enable encryption at external target level	No	No	No	No
Once enabled, disable encryption at external target level	No	No	No	No
Enable/disable compression at external target level	No	No	No	No
Delete external target	No	No	No	No
Enable secure archive at external target level	No	No	No	No
Once enabled, disable Secure Archive at external target level	No	No	No	No
Perform search for CloudRetrieve operations	No	No	No	No
Download jobs (CloudRetrieve)	No	No	No	No
Download files/folders from archive	No	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No	No
Recover files/folders to source	No	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No	No

AWS External Targets for CloudArchiveDirect - C2S

Workflows	External Targets	
	AWS C2S - S3	AWS C2S - Glacier
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

AWS External Targets for CloudArchiveDirect - Gov

Workflows	External Targets	
	AWS Gov - S3	AWS Gov - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	No	No

Workflows	External Targets	
	AWS Gov - S3	AWS Gov - S3 Glacier Deep Archive
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

GCP External Targets for CloudArchiveDirect

Workflows	External Targets	
	Google - Nearline	Google - Coldline
Enable/disable deduplication at external target level	No	No

Workflows	External Targets	
	Google - Nearline	Google - Coldline
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

Oracle External Targets for CloudArchiveDirect

Workflows	External Targets	
	Oracle Standard	Oracle Archive
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No

Workflows	External Targets	
	Oracle Standard	Oracle Archive
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

Microsoft Azure External Targets for CloudArchiveDirect

Workflows	External Targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Enable/disable deduplication at external target level	No	No	No	No
Enable/disable incremental archive at external target level	No	No	No	No
Enable encryption at external target level	No	No	No	No
Once enabled, disable encryption at external target level	No	No	No	No
Enable/disable compression at external target level	No	No	No	No

Workflows	External Targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Delete external target	No	No	No	No
Enable secure archive at external target level	No	No	No	No
Once enabled, disable Secure Archive at external target level	No	No	No	No
Perform search for CloudRetrieve operations	No	No	No	No
Download jobs (CloudRetrieve)	No	No	No	No
Download files/folders from archive	No	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No	No
Recover files/folders to source	No	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No	No

S3 - Compatible, NAS, Tape External Targets for CloudArchiveDirect

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Enable/disable deduplication at external target level	No	No	No
Enable/disable incremental archive at external target level	No	No	No
Enable encryption at external target level	No	No	No
Once enabled, disable encryption at external target level	No	No	No
Enable/disable compression at external target level	No	No	No

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Delete external target	No	No	No
Enable secure archive at external target level	No	No	No
Once enabled, disable Secure Archive at external target level	No	No	No
Perform search for CloudRetrieve operations	No	No	No
Download jobs (CloudRetrieve)	No	No	No
Download files/folders from archive	No	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No	No
Recover files/folders to source	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No	No

Supported Workflow for NAS Tiering

NAS Tiering

Workflows	External Targets			
	Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
Down-Tiering without leaving anything at source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Downtiering with leaving the symlink behind at the source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Orphan data delete	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Workflows	External Targets			
	Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
Capacity Planner	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Tiering Audit Log	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Bucket Customization	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

* = When all protocols are enabled, S3 is read-only mode and not
** = Yes for either SMB or NFS protocol

Supported Multitenancy Workflows for Hybrid Extender

This sections lists the supported multitenancy workflows for Hybrid Extender.

Protection Group and External Targets Workflows (Hybrid Extender)

- Virtual Server (VMs) and Physical Server Workflows (Hybrid Extender)
  - Protection Based Workflow
  - Archival Based Workflow
  - Replication Based Workflow
- Databases Workflows (Hybrid Extender)
  - Protection Based Workflow for MS SQL VM, FCI, and AG
  - Archival Based Workflow for MS SQL VM, FCI, and AG
  - Replication Based Workflow for MS SQL VM, FCI, and AG
  - Protection Based Workflow for MS SQL Volume-base, File-based, and Native Backup
  - Archival Based Workflow for MS SQL Volume-base, File-based, and Native Backup
  - Replication Based Workflow for MS SQL Volume-base, File-based, and Native Backup
  - Protection Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE
  - Archival Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE
  - Replication Based Workflow Oracle, SAP HANA, and SAP Sybase ASE

- NAS Workflows (Hybrid Extender)
    - Protection Based Workflow
    - Archival Based Workflow
    - Replication Based Workflow
- External Targets Using Hybrid Extender
    - AWS External Targets for CloudArchive - Standard
    - AWS External Targets for CloudArchive - C2S
    - AWS External Targets for CloudArchive - Gov
    - Microsoft Azure External Targets for CloudArchive
    - GCP External Targets for CloudArchive
    - Oracle External Targets for CloudArchive
    - S3 - Compatible, NAS, Tape External Targets for CloudArchive
- External Target for Incremental Forever Archival/CloudArchive Direct with Dedup (Hybrid Extender)
- External Target for Incremental Forever Archival/CloudArchive Direct without Dedup (Hybrid Extender)
    - AWS External Targets for CloudArchiveDirect - Standard
    - AWS External Targets for CloudArchiveDirect - C2S
    - AWS External Targets for CloudArchiveDirect - Gov
    - AWS External Targets for CloudArchiveDirect - GCP
    - AWS External Targets for CloudArchiveDirect - Oracle
    - AWS External Targets for CloudArchiveDirect - Azure
    - AWS External Targets for CloudArchiveDirect -S3 Compatible, NAS, Tape
- Supported Workflow for NAS Tiering (Hybrid Extender)

Legend for Charts Below:

Yes = Supported

No = Not Supported

N/A = Not Applicable

CA = Cluster Administrator

TA = Tenant Administrator

Create = Create for an Org

Delete = Delete from Cluster

Assign = Assign to an Org

Un-assign = Un-assign from an Org

Virtual Server (VMs) and Physical Server Workflows (Hybrid Extender)

Protection Based Workflow

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	N/A	N/A	N/A	N/A	N/A
	Auto-protect new objects	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CA,TA)
	Full backup with CBT	N/A	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CA,TA)
	Incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A	N/A	N/A
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	CloudSpin to AWS and Azure	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recovering an entity in the Job Run	No	No	No	No	No
	Recover an entire backup job run	Supported (CA,TA)	No	No	No	No
	Recover a SQL DB to the original server and instance	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A	N/A	N/A	N/A
	Recover DB to Alternate Instance with overwite	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A	N/A	N/A	N/A
	Instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No
	Crash-consistent recover for apps	Supported (CA,TA)	No	No	No	No
CLONE	Clone an object of a job run	Supported (CA,TA)	No	No	No	No
	Clone all the objects of a job run	Supported (CA,TA)	No	No	No	No
	Clone db to FCI instance	N/A	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A
TEARDOWN	Tear down objects	Supported (CA,TA)	No	No	No	No
	Tear down instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for a backed-up object	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download files/folders	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file	Supported (TA)	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	Supported (CA,TA)	N/A	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	Supported (CA,TA)	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	Supported (CA,TA)	N/A	N/A	N/A	N/A

Archival Based Workflow

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A	N/A	N/A
	Archive a full backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (CA,TA)	No	No	No	No
	Recover an entire backup job run from an archive to a source	Supported (CA,TA)	No	No	No	No
	Instant volume mount from an archive	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CA,TA)	No
	Crash-consistent recover for apps from an archive	Yes	No	No	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	No	No	No	No
	Clone all objects of a job run from an archive	Supported (CA,TA)	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	Supported (CA,TA)	No	No	No	No
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders from an archive	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for files or folders from an archive snapshot	No	No	No	Supported (CA,TA)	Supported (CA,TA)
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	No	Supported (CA,TA)	Supported (CA,TA)	No	No
	Download n VMX file from an archive	No	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	Supported (CA,TA)	N/A	N/A	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	No	No	No	No	No
	Clone all objects of a job run from an archive	No	No	No	No	No
	clone DB to FCI	N/A	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	No	No	No	No	No
	Recover all objects of job run from an archive to a source	No	No	No	No	No
	Instant volume mount from an archive	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Search files or folders from the metadata of a retrieved job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse files or folders from a retrieved snapshot	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse files or folders from the metadata of a downloaded job run	Supported (CA,TA)	No	No	No	No
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	No	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Download a VMX file from an archive	No	N/A	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	No	N/A	N/A	N/A	N/A

Replication Based Workflow

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	N/A	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	Supported (CA,TA)	No	No	No	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	Recover an entire backup job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A	N/A	N/A	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A	N/A	N/A	N/A
	Instant volume mount of replicated data	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
	Crash-consistent recover for apps from a remote cluster	Supported (CA,TA)	No	No	No	No
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A	N/A
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	Clone all objects of a job run from a remote cluster	Supported (CA,TA)	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A	No
	Crash-consistent clone for apps from a remote cluster	Supported (CA,TA)	No	No	No	

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	No	No	No	No
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	No
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for replicated object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover replicated files from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from a remote cluster	Supported (TA)	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	Supported (CA,TA)	N/A	N/A	N/A	N/A

Categories	Workflows	Sources				
		ESXi VM Servers	Physical Servers (Windows - Block Based)	Physical Servers (Windows - File Based)	Physical Servers (Linux - Block Based)	Physical Servers (Linux - File Based)
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	Supported (CA,TA)	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	Supported (CA,TA)	N/A	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A	N/A	N/A

Databases Workflows (Hybrid Extender)

Protection Based Workflow for MS SQL VM, FCI, and AG

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	Yes using VM edit workflow	Supported (CA,TA)	Yes using VM edit workflow
	Auto-protect new objects	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
BACKUP	Incremental backup with CBT	Yes as log-based incremental	Supported (CA,TA)	Supported (CA,TA)
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Incremental backup with SQL server differential	N/A	N/A	N/A
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A
	Backup, archive and replicate now	Supported at DB and server level both	Supported at DB and server level both	Supported at DB and server level both
	CloudSpin to AWS and Azure	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
RECOVER	Recovering an entity in the Job Run	Yes	Yes for DB and VM	Yes for DB and VM
	Recover an entire backup job run	Yes as crash-consistent VM not as SQL restore	No	Yes as crash-consistent VM not as SQL restore
	Recover a SQL DB to the original server and instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to the original server but a different instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover DB to Alternate Instance with overwite	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Crash-consistent recover for apps	Supported (CA,TA)	Supported (CA,TA)	Yes using IVM (physical)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLONE	Clone an object of a job run	Supported (CA,TA)	Cloning DB from FCI to standalone server / node Supported. cloning db to fci instance not supported (CA,TA)	Supported (CA,TA)
	Clone all the objects of a job run	Supported as Crash Consistent VM not as SQL Clone(CA,TA)	No	Supported as Crash Consistent VM not as SQL Clone(CA,TA)
	Clone db to FCI instance	N/A	N/A	N/A
	Crash-consistent recover for apps	Supported (CA,TA)	No	Supported using Instant volume mount (physical) (CA,TA)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
TEARDOWN	Tear down objects	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Tear down instant volume mount	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for a backed-up object	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download files/folders	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file	Supported (TA)	N/A	Supported (TA)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A

Archival Based Workflow for MS SQL VM, FCI, and AG

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
ARCHIVAL	Archive an incremental backup	Supported (Not for log backup)(CA,TA)	Supported (Not for log backup)(CA,TA)	Supported (Not for log backup)(CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A
	Archive a full backup	Supported (Not for log backup)(CA,TA)	Supported (Not for log backup)(CA,TA)	Supported (Not for log backup)(CA,TA)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (No Databases, Only Files and VM) (CA,TA)	Supported (No Databases, Only Files)(CA,TA)	Supported (No Databases, Only Files and VM) (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No	No
	Instant volume mount from an archive	No for SQL (use VM or physical)	No for SQL (use VM or physical)	No for SQL (use VM or physical)
	Crash-consistent recover for apps from an archive	Yes for files & VMs but not DBs	Yes for files but not DBs	Yes for files & VMs but not DBs
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLONE	Clone objects of a job run from an archive	No	No	No
	Clone all objects of a job run from an archive	No	No	No
	clone db to FCI	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	No	No	No
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	No	No	No
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
FILE RECOVERY	Search for files or folders from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for files or folders from an archive snapshot	No	No	Supported (CA,TA)
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	No	No	No
	Download n VMX file from an archive	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported for volume based job(CA)	Supported for volume based job(CA)	Supported for volume based job(CA)
	Download metadata for all runs of a job	Supported for volume based job(CA)	Supported for volume based job(CA)	Supported for volume based job(CA)
	Download a snapshot (archived data) for a job run	Supported for volume based job(CA)	Supported for volume based job(CA)	Supported for volume based job(CA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported for volume based job	Supported for volume based job	Supported for volume based job
	Clone all objects of a job run from an archive	No	No	No
	clone DB to FCI	N/A	N/A	N/A
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported for volume based job	Supported for volume based job	Supported for volume based job
	Recover all objects of job run from an archive to a source	No	No	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical, volume based)(CA,TA)	Not Supported for SQL (use VM or physical, volume based)(CA,TA)	Not Supported for SQL (use VM or physical, volume based)(CA,TA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported for volume based job	Supported for volume based job	Supported for volume based job
	Search files or folders from the metadata of a retrieved job run	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)
	Browse files or folders from a retrieved snapshot	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)
	Browse files or folders from the metadata of a downloaded job run	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)
	Download files and folders from an archive	Supported for volume based job(TA)	Supported for volume based job(TA)	Supported for volume based job(TA)
	Recover files from an archive to a source (to the same and different sources)	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)	Supported for volume based job(CA,TA)
	Download a VMX file from an archive	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Replication Based Workflow for MS SQL VM, FCI, and AG

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Yes as log-based incremental	Supported as Log based incremental (CA,TA)	Supported as Log based incremental (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	Yes only for DBs & not full server for physical	Yes only for DBs & not full server for physical
	Recover an entire backup job run from a remote cluster	Yes as crash-consistent, not as SQL restore	No	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount of replicated data	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Crash-consistent recover for apps from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Only DB and VM supported (CA,TA)
	Clone all objects of a job run from a remote cluster	Supported as Crash Consistent VM not as SQL Clone(CA,TA)	No	No
	clone db to FCI	N/A	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	Supported (CA,TA)	No	Only DB and VM supported (CA,TA)
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL VM	MS SQL FCI	MS SQL AG
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse for replicated object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Supported (TA)	Supported (TA)	Supported (TA)
	Recover replicated files from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from a remote cluster	Supported (TA)	N/A	Supported but only VM(TA)
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A

Protection Based Workflow for MS SQL Volume-base, File-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Auto-protect new objects	Supported (CA,TA)	Supported exclusion is not supported (CA,TA)	Supported exclusion is not supported (CA,TA)
BACKUP	Incremental backup with CBT	Yes as log-based incremental	Supported (CA,TA)	N/A
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)	N/A
	Incremental backup with SQL server differential	N/A	N/A	Yes
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	Supported (CA,TA)	Supported (CA,TA)
	Backup, archive and replicate now	Supported at DB and server level both	Supported at DB and server level both	Supported at DB and server level both
	CloudSpin to AWS and Azure	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
RECOVER	Recovering an entity in the Job Run	Yes for DBs but not full server	Yes at DB level	Yes at DB level
	Recover an entire backup job run	No	No	No
	Recover a SQL DB to the original server and instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to the original server but a different instance	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover DB to Alternate Instance with overwite	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount	Supported (CA,TA)	N/A	N/A
	Crash-consistent recover for apps	Yes using IVM (physical)	N/A	N/A
CLONE	Clone an object of a job run	Supported (CA,TA)	Supported (CA,TA)	No
	Clone all the objects of a job run	No	N/A	N/A
	Clone db to FCI instance	N/A	Supported (CA,TA)	N/A
	Crash-consistent recover for apps	No	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
DB Migration	SQL DB Migration without auto sync	N/A	Supported (CA,TA)	N/A
	SQL DB Migration with auto sync	N/A	Supported (CA,TA)	N/A
TEARDOWN	Tear down objects	Supported (CA,TA)	Supported (CA,TA)	N/A
	Tear down instant volume mount	Supported (CA,TA)	N/A	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported at server level (CA,TA)	Supported at server level (CA,TA)
FILE RECOVERY	Search for files or folders	Supported (CA,TA)	N/A	N/A
	Browse for a backed-up object	Supported (CA,TA)	N/A	N/A
	Download files/folders	Supported (TA)	N/A	N/A
	Recover files	Supported (CA,TA)	N/A	N/A
	Download a VMX file	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A

Archival Based Workflow for MS SQL Volume-base, File-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
ARCHIVAL	Archive an incremental backup	Supported (Not for log backup)(CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	Supported (CA,TA)
	Archive a full backup	Supported (Not for log backup)(CA,TA)	Supported (CA,TA)	Supported (CA,TA)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (No Databases, Only Files)(CA,TA)	Supported at DB level (CA,TA)	Supported at DB level (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No	No
	Instant volume mount from an archive	No for SQL (use VM or physical)	N/A	N/A
	Crash-consistent recover for apps from an archive	Yes for files but not DBs	N/A	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
CLONE	Clone objects of a job run from an archive	No	Yes	No
	Clone all objects of a job run from an archive	No	No	No
	clone db to FCI	N/A	Yes	N/A
	Crash-consistent clone for apps from an archive	No	N/A	N/A
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	No	Yes	Yes
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Yes	Yes

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
FILE RECOVERY	Search for files or folders from an archive	Supported (CA,TA)	N/A	N/A
	Browse for files or folders from an archive snapshot	Supported (CA,TA)	N/A	N/A
	Download files and folders from an archive	Supported (TA)	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	No	N/A	N/A
	Download n VMX file from an archive	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported for volume based job(CA)	No	No
	Download metadata for all runs of a job	Supported for volume based job(CA)	No	No
	Download a snapshot (archived data) for a job run	Supported for volume based job(CA)	No	No
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported for volume based job	No	No
	Clone all objects of a job run from an archive	No	No	No
	clone DB to FCI	N/A	Yes	N/A

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native (VDI)
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported for volume based job	No	No
	Recover all objects of job run from an archive to a source	No	No	No
	Instant volume mount from an archive	Not Supported for SQL (use VM or physical, volume based)(CA,TA)	No	No
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported for volume based job	N/A	N/A
	Search files or folders from the metadata of a retrieved job run	Supported for volume based job(CA,TA)	N/A	N/A
	Browse files or folders from a retrieved snapshot	Supported for volume based job(CA,TA)	N/A	N/A
	Browse files or folders from the metadata of a downloaded job run	Supported for volume based job(CA,TA)	N/A	N/A
	Download files and folders from an archive	Supported for volume based job(TA)	N/A	N/A
	Recover files from an archive to a source (to the same and different sources)	Supported for volume based job(CA,TA)	N/A	N/A
	Download a VMX file from an archive	N/A	N/A	N/A
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A

Replication Based Workflow for MS SQL Volume-base, File-based, and Native Backup

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Yes as log-based incremental	Supported (CA,TA)	N/A
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	Supported (CA,TA)
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No
RECOVER	Recover a single object of a job run from a remote cluster	Yes only for DBs & not full server for physical	Yes	Supported (CA,TA)
	Recover an entire backup job run from a remote cluster	No	No	No
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Instant volume mount of replicated data	Supported (CA,TA)	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	No	No	No

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Only DB is supported (CA,TA)	No
	Clone all objects of a job run from a remote cluster	No	No	No
	clone db to FCI	N/A	Supported (CA,TA)	N/A
	Crash-consistent clone for apps from a remote cluster	No	No	No
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Only DB is supported (CA,TA)	No
	Tear down instant volume mount on a remote cluster	Supported (CA,TA)	N/A	N/A
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources		
		MS SQL Volume Based	MS SQL File Based	MS SQL Native
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (CA,TA)	N/A	N/A
	Browse for replicated object on a remote cluster	Supported (CA,TA)	N/A	N/A
	Download replicated files and folders from a remote cluster	Supported (TA)	N/A	N/A
	Recover replicated files from a remote cluster	Supported (CA,TA)	N/A	N/A
	Download a VMX file from a remote cluster	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A

Protection Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	N/A	Supported
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Supported
	Refresh Application registration	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Supported
	Auto-protect new objects	No	No	N/A	Not Supported
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Full backup with CBT	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Incremental backup with SQL server differential	N/A	N/A	N/A	Not Applicable
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported
	Ability to run a log backup in parallel to a full backup.	No	No	Supported (CA,TA)	Supported
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)	Yes except archive	Supported
	CloudSpin to AWS and Azure	N/A	N/A	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	Supported (CA,TA)	N/A	Supported
	Recover an entire backup job run	No	No	Supported (CA,TA)	Not Supported
	Recover a SQL DB to the original server and instance	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A	Not Applicable
	Recover a SQL DB to a different server that has the same SQL version	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Recover DB to Alternate Instance with overwite	Supported (CA,TA)	N/A	N/A	Not Applicable
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	No	No	N/A	Not Applicable
	Instant volume mount	N/A	N/A	N/A	Not Applicable
	Crash-consistent recover for apps	N/A	N/A	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLONE	Clone an object of a job run	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Clone all the objects of a job run	No	No	N/A	Not Applicable
	Clone db to FCI instance	N/A	N/A	N/A	Not Applicable
	Crash-consistent recover for apps	N/A	N/A	N/A	Not Applicable
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	Not Applicable
	SQL DB Migration with auto sync	N/A	N/A	N/A	Not Applicable
TEARDOWN	Tear down objects	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Tear down instant volume mount	N/A	N/A	N/A	Supported
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (Cloning is not supported)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (Cloning is not supported) (CA,TA)	Supported
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
FILE RECOVERY	Search for files or folders	N/A	N/A	N/A	Not Applicable
	Browse for a backed-up object	N/A	N/A	N/A	Not Applicable
	Download files/folders	N/A	N/A	N/A	Not Applicable
	Recover files	N/A	N/A	N/A	Not Applicable
	Download a VMX file	N/A	N/A	N/A	Not Applicable
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	Not Applicable
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	Not Applicable

Archival Based Workflow for Oracle, SAP HANA, and SAP Sybase ASE

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A	Not Applicable
	Archive a full backup	Supported (CA,TA)	Supported (CA,TA)	No	Supported
RECOVER	Recover a single object of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Recover an entire backup job run from an archive to a source	No	No	No	Supported
	Instant volume mount from an archive	N/A	N/A	No	Not Applicable
	Crash-consistent recover for apps from an archive	N/A	N/A	N/A	Not Applicable
DB Migration	SQL DB Migration without auto sync	N/A	N/A	No	Not Applicable
	SQL DB Migration with auto sync	N/A	N/A	No	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Clone all objects of a job run from an archive	Not supported	Not supported	No	Not Applicable
	clone db to FCI	N/A	N/A	N/A	Not Applicable
	Crash-consistent clone for apps from an archive	N/A	N/A	N/A	Not Applicable
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Cancel a running clone or recover task that is retrieving from an archive	Not supported	Not supported	No	Supported (Cloning is not supported)
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	No	Supported

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
FILE RECOVERY	Search for files or folders from an archive	N/A	N/A	No	Not Applicable
	Browse for files or folders from an archive snapshot	N/A	N/A	No	Not Applicable
	Download files and folders from an archive	N/A	N/A	No	Not Applicable
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A	No	Not Applicable
	Download n VMX file from an archive	N/A	N/A	No	Not Applicable
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	No	Not supported	No	Not Supported
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)	No	Not Applicable
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)	No	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	No	Not Supported
	Clone all objects of a job run from an archive	No	Not supported	No	Not Supported
	clone DB to FCI	N/A	N/A	N/A	Not Supported
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Recover all objects of job run from an archive to a source	No	No	No	Supported
	Instant volume mount from an archive	N/A	N/A	No	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	N/A	N/A	No	Not Applicable
	Search files or folders from the metadata of a retrieved job run	N/A	N/A	No	Not Applicable
	Browse files or folders from a retrieved snapshot	N/A	N/A	No	Not Applicable
	Browse files or folders from the metadata of a downloaded job run	N/A	N/A	No	Not Applicable
	Download files and folders from an archive	N/A	N/A	No	Not Applicable
	Recover files from an archive to a source (to the same and different sources)	N/A	N/A	No	Not Applicable
	Download a VMX file from an archive	N/A	N/A	N/A	Not Applicable
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable

Replication Based Workflow Oracle, SAP HANA, and SAP Sybase ASE

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CBT is not applicable)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A	Not Applicable
	Replicate a full backup with CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Not Applicable
	Replicate a full backup without CBT to a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Supported (CBT is not applicable)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	N/A	Not Supported

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
RECOVER	Recover a single object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A	Supported
	Recover an entire backup job run from a remote cluster	Not supported	No	N/A	Not Supported
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	No	No	N/A	Not Applicable
	Instant volume mount of replicated data	N/A	N/A	N/A	Not Applicable
	Crash-consistent recover for apps from a remote cluster	N/A	N/A	N/A	Not Supported
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	Not Applicable
	SQL DB Migration with auto sync	N/A	N/A	N/A	Not Applicable

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
CLONE	Clone an object of a job run from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Applicable
	Clone all objects of a job run from a remote cluster	No	No	N/A	Not Applicable
	clone db to FCI	N/A	N/A	N/A	Not Applicable
	Crash-consistent clone for apps from a remote cluster	N/A	N/A	N/A	Not Applicable
TEARDOWN	Tear down object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	N/A	Not Supported
	Tear down instant volume mount on a remote cluster	N/A	N/A	N/A	Not Applicable
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	No	Supported
	Cancel a clone or recover task on a remote cluster	No	No	No	Supported (Clone not supported)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	No	Supported

Categories	Workflows	Sources			
		Oracle Standalone	Oracle RAC	SAP HANA	SAP Sybase ASE
FILE RECOVERY	Search for files or folders on a remote cluster	N/A	N/A	No	Not Applicable
	Browse for replicated object on a remote cluster	N/A	N/A	No	Not Applicable
	Download replicated files and folders from a remote cluster	N/A	N/A	No	Not Applicable
	Recover replicated files from a remote cluster	N/A	N/A	No	Not Applicable
	Download a VMX file from a remote cluster	N/A	N/A	N/A	Not Applicable
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	Not Applicable
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	Not Applicable
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	Not Applicable
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A	Not Applicable
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A	Not Applicable

NAS Workflows (Hybrid Extender)

Protection Based Workflow

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
REGISTRATION	Register a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh object hierarchy	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Refresh Application registration	N/A	N/A	N/A	N/A
	Auto-protect new objects	N/A	N/A	No	No
BACKUP	Incremental backup with CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Full backup with CBT	N/A	N/A	N/A	N/A
	Incremental backup with SQL server differential	N/A	N/A	N/A	N/A
	Full backup without CBT	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Ability to run a log backup in parallel to a full backup.	N/A	N/A	N/A	N/A
	Backup, archive and replicate now	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	CloudSpin to AWS and Azure	N/A	N/A	N/A	N/A

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
RECOVER	Recovering an entity in the Job Run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run	No	No	No	No
	Recover a SQL DB to the original server and instance	N/A	N/A	N/A	N/A
	Recover a SQL DB to the original server but a different instance	N/A	N/A	N/A	N/A
	Recover a SQL DB to a different server that has the same SQL version	N/A	N/A	N/A	N/A
	Recover DB to Alternate Instance with overwite	N/A	N/A	N/A	N/A
	Recover a SQL DB to a different server that has a different SQL version (Based on MS requirements)	N/A	N/A	N/A	N/A
	Instant volume mount	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A	N/A	N/A

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
CLONE	Clone an object of a job run	No	No	No	No
	Clone all the objects of a job run	No	No	No	No
	Clone db to FCI instance	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps	N/A	N/A	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A
TEARDOWN	Tear down objects	N/A	N/A	N/A	N/A
	Tear down instant volume mount	N/A	N/A	N/A	N/A
CANCEL	Cancel a running job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover tasks	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete a local snapshot of a job run	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
FILE RECOVERY	Search for files or folders	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Browse for a backed-up object	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download files/folders	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	N/A

Archival Based Workflow

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
ARCHIVAL	Archive an incremental backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Archiving an incremental backup with SQL server differential	N/A	N/A	N/A	N/A
	Archive a full backup	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
RECOVER	Recover a single object of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover an entire backup job run from an archive to a source	No	No	No	No
	Instant volume mount from an archive	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps from an archive	N/A	N/A	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from an archive	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A
	Crash-consistent clone for apps from an archive	N/A	N/A	N/A	N/A
CANCEL	Cancel a running archival task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a running clone or recover task that is retrieving from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete an archived snapshot of a job run from the external target	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
FILE RECOVERY	Search for files or folders from an archive	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Browse for files or folders from an archive snapshot	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Download files and folders from an archive	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover files from an archive to a source (to the same and different sources)	No	No	No	No
	Download n VMX file from an archive	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A
CLOUD RETRIEVE - SEARCH & RETRIEVE	Search for an object in an external target	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download metadata for all runs of a job	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
	Download a snapshot (archived data) for a job run	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
CLOUD RETRIEVE - CLONE	Clone objects of a job run from an archive	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Clone all objects of a job run from an archive	No	No	No	No
	clone DB to FCI	N/A	N/A	N/A	N/A
CLOUD RETRIEVE - RECOVER	Recover objects of a job run from an archive to a source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Recover all objects of job run from an archive to a source	No	No	No	No
	Instant volume mount from an archive	N/A	N/A	N/A	N/A

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
CLOUD RETRIEVE - FILE RECOVERY	Search files or folders from downloaded snapshot	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Search files or folders from the metadata of a retrieved job run	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Browse files or folders from a retrieved snapshot	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Browse files or folders from the metadata of a downloaded job run	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Download files and folders from an archive	Supported (Protection Run must be indexed)(TA)	Supported (Protection Run must be indexed)(TA)	Supported (Protection Run must be indexed)(TA)	Supported (Protection Run must be indexed)(TA)
	Recover files from an archive to a source (to the same and different sources)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from an archive	N/A	N/A	N/A	N/A

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
CLOUD RETRIEVE - VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A

Replication Based Workflow

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
REPLICATION	Replicate an incremental backup with CBT to a remote cluster	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)
	Replicate an incremental backup with SQL server differential to a remote cluster	N/A	N/A	N/A	N/A
	Replicate a full backup with CBT to a remote cluster	N/A	N/A	N/A	N/A
	Replicate a full backup without CBT to a remote cluster	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)	Supported (CBT is not applicable) (CA,TA)
FAILOVER	Failover an inactive Job on a remote cluster	No	No	No	No

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
RECOVER	Recover a single object of a job run from a remote cluster	Single Vol recovery supported (CA,TA)	Single Vol recovery supported (CA,TA)	Single Vol recovery supported (CA,TA)	Single Vol recovery supported (CA,TA)
	Recover an entire backup job run from a remote cluster	Single Vol recovery supported (CA,TA)	Single Vol recovery supported (CA,TA)	Single Vol recovery supported (CA,TA)	Single Vol recovery supported (CA,TA)
	Recover a SQL DB from a remote cluster to another SQL server with the same SQL version	N/A	N/A	N/A	N/A
	Recover a SQL DB from a remote cluster to different SQL server with a different SQL version	N/A	N/A	N/A	N/A
	Instant volume mount of replicated data	N/A	N/A	N/A	N/A
	Crash-consistent recover for apps from a remote cluster	N/A	N/A	N/A	N/A
DB Migration	SQL DB Migration without auto sync	N/A	N/A	N/A	N/A
	SQL DB Migration with auto sync	N/A	N/A	N/A	N/A

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
CLONE	Clone an object of a job run from a remote cluster	No	No	No	No
	Clone all objects of a job run from a remote cluster	No	No	No	No
	clone db to FCI	N/A	N/A	N/A	N/A
	Crash-consistent clone for apps from a remote cluster	N/A	N/A	N/A	N/A
TEARDOWN	Tear down object on a remote cluster	N/A	N/A	N/A	N/A
	Tear down instant volume mount on a remote cluster	N/A	N/A	N/A	N/A
CANCEL	Cancel a running replication task	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Cancel a clone or recover task on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
DELETE SNAPSHOT	Delete replicated snapshot of a job run on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Categories	Workflows	Sources			
		Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
FILE RECOVERY	Search for files or folders on a remote cluster	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)	Supported (Protection Run must be indexed) (CA,TA)
	Browse for replicated object on a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download replicated files and folders from a remote cluster	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)
	Recover replicated files from a remote cluster	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Download a VMX file from a remote cluster	N/A	N/A	N/A	N/A
VIRTUAL DISKS RECOVERY	Recover Virtual Disks	N/A	N/A	N/A	N/A
CONVERT AND DEPLOY	Register Azure source and clone VMs to Azure	N/A	N/A	N/A	N/A
	Register AWS source and clone VMs to AWS	N/A	N/A	N/A	N/A
CLOUD REPLICATION	Create a snapshot in AWS and replicate to another region	N/A	N/A	N/A	N/A
	Create a snapshot in Azure and replicate to another region	N/A	N/A	N/A	N/A

External Targets Using Hybrid Extender

AWS External Targets for CloudArchive - Standard

Workflows	External Targets					
	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone	AWS - Glacier (Legacy)	AWS - S3-Glacier	AWS - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	Yes	Yes	Yes	Yes	Yes	Yes
Enable/disable incremental archive at external target level	Yes	Yes	Yes	Yes	Yes	Yes
Enable encryption at external target level	Yes	Yes	Yes	Yes	Yes	Yes
Once enabled, disable encryption at external target level	No	No	No	No	No	No
Enable/disable compression at external target level	Yes	Yes	Yes	Yes	Yes	Yes
Delete external target	Yes	Yes	Yes	Yes	Yes	Yes
Enable secure archive at external target level	Yes	Yes	Yes	Yes	Yes	Yes
Once enabled, disable Secure Archive at external target level	No	No	No	No	No	No
Perform search for CloudRetrieve operations	Yes	Yes	Yes	Yes	Yes	Yes
Download jobs (CloudRetrieve)	Yes	Yes	Yes	Yes	Yes	Yes

Workflows	External Targets					
	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone	AWS - Glacier (Legacy)	AWS - S3-Glacier	AWS - S3 Glacier Deep Archive
Download files/folders from archive	Yes	Yes	Yes	Yes	Yes	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes	Yes	Yes	Yes
Recover files/folders to source	Yes	Yes	Yes	No	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes	No	No	No

AWS External Targets for CloudArchive - C2S

Workflows	External Targets	
	AWS C2S - S3	AWS C2S - Glacier
Enable/disable deduplication at external target level	Yes	Yes
Enable/disable incremental archive at external target level	Yes	Yes
Enable encryption at external target level	Yes	Yes
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	Yes	Yes
Delete external target	Yes	Yes
Enable secure archive at external target level	Yes	Yes

Workflows	External Targets	
	AWS C2S - S3	AWS C2S - Glacier
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	Yes	Yes
Download jobs (CloudRetrieve)	Yes	Yes
Download files/folders from archive	Yes	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes
Recover files/folders to source	Yes	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	No

AWS External Targets for CloudArchive - Gov

Workflows	External Targets	
	AWS Gov - S3	AWS Gov - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	Yes	Yes
Enable/disable incremental archive at external target level	Yes	Yes
Enable encryption at external target level	Yes	Yes
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	Yes	Yes
Delete external target	Yes	Yes
Enable secure archive at external target level	Yes	Yes
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	Yes	Yes

Workflows	External Targets	
	AWS Gov - S3	AWS Gov - S3 Glacier Deep Archive
Download jobs (CloudRetrieve)	Yes	Yes
Download files/folders from archive	Yes	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes
Recover files/folders to source	Yes	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	No

Microsoft Azure External Targets for CloudArchive

Workflows	External Targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Enable/disable deduplication at external target level	Yes	Yes	Yes	Yes
Enable/disable incremental archive at external target level	Yes	Yes	Yes	Yes
Enable encryption at external target level	Yes	Yes	Yes	Yes
Once enabled, disable encryption at external target level	No	No	No	No
Enable/disable compression at external target level	Yes	Yes	Yes	Yes
Delete external target	Yes	Yes	Yes	Yes
Enable secure archive at external target level	Yes	Yes	Yes	Yes
Once enabled, disable Secure Archive at external target level	No	No	No	No
Perform search for CloudRetrieve operations	Yes	Yes	Yes	Yes

Workflows	External Targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Download jobs (CloudRetrieve)	Yes	Yes	Yes	Yes
Download files/folders from archive	Yes	Yes	No	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes	No	Yes
Recover files/folders to source	Yes	Yes	No	Yes
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	Yes	No	Yes

GCP External Targets for CloudArchive

Workflows	External Targets		
	Google - Standard	Google - Nearline	Google - Coldline
Enable/disable deduplication at external target level	Yes	Yes	Yes
Enable/disable incremental archive at external target level	Yes	Yes	Yes
Enable encryption at external target level	Yes	Yes	Yes
Once enabled, disable encryption at external target level	No	No	No
Enable/disable compression at external target level	Yes	Yes	Yes
Delete external target	Yes	Yes	Yes
Enable secure archive at external target level	Yes	Yes	Yes
Once enabled, disable Secure Archive at external target level	No	No	No
Perform search for CloudRetrieve operations	Yes	Yes	Yes

Workflows	External Targets		
	Google - Standard	Google - Nearline	Google - Coldline
Download jobs (CloudRetrieve)	Yes	Yes	Yes
Download files/folders from archive	Yes	Yes	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes
Recover files/folders to source	Yes	Yes	Yes
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes

Oracle External Targets for CloudArchive

Workflows	External Targets	
	Oracle Standard	Oracle Archive
Enable/disable deduplication at external target level	Yes	Yes
Enable/disable incremental archive at external target level	Yes	Yes
Enable encryption at external target level	Yes	Yes
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	Yes	Yes
Delete external target	Yes	Yes
Enable secure archive at external target level	Yes	Yes
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	Yes	Yes
Download jobs (CloudRetrieve)	Yes	Yes

Workflows	External Targets	
	Oracle Standard	Oracle Archive
Download files/folders from archive	Yes	No
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	No
Recover files/folders to source	Yes	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	No

S3 - Compatible, NAS, Tape External Targets for CloudArchive

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Enable/disable deduplication at external target level	Yes	No	Yes
Enable/disable incremental archive at external target level	Yes	No	Yes
Enable encryption at external target level	Yes	Yes	Yes
Once enabled, disable encryption at external target level	No	No	No
Enable/disable compression at external target level	Yes	Yes	Yes
Delete external target	Yes	Yes	Yes
Enable secure archive at external target level	Yes	Yes	Yes
Once enabled, disable Secure Archive at external target level	No	No	No
Perform search for CloudRetrieve operations	Yes	No	Yes
Download jobs (CloudRetrieve)	Yes	No	Yes
Download files/folders from archive	Yes	No	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	No	Yes

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Recover files/folders to source	Yes	No	Yes
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	No	Yes

External Target for Incremental Forever Archival/CloudArchive Direct with Dedup (Hybrid Extender)

External Targets for CloudArchiveDirect

Workflows	External Targets						
	NAS	S3 - Compatible (E.g., Cohesity)	Azure Hot Blob	Azure Cool Blob	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone
Enable encryption at external target level	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Once enabled, disable encryption at external target level	No	No	No	No	No	No	No
Enable/disable compression at external target level	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Delete external target	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Enable secure archive at external target level	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Workflows	External Targets						
	NAS	S3 - Compatible (E.g., Cohesity)	Azure Hot Blob	Azure Cool Blob	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone
Once enabled, disable Secure Archive at external target level	No	No	No	No	No	No	No
Once enabled, going back to legacy archival	No	No	No	No	No	No	No
Perform search for CloudRetrieve operations	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Download jobs (CloudRetrieve)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Download files/folders from archive	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Recover files/folders to source	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes	Yes	Yes	Yes	Yes

External Target for Incremental Forever Archival/CloudArchive Direct without Dedup (Hybrid Extender)

AWS External Targets for CloudArchiveDirect - Standard

Workflows	External Targets				
	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone	AWS - Glacier	AWS - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	No	No	No	No	No
Enable/disable incremental archive at external target level	Yes	Yes	Yes	No	No
Enable encryption at external target level	Yes	Yes	Yes	No	No
Once enabled, disable encryption at external target level	No	No	No	No	No
Enable/disable compression at external target level	Yes	Yes	Yes	No	No
Delete external target	No	No	No	No	No
Enable secure archive at external target level	Yes	Yes	Yes	No	No
Once enabled, disable Secure Archive at external target level	No	No	No	No	No
Perform search for CloudRetrieve operations	Yes	Yes	Yes	No	No
Download jobs (CloudRetrieve)	Yes	Yes	Yes	No	No
Download files/folders from archive	Yes	Yes	Yes	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes	No	No
Recover files/folders to source	Yes	Yes	Yes	No	No

Workflows	External Targets				
	AWS - S3	AWS - S3-Intelligent	AWS - S3-One-Zone	AWS - Glacier	AWS - S3 Glacier Deep Archive
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	Yes	Yes	No	No

AWS External Targets for CloudArchiveDirect - C2S

Workflows	External Targets	
	AWS C2S - S3	AWS C2S - Glacier
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	No	No
Enable encryption at external target level	No	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	No	No
Delete external target	No	No
Enable secure archive at external target level	No	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	No	No
Download jobs (CloudRetrieve)	No	No
Download files/folders from archive	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	No	No
Recover files/folders to source	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	No	No

AWS External Targets for CloudArchiveDirect - Gov

Workflows	External Targets	
	AWS Gov - S3	AWS Gov - S3 Glacier Deep Archive
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	Yes	No
Enable encryption at external target level	Yes	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	Yes	No
Delete external target	No	No
Enable secure archive at external target level	Yes	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	Yes	No
Download jobs (CloudRetrieve)	Yes	No
Download files/folders from archive	Yes	No
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	No
Recover files/folders to source	Yes	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	No

AWS External Targets for CloudArchiveDirect - GCP

Workflows	External Targets	
	Google - Nearline	Google - Coldline
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	Yes	Yes
Enable encryption at external target level	Yes	Yes
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	Yes	Yes
Delete external target	No	No
Enable secure archive at external target level	Yes	Yes
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	Yes	Yes
Download jobs (CloudRetrieve)	Yes	Yes
Download files/folders from archive	Yes	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes
Recover files/folders to source	Yes	Yes
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	Yes

AWS External Targets for CloudArchiveDirect - Oracle

Workflows	External Targets	
	Oracle Standard	Oracle Archive
Enable/disable deduplication at external target level	No	No
Enable/disable incremental archive at external target level	Yes	No
Enable encryption at external target level	Yes	No
Once enabled, disable encryption at external target level	No	No
Enable/disable compression at external target level	Yes	No
Delete external target	No	No
Enable secure archive at external target level	Yes	No
Once enabled, disable Secure Archive at external target level	No	No
Perform search for CloudRetrieve operations	Yes	No
Download jobs (CloudRetrieve)	Yes	No
Download files/folders from archive	Yes	No
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	No
Recover files/folders to source	Yes	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	No

AWS External Targets for CloudArchiveDirect - Azure

Workflows	External Targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Enable/disable deduplication at external target level	No	No	No	No

Workflows	External Targets			
	Azure Hot Blob	Azure Cool Blob	Azure Archive	Azure Gov
Enable/disable incremental archive at external target level	Yes	Yes	No	No
Enable encryption at external target level	Yes	Yes	No	No
Once enabled, disable encryption at external target level	No	No	No	No
Enable/disable compression at external target level	Yes	Yes	No	No
Delete external target	No	No	No	No
Enable secure archive at external target level	Yes	Yes	No	No
Once enabled, disable Secure Archive at external target level	No	No	No	No
Perform search for CloudRetrieve operations	Yes	Yes	No	No
Download jobs (CloudRetrieve)	Yes	Yes	No	No
Download files/folders from archive	Yes	Yes	No	No
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	Yes	No	No
Recover files/folders to source	Yes	Yes	No	No
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	Yes	No	No

AWS External Targets for CloudArchiveDirect -S3 Compatible, NAS, Tape

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Enable/disable deduplication at external target level	No	No	No

Workflows	External Targets		
	NAS	QStar Tape	S3 - Compatible (E.g., Cohesity)
Enable/disable incremental archive at external target level	Yes	No	Yes
Enable encryption at external target level	Yes	No	Yes
Once enabled, disable encryption at external target level	No	No	No
Enable/disable compression at external target level	Yes	No	Yes
Delete external target	No	No	No
Enable secure archive at external target level	Yes	No	Yes
Once enabled, disable Secure Archive at external target level	No	No	No
Perform search for CloudRetrieve operations	Yes	No	Yes
Download jobs (CloudRetrieve)	Yes	No	Yes
Download files/folders from archive	Yes	No	Yes
Download files/folders from archive for downloaded CloudRetrieve jobs	Yes	No	Yes
Recover files/folders to source	Yes	No	Yes
Recover files/folders to source from archive for downloaded CloudRetrieve jobs	Yes	No	Yes

Supported Workflow for NAS Tiering (Hybrid Extender)

NAS Tiering

Workflows	External Targets			
	Generic NAS - NFS	Generic NAS - SMB	Isilon - NFS	Isilon - SMB
Down-Tiering without leaving anything at source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Downtiering with leaving the symlink behind at the source	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Orphan data delete	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Capacity Planner	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Tiering Audit Log	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Bucket Customization	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Supported Cloud Spin and Other Workflows

This sections lists the supported Cloud Spin and other supported workflow:

- Cloud Spin Workflows
- **Other Supported Workflows**
  - Cohesity Data Cloud (Self-Managed)
  - User Interface
  - Replication
  - View Tenant-ID
  - Authentication
  - Authorization
  - Supported Cloud Spin and Other Workflows
  - RBAC

- Storage Domain Settings
- View Workflows

Legend for Charts Below:

Yes = Supported

No = Not Supported

N/A = Not Applicable

CA = Cluster Administrator

TA = Tenant Administrator

Create = Create for an Org

Delete = Delete from Cluster

Assign = Assign to an Org

Un-assign = Un-assign from an Org

Cloud Spin Workflows

AWS/AWS Gov Workflows

Workflows	VMware	OS Supported by VMware	HyperV 2016 VM Servers	HyperV 2012 VM Servers
Full cloud spin (Policy-based)	Supported (CA,TA)	Windows 10 Desktop Edition, Microsoft Windows Server 2012, Microsoft Windows Server 2016, and Microsoft Windows Server 2019, Windows Server 2022 CentOS 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit RHEL 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit SUSE Linux Enterprise Server 11 64 bit, 12 64 bit Ubuntu 14.x 64-bit, 16 64-bit, 18 64-bit, 20 64-bit OEL 5.11 - 5.8 64-bit, 6.x 64-bit, 7.0 - 7.9 64-bit Debian 6.0.0-6.0.8 64 bit, 7.0.0-7.8.0 64 bit, 8.0.0 64 bit Fedora Server 19-21 64 bit	No	No
Launch instance	Supported (CA,TA)		No	No
Tear-down launched instance (Full)	Supported (CA,TA)		No	No
Incremental Cloudspin (Policy-based)	Supported (CA,TA)		No	No
Launch instance (Incremental)	Supported (CA,TA)		No	No
Tear-down launched instance (Incremental)	Supported (CA,TA)		No	No
On-demand cloudspin / DR Failover from replication cluster	Supported (CA,TA)		No	No
Tear-down deployed instance	Supported (CA,TA)		No	No
DR Failback using Native Snapshot	Supported (CA,TA)		No	No
DR Failback to using Cohesity Protection Service	Supported (CA,TA)		No	No

Azure Workflows

Workflows	VMware	OS Supported by VMware	HyperV 2016 VM Servers	HyperV 2012 VM Servers	OS Supported by Hyper-V
Full cloud spin (Policy-based)	Supported (CA,TA)	Windows Server 2016 Core, Windows Server 2019 Core, Windows 10 Desktop Edition, Windows 2022 64-bit, Windows 2019 64-bit, Windows 2016 64-bit, Windows 2012, 2012 R2 64-bit CentOS 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit RHEL 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit	Supported (for Gen1 and Gen2 VMs) (CA,TA)	Supported (for Gen1 and Gen2 VMs) (CA,TA)	Windows Server 2016 Core, Windows Server 2019 Core, Windows 10 Desktop Edition, Windows 2022 64-bit, Windows 2019 64-bit, Windows 2016 64-bit, Windows 2012, 2012 R2 64-bit CentOS 5.10 64-bit, 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit RHEL 5.8 - 5.11 64-bit, 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit Ubuntu 14.x 64-bit, 16.x 64-bit, 18.x 64-bit, 19.x 64-bit, 20.x 64-bit
Launch instance	Supported (CA,TA)		Supported (for Gen1 and Gen2 VMs) (CA,TA)	Supported (for Gen1 and Gen2 VMs) (CA,TA)	
Tear-down launched instance (Full)	Supported (CA,TA)		Supported (for Gen1 and Gen2 VMs) (CA,TA)	Supported (for Gen1 and Gen2 VMs) (CA,TA)	
Incremental Cloudspin (Policy-based)	No		No	No	
Launch instance (Incremental)	No		No	No	
Tear-down launched instance (Incremental)	No		No	No	
On-demand cloudspin / DR Failover from replication cluster	Supported (CA,TA)		Supported (for Gen1 and Gen2 VMs) (CA,TA)	Supported (for Gen1 and Gen2 VMs) (CA,TA)	
Tear-down deployed instance	Supported (CA,TA)		Supported (for Gen1 and Gen2 VMs) (CA,TA)	Supported (for Gen1 and Gen2 VMs) (CA,TA)	

Workflows	VMware	OS Supported by VMware	HyperV 2016 VM Servers	HyperV 2012 VM Servers	OS Supported by Hyper-V
DR Failback using Native Snapshot	No		No	No	
DR Failback to using Cohesity Protection Service	Supported (CA,TA)		No	No	

GCP Workflows

Workflows	VMware	HyperV 2016 VM Servers	HyperV 2012 VM Servers
Full cloud spin (Policy-based)	No	No	No
Launch instance	No	No	No
Tear-down launched instance (Full)	No	No	No
Incremental Cloudspin (Policy-based)	No	No	No
Launch instance (Incremental)	No	No	No
Tear-down launched instance (Incremental)	No	No	No
On-demand cloudspin / DR Failover from replication cluster	No	No	No
Tear-down deployed instance	No	No	No
DR Failback using Native Snapshot	No	No	No
DR Failback to using Cohesity Protection Service	No	No	No

Cohesity Data Cloud (Self-Managed)

Category	Workflows	Cluster Admin	Tenant Admin
Cluster Configuration	Add/Edit DNS servers	Supported	Not Applicable
	Add/Edit NTP servers	Supported	Not Applicable
	Add/Edit Domain Names	Supported	Not Applicable
	Add/Edit Web Proxy	Supported	Not Applicable
	Add/Edit Failures Tolerated	Supported	Not Applicable
	Change system admin password	Supported	Not Applicable
	Enable organization	Supported	Not Applicable
	Local Storage Utilization Threshold	Supported	Not Applicable
	Storage utilization threshold	Not Applicable	Not Applicable
	Add/Edit System Admin Email Address	Supported	Supported
	Enable SMTP	Supported	Not Applicable
	Enable Two factor authentication	Supported	Not Applicable
	Local Groups	Supported	Not Applicable

Category	Workflows	Cluster Admin	Tenant Admin
	Enable/Disable Local Documentation	Supported	Not Applicable
	Enable/Disable Support Channel (RT access enable/disable)	Supported	Not Applicable
Edit/Update Cluster Partition	Partition - Name/Hostname/VIPs	Not Applicable	Not Applicable
	Partition - VLAN and VLAN VIPs	Not Applicable	Not Applicable
	Partition - Custom Host Mapping	Not Applicable	Not Applicable
	Partition - Add Node	Not Applicable	Not Applicable
Key Management Server	Add/Edit Key Management Server	Supported	Not Applicable
Upgrade	Using Provided Download URL	Supported	Not Applicable
	By Uploading a package file	Supported	Not Applicable
	Dark-site Upgrade	Supported	Not Applicable

Category	Workflows	Cluster Admin	Tenant Admin
VLAN	VLAN - Add/Edit/Delete VLAN & VLAN VIPs	Supported	Not Applicable
	Backup VM on VCenter over VLAN	Supported	Supported
	Recover VM on Same VLAN	Supported	Supported
	Recover File & Folder on same VLAN	Supported	Supported
	Recover Instant Volume Mount on same VLAN	Supported	Supported
Scale Out	Add Node (Single/Multiple)	Supported	Not Applicable
	VIPs addition during Add node	Supported	Not Applicable
	Node Status (Active/Unreachable/Removal)	Supported	Not Applicable
Alerts	Alerts Summary - Filter via status, category and date	Supported	Supported
	Alerts Summary - Marking alerts as resolved	Supported	Supported
	Resolution Summary - Read only details of resolved alerts	Supported	Supported
	Settings - Add/Edit Alerts Notification Rule	Supported	Supported

Category	Workflows	Cluster Admin	Tenant Admin
SNMP	Enable/Disable SNMP	Supported	Not Applicable
	Download MIB	Supported	Not Applicable
	Add Read User	Supported	Not Applicable
	Add Write User	Supported	Not Applicable
	Trap user (edit/update)	Supported	Not Applicable
	SNMP Veriosn (edit between v2 and v3)	Supported	Not Applicable
	Add/Edit/Update NMS or Hostname	Supported	Not Applicable

User Interface

Reports

Feature	Workflows		Cluster Admin	Tenant Admin
Reports	Available Local Snapshots Per Job	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported
	Agent Status Summary	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported
	Backup Summary	Report	Supported	Supported
		Filters	Supported	Supported
		Email	Supported	Supported
		CSV	Supported	Supported

Feature	Workflows		Cluster Admin	Tenant Admin
	Clone and Recover VM Tasks Summary	Report	Supported	Supported

Feature	Workflows		Cluster Admin	Tenant Admin
		Filters	Supported	Supported

Feature	Workflows		Cluster Admin	Tenant Admin
		Email	Not-Supported	Not-Supported
		CSV	Not-Supported	Not-Supported
	Cluster Wide Storage	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Not-Supported	Not-Supported
		CSV	Not-Supported	Not-Supported
	Data Transferred to External Targets	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported

Feature	Workflows		Cluster Admin	Tenant Admin
	Directory Quotas	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported
	Failed Objects	Report	Supported	Supported
		Filters	Supported	Supported
		Email	Supported	Supported
		CSV	Supported	Supported
	GDPR Object Summary	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported

Feature	Workflows		Cluster Admin	Tenant Admin
	Objects Protected by Multiple Protection Jobs	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Not-Supported	Not-Supported
		CSV	Not-Supported	Not-Supported
	Protected Objects Heatmap	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Not-Supported	Not-Supported
	Protection Details Per Object	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported

Feature	Workflows		Cluster Admin	Tenant Admin
	Protection Jobs Inventory and Schedule	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported
	Protection Runs Summary	Report	Supported	Supported
		Filters	Supported	Supported
		Email	Supported	Supported
		CSV	Supported	Supported
	Protection Summary by Object Type	Report	Supported	Supported
		Filters	Supported	Supported
		Email	Supported	Supported
		CSV	Supported	Supported
	Source Growth and Variance	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported

Feature	Workflows		Cluster Admin	Tenant Admin
	Sources Protected	Report	Supported	Not-Supported
		Filters	Not-Supported	Not-Supported
		Email	Not-Supported	Not-Supported
		CSV	Not-Supported	Not-Supported
	Storage Consumed by Backups	Report	Supported	Supported
		Filters	Supported	Supported
		Email	Supported	Supported
		CSV	Supported	Supported
	Storage Consumed by File Categories	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported

Feature	Workflows		Cluster Admin	Tenant Admin
	Storage Consumed by Servers	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported
	Storage Consumed by Storage Domains	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported
	Storage Consumed by Organizations	Report	Supported	Supported
		Filters	Supported	Supported
		Email	Supported	Supported
		CSV	Supported	Supported

Feature	Workflows		Cluster Admin	Tenant Admin
	Top Jobs	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Not-Supported	Not-Supported
		CSV	Not-Supported	Not-Supported
	Unprotected/Protected VMs	Report	Supported	Supported
		Filters	Supported	Supported
		Email	Supported	Supported
		CSV	Supported	Supported
	User Quota	Report	Supported	Not-Supported
		Filters	Supported	Not-Supported
		Email	Supported	Not-Supported
		CSV	Supported	Not-Supported

Feature	Workflows	Cluster Admin	Tenant Admin
Cluster Audit Logs	Verify that relevant audit logs show up	Supported	Supported
	Retention settings	Supported	Not-Supported
	Download CSV	Supported	Supported
	Filters	Supported	Supported
	Pagination	Supported	Supported

Replication

Config Flow	Options	Create	Read	Update	Delete
Remote Cluster		Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Cluster Connection	Nodes	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	User	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Password	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Network interface	Automatic	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
	Selecting Manually	Not Applicable	Not Applicable	Not Applicable	Not Applicable
Cluster Option	Remote Access	Not Applicable	Not Applicable	Not Applicable	Not Applicable
	Replication	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

Config Flow	Options	Create	Read	Update	Delete
Storage Domain Pairing		Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Distribute Load		Not Applicable	Supported (CA,TA)	Supported (CA,TA)	Not Applicable
Compression		Not Applicable	Supported (CA,TA)	Supported (CA,TA)	Not Applicable
Encryption		Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Global Rate Limit		Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Quiet Times		Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Bandwidth Throttling		Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)
Remote View-Overwrite		Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)

View Tenant-ID

Advanced Diags

Schemas	Cluster Admin	Tenant Admin
Cluster physical stats	Supported	Not Applicable
Cluster backup and restore stats	Supported	Not Applicable
Node physical stats	Supported	Not Applicable
Disk physical stats	Supported	Not Applicable
Apollo storage domain stats	Supported	Not Applicable
Apollo fault tolerance stats	Supported	Not Applicable

Schemas	Cluster Admin	Tenant Admin
Target replication stats	Supported	Not Applicable
Backup job repl stats	Supported	Not Applicable
Backup job target node repl stats	Supported	Not Applicable
Backup job target repl stats	Supported	Not Applicable

Reports

Schemas	Cluster Admin	Tenant Admin
Agent Status	Supported	Not Applicable
Available Local Snapshots	Supported	Not Applicable
Backup Summary	Supported	Supported
Clone and Recover VM Tasks Summary	Supported	Supported
Cluster Wide Storage	Supported	Not Applicable
Data Transferred To External Targets	Supported	Not Applicable
Failed Objects	Supported	Supported
GDPR objects Summary	Supported	Not Applicable
Objects protected by Multiple Protection Jobs	Supported	Not Applicable
Protection Details per Object	Supported	Not Applicable
Protection Jobs Inventory and Schedule	Supported	Not Applicable
Protection Runs Summary	Supported	Supported
Protection Summary by Object Type	Supported	Supported
Source Growth & Variance	Supported	Not Applicable
Sources Protected	Supported	Not Applicable

Schemas	Cluster Admin	Tenant Admin
Storage Consumed by Backups	Supported	Supported
Storage Consumed by File Categories	Supported	Not Applicable
Storage Consumed by Servers	Supported	Not Applicable
Storage Consumed by Storage Domains	Supported	Not Applicable
Storage Consumed byOrganizations	Supported	Supported
Top Jobs	Supported	Not Applicable
Unprotectd/Protected VMs	Supported	Supported
User Quotas	Supported	Not Applicable
Directory Quotas	Supported	Not Applicable
Protected Objects Heatmap	Supported	Not Applicable

> **Note:** None of these settings should be available for Tenant user for any CRUD operation as these are cluster level settings. Only setting which should be available is viewbox pairing.

Authentication

Cluster Admin

Active Directory		SMB Auth		NFS Auth	S3 Auth	Access Auth
		Kerberos	**NTLM**			
	No AD	Not Supported	Not Supported	Not Supported	Not Applicable	Not Supported
	Preferred DC	Supported	Supported	Not Supported	Not Applicable	Supported
	Exclude trusted domains	Supported	Supported	Not Supported	Not Applicable	Supported
	Ignore trusted domains	Supported	Supported	Not Supported	Not Applicable	Supported
	Set ID-Map config	Supported	Supported	Not Supported	Not Applicable	Not Applicable
	Disable Closest DC	Supported	Supported	Not Supported	Not Applicable	Supported
Active Directory Domains						
	Multiple Domains	Supported	Supported	Not Supported	Not Applicable	Supported
	Single Domain	Supported	Supported	Not Supported	Not Applicable	Supported
LDAP						
	Single/Multiple LDAP domains	Not Supported	Not Supported	Not Supported	Not Applicable	Not Applicable
Local Principals						

Active Directory		SMB Auth		NFS Auth	S3 Auth	Access Auth
	Local Users	Not Supported	Supported	Not Supported	Supported	Supported
	Local Groups	Not Supported	Not Supported	Not Supported	Not Applicable	Not Applicable

Tenant Admin

Active Directory		SMB Auth		NFS Auth	S3 Auth	Access Auth
		Kerberos	**NTLM**			
	No AD	Not Supported	Not Supported	Not Supported	Not Applicable	Not Supported
	Preferred DC	Supported	Supported	Not Supported	Not Applicable	Supported
	Exclude trusted domains	Supported	Supported	Not Supported	Not Applicable	Supported
	Ignore trusted domains	Supported	Supported	Not Supported	Not Applicable	Supported
	Set ID-Map config	Supported	Supported	Not Supported	Not Applicable	Not Applicable
	Disable Closest DC	Supported	Supported	Not Supported	Not Applicable	Supported
Active Directory Domains						
	Multiple Domains	Supported	Not Supported	Not Supported	Not Applicable	Supported
	Single Domain	Supported	Supported	Not Supported	Not Applicable	Supported

Active Directory		SMB Auth		NFS Auth	S3 Auth	Access Auth
LDAP						
	Single/Multiple LDAP domains	Not Supported	Not Supported	Not Supported	Not Applicable	Not Applicable
Local Principals						
	Local Users	Not Supported	Supported	Not Supported	Supported	Supported
	Local Groups	Not Supported	Not Supported	Not Supported	Not Applicable	Not Applicable

Authorization

Cluster Admin

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
		AD Principals	Local Principals	LDAP Principals	AD Principals	Local Principals	LDAP Principals	Local Principals

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
Native	None	Supported	Supported	Not Supported	Not Supported	Not Supported	Supported	Not Applicable
	Centrify	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Applicable
	RFC2307	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Applicable
	RID	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Applicable
	FiNot Supported ed	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Applicable
	LDAP	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Supported	Not Applicable

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
Unified	None	Supported	Supported	Not Supported	Not Supported	Not Supported	Supported	Not Applicable
	Centrify	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	RFC2307	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	RID	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	FiNot Supported ed	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	LDAP	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Supported	Not Applicable

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
NTFS	None	Supported	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Applicable
	Centrify	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	RFC2307	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	RID	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	FiNot Supported	Supported	Not Supported	Not Supported	Supported	Not Supported	Not Supported	Not Applicable
	LDAP	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Supported	Not Applicable

Tenant Admin

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
		AD Principals	Local Principals	LDAP Principals	AD Principals	Local Principals	LDAP Principals	Local Principals

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
Native	None	Supported	Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable
	Centrify	Supported	Not Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable
	RFC2307	Supported	Not Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable
	RID	Supported	Not Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable
	FiNot Supported	Supported	Not Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable
	LDAP	Not Supported	Not Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
Unified	None	Supported	Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable
	Centrify	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	RFC2307	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	RID	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	FiNot Supported	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	LDAP	Not Supported	Not Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable

View Security Style	AD ID-Map Type	SMB Authorization			NFS Authorization			S3 Authorization
NTFS	None	Supported	Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable
	Centrify	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	RFC2307	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	RID	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	FiNot Supported	Supported	Not Supported	Not Applicable	Supported	Not Supported	Not Supported	Not Applicable
	LDAP	Not Supported	Not Supported	Not Applicable	Not Supported	Not Supported	Not Supported	Not Applicable

Org Management

Property	Workflows					
	Create	Read	Update	Delete	Assign	Un-assign
Organization	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Not Supported	Not Supported
Storage Domain	Supported (CA)	Supported (CA)	Supported (CA)	Not Supported	Supported (CA)	Not Supported
View	Supported (CA,TA)	Supported (CA,TA)	Supported (TA)	Supported (TA)	Supported (CA)	Not Supported

Property	Workflows					
Shares	Supported (TA)	Supported (CA,TA)	Supported (TA)	Supported (TA)	Not Supported	Not Supported
Local User	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA)	Supported (CA)
Local Group	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
Role	Supported (CA)	Supported (CA,TA)	Supported (CA)	Supported (CA)	Not Supported	Not Supported
Protection Source	Supported (CA,TA)	Supported (CA,TA)	Supported (TA)	Supported (TA)	Supported (CA)	Not Supported
Protection Source Sub-Tree	Not Supported	Supported (CA,TA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
Protection Policy	Supported (CA)	Supported (CA,TA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
Protection Job	Supported (CA,TA)	Supported (CA,TA)	Supported (TA)	Supported (TA)	Supported (CA)	Not Supported
VLAN	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)	Supported (CA)
AD	Supported (CA,TA)	Supported (CA,TA)	Supported (TA)	Supported (TA)	Supported (CA)	Not Supported
LDAP	Supported (CA,TA)	Supported (CA,TA)	Supported (TA)	Supported (TA)	Supported (CA)	Not Supported
AD User	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA)	Supported (CA)
AD Group	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA,TA)	Supported (CA)	Supported (CA)

Property	Workflows					
Remote Cluster - Replication	Supported (TA)	Supported (TA)	Supported (TA)	Supported (TA)	Not Supported	Not Supported
Remote Cluster - Remote Access	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
External Target	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
SNMP	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
SSO	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported

RBAC

Cluster Admin

Principal Type	SMB	NFS	S3	Cluster Management Roles
Local Principals	Supported	Not Supported	Not Supported	Supported
AD Principals	Supported	Not Supported	Not Supported	Supported
LDAP Principals	Not Supported	Not Supported	Not Supported	Not Supported

Tenant Admin

Principal Type	SMB (SMB privileges are handled as Roles)	NFS	S3	Cluster Management Roles
Local Principals	Supported	Not Supported	Not Supported	Supported
AD Principals	Supported	Not Supported	Not Supported	Supported
LDAP Principals	Not Supported	Not Supported	Not Supported	Not Supported

Storage Domain Settings

Cluster Admin

Storage Domain Settings	Sub-Options	SMB	NFS	S3
Authentication Provider	AD	Supported	Not Applicable	Not Applicable
	LDAP	Supported	Not Applicable	Not Applicable
	AD + LDAP	Supported	Not Applicable	Not Applicable
Physical Quota & Alerts		Supported	Supported	Not Applicable
Logical Quotas		Supported	Supported	Not Applicable
Cloud Tier	Down Tier data to cloud	Supported	Supported	Supported

Tenant Admin

Storage Domain Settings	Sub-Options	SMB	NFS	S3
Authentication Provider	AD	Supported	Not Applicable	Not Applicable
	LDAP	Supported	Not Applicable	Not Applicable
	AD + LDAP	Supported	Not Applicable	Not Applicable
Physical Quota & Alerts		Not Applicable	Not Applicable	Not Applicable
Logical Quotas		Not Applicable	Not Applicable	Not Applicable
Cloud Tier	Down Tier data to cloud	Not Applicable	Not Applicable	Not Applicable

Storage Domain Efficiency and Fault Tolerance

Cluster Admin

Storage efficiency				Security	Fault Tolerance	
Deduplication		Compression		Encryption	Replication Factor	Erasure Coding
Inline	Post-Process	Inline	Post-Process			
Supported	Supported	Supported	Supported	Supported	Supported	Supported

Tenant Admin

Storage efficiency				Security	Fault Tolerance	
Deduplication		Compression		Encryption	Replication Factor	Erasure Coding
Inline	Post-Process	Inline	Post-Process			
Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported

View Workflows

Cluster Admin

Workflows	Description	All Protocols	SMB	NFS	S3
QoS Policy	Edit QoS Policy (independant of protocol)	Not Supported	Not Supported	Not Supported	Not Supported
Case Sensitive	Cannot be disabled after view create	Supported	Supported	Supported	Supported

Workflows	Description	All Protocols	SMB	NFS	S3
Security Settings	Override Global Allowlist	Supported	Supported	Supported	Not applicable
	SMB Share Owner	Supported	Supported	Not Supported	Not Supported
	View Permissions for Local or AD Principals (Allow/Disallow access to files/folders)	Supported	Supported	Not Supported	Not Supported
	Supported Security Style (Native, Unified, NTFS)	Supported	Not Supported	Not Supported	Not Supported
Dedupe & Compression	Override/Inherit from Storage domain	Supported	Not Supported	Not Supported	Not Supported
Quota	Inherit Quotas from Storage Domain	Supported	Not Supported	Not Supported	Not Supported
	Override Logical or Alerts Quota	Supported	Not Supported	Not Supported	Not Supported
File Datalock	Set Datalock on files (Can be disabled only by user with DataSecurity role)	Supported	Not Supported	Not Supported	Not Supported
File System/File Filtering	Filter filtering based on file types (independent of protocol)	Not Supported	Supported	Supported	Not Supported
NFS Options	Browsable NFS Shares (using showmount)	Supported	Not Supported	Supported	Not Supported

Workflows	Description	All Protocols	SMB	NFS	S3
SMB Options	Browsable SMB Shares (using net use)	Supported	Supported	Not Supported	Not Supported
	Access Based Enumeration	Supported	Supported	Not Supported	Not Supported
	SMB Encryption	Not Supported	Not Supported	Not Supported	Not Supported
Antivirus	Enable/Disable Antivirus protection on a view	Supported*	Supported	Not Supported	Not Supported
Protection Policies	View protection policies	Supported	Supported	Supported	Supported
User Quotas	Enable/Disable User Quotas on a view	Supported*	Supported	Supported	Not Supported
	Set Quota policies for Users (Local/AD users)	Supported	Supported	Supported	Not Supported
View Aliases/ Shares		Supported	Supported	Supported	Not Supported

Tenant Admin

Workflows	Description	All Protocols	SMB	NFS	S3
QoS Policy	Edit QoS Policy (independant of protocol)	Not Applicable	Not Applicable	Not Applicable	Not Applicable
Case Sensitive	Cannot be disabled after view create	Supported	Supported	Supported	Not Applicable

Workflows	Description	All Protocols	SMB	NFS	S3
Security Settings	Override Global Allowlist-deprioritize	Supported	Supported	Supported	Not Applicable
	SMB Share Owner	Supported	Supported	Not Supported	Not Supported
	View Permissions for Local or AD Principals (Allow/Disallow access to files/folders)	Supported	Supported	Not Supported	Not Supported
	Supported Security Style (Native, Unified, NTFS)	Supported	Not Supported	Not Supported	Not Supported
Dedupe & Compression	Override/Inherit from Storage domain	Supported	Not Supported	Not Supported	Not Supported
Quota	Inherit Quotas from Storage Domain	Supported	Not Supported	Not Supported	Not Supported
	Override Logical or Alerts Quota	Supported	Not Supported	Not Supported	Not Supported
File Datalock	Set Datalock on files (Can be disabled only by user with DataSecurity role)	Supported	Not Supported	Not Supported	Not Supported
File System/File Filtering	Filter filtering based on file types (independent of protocol)	Not Supported	Supported	Supported	Not Supported
NFS Options	Browsable NFS Shares (using showmount)	Supported	Not Supported	Supported	Not Supported

Workflows	Description	All Protocols	SMB	NFS	S3
SMB Options	Browsable SMB Shares (using net use)	Supported	Supported	Not Supported	Not Supported
	Access Based Enumeration	Supported	Supported	Not Supported	Not Supported
	SMB Encryption	Not Supported	Not Supported	Not Supported	Not Supported
Antivirus	Enable/Disable Antivirus protection on a view	Supported*	Supported	Not Supported	Not Supported
Protection Policies	View protection policies	Supported	Supported	Supported	Supported
User Quotas De-Prioritize	Enable/Disable User Quotas on a view	Supported*	Supported	Supported	Not Supported
	Set Quota policies for Users (Local/AD users)	Supported	Supported	Supported	Not Supported
View Aliases/ Shares		Supported	Supported	Supported	Not Supported

* = When all protocols are enabled, S3 is read-only mode and not
** = Yes for either SMB or NFS protocol

Tenant-Level View Allowlist

In a multitenant environment, a tenant can add subnets to the allowlist and only the IP addresses within this subnet can access all the views of the tenant. This feature enables the tenant to manage all endpoints that have access to their views seamlessly.

Points to Note

- Tenant-level allowlist complements global allowlist and view-level allowlist.
- The tenant administrator must have the **Manage Cohesity Views** privilege. For more information, see Privileges for Service Provider Admin and Organization Admin.

To allowlist views at the tenant level, complete the following steps:

1. In the Cohesity Dashboard, click the user icon ( ) in the top-right corner and select **Impersonate**.
2. In the **Impersonate** pop-up window, select the required **Organization** and click **Switch**.
3. Select **File Services > Views > Global Settings**.
4. In the **Add Allowlist** pop-up window, enter a subnet IP in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).
5. Specify NFS and SMB permissions.
6. Specify whether NFS root squash must be enabled.
7. Optionally, enter a description for the allowlist.
8. Click **Add**.
9. Repeat the preceding steps for specifying additional subnets at the tenant level.

   For more information about adding subnets to allowlist, see Add Subnets to Allowlists.

Key Management Service (KMS) in Multitenancy

In a multitenancy (Organizations) environment, each tenant must use its own Key Management Service (KMS). This topic describes the required KMS configuration for tenants and the configuration that is not supported.

A KMS manages the encryption keys used to protect tenant data. In a multitenancy environment, every tenant is linked to a single vault, and that vault is linked to a unique KMS server. This one-to-one mapping keeps each tenant's encryption keys separate from every other tenant's keys.

This model applies to any external KMS, including cloud KMS such as GCP, Azure, or AWS.

Assign a Unique KMS Server to Each Tenant

Attach each tenant to its own KMS server and vault. Assigning a separate KMS server per tenant is a strict requirement. It allows tenant-level data partitioning and supports migrating a tenant in the future.

- When a tenant is provisioned, a new KMS server is created for that tenant, a new vault is attached to the KMS server, and both are assigned to the tenant.
- When you register a new tenant, attach it to a new KMS server. Do not reuse an existing tenant's KMS server for a new tenant.
- Each tenant's vault is linked to exactly one KMS server, and each KMS server serves exactly one tenant.

> **Important:** Attaching multiple tenants to the same KMS server is not a supported configuration. It is unsupported both from a security perspective and from a migration and availability perspective.

Chapter 8: Manage Protection Policy

A protection policy is a reusable set of settings that define how and when objects are protected, replicated and archived. You can select a standard policy to use when configuring a Protection Group. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived or replicated.

For instructions on creating and managing policies, see Create or Edit a Standard Policy.

Some settings are defined in the Protection Group and some settings are defined in the policy. For example, for a Protection Group that protects servers, the schedule is defined at the policy level but indexing is defined at the Protection Group level as shown in the following figure.

Policies are reusable. You can define a single policy that is used by different Protection Groups as shown in the following figure.

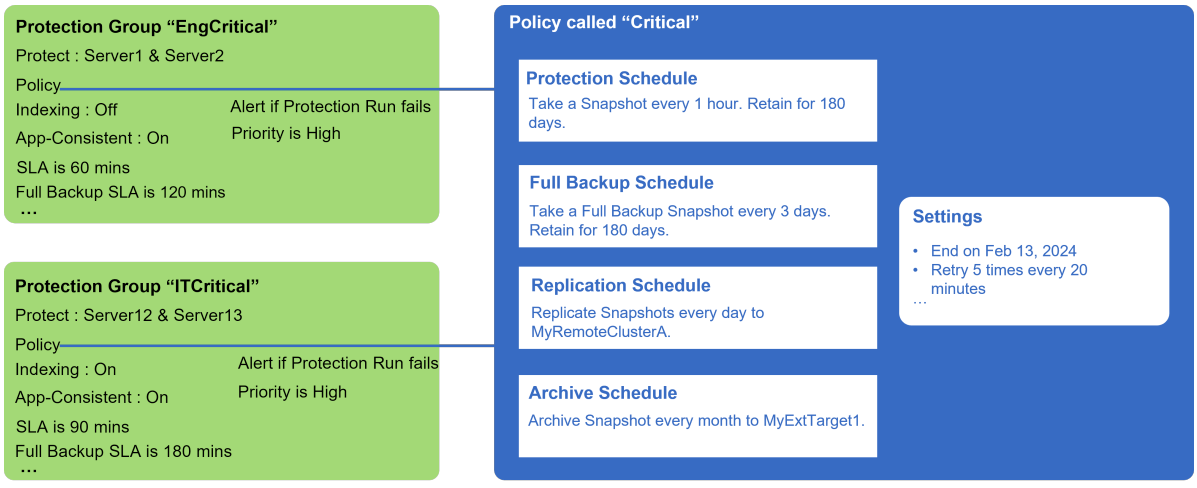

Create or Edit a Standard Policy

A protection policy is a reusable set of settings that define how and when objects are protected, replicated and archived. You can select a standard policy to use when configuring

a Protection Group. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived or replicated.

Use the default standard policies or create your own custom policies. Policies save time because you do not need to repetitively enter settings.

> **Note:** Policy creation and editing operations treat a month as having 31 days instead of 30 days, which was the definition in releases older than 6.8.1.

To create or edit a standard policy:

On the **Create Policy** page, select the target type and configure the advanced settings by clicking the respective icons from the floating menu when creating a policy.

On the **Create Policy** page you can change underlined numeric fields by over typing them or using the up and down arrow keys on your keyboard.

1. Navigate to **Data Protection > Policies**.
2. Click **Create Policy** located at the top right of the page..

   To edit an existing policy in the list, click the actions menu (⋮) and select **Edit** (✎).

   > **Note:** For Cohesity 6.6 and higher versions, the policy **Description** cannot be set or edited. For more details on how to edit the field, see the Unable to update the Description field in a Protection Policy KB article.

3. **Policy Name:** Specify a protection policy name. The name can contain alphanumerics, underscores, hyphens, periods, and spaces. This field is required and can be changed later.
4. Specify when backups are captured:
   - **Backup every xx**: Specify the minutes, hours, weeks, months, or years to set the frequency for how often backups are captured by the Protection Group. This defines the protection schedule for capturing backups.

     The policy schedule can be configured to run on specific days. From the **Every** drop-down, select **Week** and in the **On** field, select the days of the week you want to run the backup. For example, you can schedule a backup to run on Monday, Wednesday, and Friday, and not on Tuesday and Thursday.

> **Note:** The policy schedule can be configured to run on specific dates as well (calendar-based scheduling). From the **Every** drop-down, select **Month**, and in the **On** field, select the **Date**. You can then choose the date of the month you want to run the backup. For example, you can schedule a backup to run on the 5th of every month.

5. Under **Primary Copy**, specify the following:
   a. **Keep On**: Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster.

   > **Note:** If you are selecting any external targets, then this policy can be only used for performing CloudArchive Direct wherein the primary copy of the data is stored on the external targets. For more information, see CloudArchive Direct.

   b. **Retain for xx days**: Specify how many days, weeks, months, or years backups are retained on the Cohesity cluster before they are deleted.

   > **Note:**
   > - The **Retain** option can be set to a maximum value of 365000.
   > - Retentions at 7-day boundaries turn into weeks.

6. **Optional**: Enable **Data Movement** if required.

   Data Movement: With this option, you can move the data (downtier) from the selected AWS Standard Storage Tier to an AWS Archive Storage Tier target after a specified time period. Data Movement helps in space reclamation on the external target. You can select the S3 archive target in the **Move Data** section and set a time interval to transfer the data from the S3 standard tier to the S3 archive bucket. This interval can be defined in days, weeks, months, or years.

   > **Note:** As an AWS consideration, you cannot downtier from S3 Glacier to S3 Deep Archive. After downtiering to either S3 Glacier or S3 Deep Archive, you cannot perform any further downtiering.

7. Click **More Options** to configure the additional policy settings.

8. **DataLock:** Toggle on to lock the policy and any snapshots that belong to Protection Groups associated with the policy. DataLock is typically used for compliance and regulatory purposes. For more information, see DataLock.

   If you have enabled the **DataLock** option at the policy level, then you have the option to lock the backup retained on the Cohesity cluster. In the **Lock** option field available in the **Primary Copy** pane, specify how many days, weeks, months, or years the retained backups must be locked. You cannot delete the snapshots during the specified lock period.

9. **Periodic Full Backup**: Click the **Periodic Full Backup** icon from the floating menu to add a periodic full backup. You can schedule a policy to capture the periodic full backup for:
   - **Day** - Captures only the first successful run that started on that particular day.
   - **Week** - From the weekly calendar, select the required days of the week to run the backup.
   - **Month** - Use the **On** dropdown to select when the backup should run:
     - **First** - First selected day of the month.
     - **Second** - Second selected day of the month.
     - **Third** - Third selected day of the month.
     - **Fourth** - Fourth selected day of the month.
     - **Last** - Last selected day of the month.

       Then, choose the specific day of the week from the adjacent dropdown.
     - **Date** - From the pre-built calendar, select a required date for the backup.
   - **Year** - From the **On** dropdown, select one of the following:
     - **First Day** - First day of the year.
     - **Last Day** - Last day of the year.

   You can also schedule multiple periodic full backups with different frequencies and dates.

Note:

- The policy schedule can be configured to run on specific dates as well (calendar-based scheduling). From the **Every** drop-down, select **Month**, and in the **On** field, select the **Date**. You can then choose the date of the month you want to run the backup. For example, you can schedule a backup to run on the 5th of every month.
- If you specify multiple periodic full backup schedules for the same set of snapshots with different retention periods, the snapshots are retained for the longest specified period.
- Cohesity recommends a first full and incremental forever approach to back up all workloads (sources) with the exception for Oracle (SBT based backup) and SQL (VDI-based) where Cohesity recommends periodic full backups.

10. **Continuous Data Protection**: (*Applies only for VMware*) To achieve near zero RPO, click the **Continuous Data Protection** option floating menu and specify the time unit (hours) to capture and retain logs on the Cohesity cluster. This value defines the CDP schedule for capturing logs from the VMs.

    You can use this option to restore the VM applications to a point-in-time as opposed to periodic snapshots available with regular backups. For more information, see Cohesity CDP for VMware.

11. **Quiet Times:** Select this option from the **Backup Options** floating menu to add a quiet time. Quiet Times define time periods when new protection runs are not started. For example, you may want to configure hourly backups that run during weekdays but not on Saturday or Sunday. Quiet Times only prevent new protection runs from starting during the specified time period. Protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet times. Similarly, for jobs set on an interval basis, such as incremental jobs running every 6 hours, the job will start immediately after the Quiet Time ends. If the quiet time window ends 6 or more hours after the start time of a protection job, that run is skipped. For example, if jobs are scheduled at 6 AM, 12 PM, 6 PM, and 12 AM, and a Quiet Time is set from 11:30 AM to 1:30 PM, the job originally planned for noon will instead commence at 1:31 PM. Once the run commences at 1:31 PM, the next run will take place at 7:30 PM and not 6 PM (6 hours difference).

You can add multiple quiet times. For example, you could define a quiet time period for Saturday and Sunday and another quiet time period every morning from 9:00 AM to 11:00 AM. The quiet time window is considered based on the timezone specified in the protection group to which the protection policy is applied. For example, consider a policy defining a quiet time from 9 AM to 11 AM. There are two protection jobs: one in the IST timezone and the other in the PST timezone having the same policy. For the first job (in IST), protection runs will not be triggered during the 9 AM to 11 AM IST window. However, for the second job (in PST), protection runs can still be triggered because the 9 AM to 11 AM quiet time does not overlap with the same hours in the PST timezone, despite both jobs having the same policy.

> **Note:** If you update the Quiet Time periods in a policy for a running Protection Group, only new protection runs using the policy are affected and any existing Protection runs that have already started will use the original Quiet Time periods.

12. **Retry Options:** Select **Customize Retries** from the **Backup Options** floating menu to customize the retry settings for capturing snapshots. By default, the Cohesity cluster attempts to capture Snapshots three times before the protection run fails. The default time between retries is five minutes. You can optionally customize the number of retries and how long to wait between each attempt.
13. **BMR Backup (Physical Server)**: Select **BMR Backup** from the **Backup Options** floating menu to provide a backup schedule and retention period for Bare Machine Recovery (BMR) system data on a physical server.
    a. **Backup every xx**: Select the time unit for how often BMR backups are captured by the Protection Group.
    b. **Retain for xx days**: Specify how many days the BMR backups are retained on the Cohesity cluster before they are deleted.

       > **Note:** The **Retain** option can be set to a maximum value of 365000.

    c. **Lock**: Specify how many days, weeks, months, or years the retained BMR backups should be locked. You cannot delete the snapshots during the specified period.

> **Note:** This option is available only if you have enabled the **DataLock** option at the policy level.

14. **Log Backup (Databases):** Select **Log Backup** from the **Backup Options** floating menu to add a database logs schedule if you plan on restoring databases to a specific point in time between two full database server backups. If this policy has an archival schedule, this retention period also applies to the archived logs.

    For example, if you back up hourly at 5 minutes after the hour but you want the ability to restore a database to its state at 10:15 AM, specify a database logs schedule. The log backup saves all the transaction logs of all the databases so you can roll back to any point in time. Transaction logs are cumulative, so all the logs are backed up. When a database is added to a database server protected by Cohesity, its database logs are not backed up until a full database server backup has occurred.

    a. **Backup every xx:** Select the time unit for how often log backups are captured by the Protection Group.

    b. **Retain for xx days:** Specify how many days log backups are retained on the Cohesity cluster before they are deleted. If this policy has a Replication schedule, this retention period also applies to the logs replicated to the target cluster.

    > **Note:** The **Retain** option can be set to a maximum value of 365000.

    c. **Lock**: Specify how many days, weeks, months, or years the retained log backups should be locked. You cannot delete the snapshots during the specified period.

    > **Note:** This option is available only if you have enabled the **DataLock** option at the policy level.

    For SQL Protection Groups, if the log backups are configured in the Protection Group, all protection runs are replicated ignoring the set policy. For more information, see the Each SQL log backup run is replicated irrespective of policy KB article.

15. **Storage Array Snapshot**: Select **Storage Array Snapshot** from the **Backup Options** floating menu to provide a backup schedule and retention period for the Safeguarded snapshots for IBM Storage FlashSystem volume and volume groups. These are point-in-time cyber-resilient copies of IBM Storage FlashSystem volume groups.

    a. **Backup every xx**: Select the time unit for how often The Protection Group captures safeguarded snapshots.

    b. **Retain for xx days**: Specify how many days the Safeguarded snapshots are retained on the Cohesity cluster before they are deleted.

16. **Extended Retention**: Click the **Primary Copy** tile and select this option from the **Backup Options** floating menu to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
    a. Specify when to take the first snapshot and how long to retain it. For more information, see Extended Retention.
    b. **Lock**: Specify how many days, weeks, months, or years the retained backups should be locked. You cannot delete the snapshots during the specified period.

    > **Note:** This option is available only if you have enabled the **DataLock** option at the policy level.

    c. **Full Backup Only**: Toggle on to apply this setting for full backup. By default, the toggle is off and extended retention will not be used.

    > **Note:** You can enable **Full Backup Only** when you have a periodic full schedule in the protection policy.

17. **Add Replication:** Select this option to replicate snapshots to another Cohesity cluster. The copied snapshots are stored in the storage domain of the remote cluster. (The storage domain, however, is specified in the Protection Group.) The replication schedule for copying Snapshots cannot be more often than the protection schedule for capturing Snapshots. For example if Snapshots are captured hourly, the Cohesity cluster cannot replicate them every 30 minutes but it can replicate them hourly or any value over one hour.

    > **Note:** For CloudArchive Direct, you should not add Replication.

    If database logs are backed up, they will also be replicated (there is not a separate replication schedule for logs). The retention period specified in the **Log Backup (Databases)** setting applies to both the source and target cluster.
    a. **Replicate to:** Select either a **Remote Cluster** or **AWS**.

    From the drop-down, select a cluster source to replicate snapshots. You can also select an AWS S3 storage source if you are backing up AWS VMs using the Cloud Snapshot Manager APIs of the cloud service.

    > **Note:** If you choose a remote cluster, use the **Replication Target** to specify the target cluster to copy the snapshots.

You can also register a new cluster to replicate to by selecting **Register Remote Cluster**. For detailed instructions, see Create a Connection to the Remote Cohesity Cluster.

b. **After every run:** Specify how often snapshots are replicated. The available replication schedule frequencies are:
   - **Hour**: Replicates the first successful run that started on that particular hour.
   - **Daily**: Replicates the first successful run that started on that particular day.
   - **Weekly**: From the weekly calendar, select the required days in a week. You can also set the replication schedule for n-number of weeks.
   - **Monthly**: Use the **On** dropdown to select the replication schedule:
     - **First** - First selected day of the month.
     - **Second** - Second selected day of the month.
     - **Third** - Third selected day of the month.
     - **Fourth** - Fourth selected day of the month.
     - **Last** - Last selected day of the month.

       Then, choose the specific day of the week from the adjacent dropdown.
     - **Date** - From the pre-built calendar, select the required date.
   - **Yearly**: From the **On** dropdown, select the available options:
     - **First Day** - First day of the year.
     - **Last Day** - Last day of the year.
     - **Date** - From the pre-built calendar, select a required date of the year.

> **Note:** Log backups are always replicated whenever a log backup run is finished irrespective of the replication schedule.

c. **Retain for xx days:**

   Full/Incremental Backups

   Specify how long the snapshots are retained on the remote cluster.

   The Retain option can be set to a maximum value of 365000.

   Log Backups

   Specify how long the log snapshots are retained on the external target.

The Retain option can be set to a maximum value of 365000.

> **Note:** All runs are replicated continuously as log backups require a continuous chain of backups.

d. **Lock**: Specify how many days, weeks, months, or years the retained snapshots should be locked. You cannot delete the snapshots during the specified period.

> **Note:** This option is available only if you have enabled the **DataLock** option at the policy level.

e. To add an additional replication schedule, repeat these steps. To replicate to multiple Cohesity clusters, create multiple replication schedules with different replication targets.

> **Important:** If you specify multiple replication schedules with the same replication target that replicate the same set of snapshots but with different retention periods, the snapshots are retained for the longest time period. For example if you define two replication schedules with the same replication target that retain the same set of snapshots but one schedule specifies a 90-day retention period and other schedule specifies a 180-day retention period, the set of snapshots are retained for 180 days.

For replication, the extended retention schedule is as follows:

- **Weekly**: The first successful run that started on a Sunday or later.
- **Monthly**: The first successful run that started on the first day of the month or later.
- **Yearly**: The first successful run that started on the first day of the year or later.

18. **Add Archival:** Select this option to add an archive schedule.
    a. **Archive to:** Select an existing external target. You can also register a new external target. For detailed instructions, see Register an External Target.
    b. **Archive only fully successful Runs:** Check this box if you want archive to occur only if the snapshots of all the objects protected by the Protection Group were created. If the creation of any snapshots fails, no snapshots are archived.

c. **Full Backup Only**: Toggle on to apply this setting for full backup. By default, the toggle is off and extended retention will not be used.

> **Note:** You can enable **Full Backup Only** when you have a periodic full schedule in the protection policy.

d. **Every #:** Specify the archive schedule for making copies of snapshots created by this Protection Group and storing them on a registered external target. (The storage domain, however, is specified in the Protection Group.) The archive schedule for copying snapshots cannot be more often than the protection schedule for capturing snapshots. For example if snapshots are captured daily, the Cohesity cluster cannot archive them hourly but it can archive them daily or any value less often than daily. The available archival schedule frequencies are:
   - **Hour** - Archives only the first successful run that started on that particular hour.
   - **Daily** - Archives only the first successful run that started on that particular day.
   - **Weekly** - From the weekly calendar, select the required days in a week. You can also set the archival schedule for n-number of weeks.
   - **Monthly** - Use the **On** dropdown to select the archival schedule:
     - **First** - First selected day of the month.
     - **Second** - Second selected day of the month.
     - **Third** - Third selected day of the month.
     - **Fourth** - Fourth selected day of the month.
     - **Last** - Last selected day of the month.

       Then, choose the specific day of the week from the adjacent dropdown.
     - **Date** - From the pre-built calendar, select the required date.
   - **Yearly** - From the **On** dropdown, select the required available options:
     - **First Day** - First day of the year.
     - **Last Day** - Last day of the year.
     - **Date** - From the pre-built calendar, select a required date of the year.

e. **Retain for xx days:**

   Full/Incremental Backups

   Specify how long the snapshots are retained on the external target.

> **Note:**
> - Cloud archive snapshots are expired per the retention policy. However, the garbage collection of the expired cloud archive snapshots happens in the background. The expired cloud archive snapshots will be garbage collected from the External Target once there are no active Cloud archive snapshots sharing data with them.
> - You cannot delete the latest archive snapshot as it might be needed in the future. However, if you still want to delete the archive snapshot, you must contact Cohesity Support.
> - The retention period is applicable only to the database backup and not the log backup.
> - The **Retain** option can be set to a maximum value of 365000.

Log Backups

Specify how long the log backup snapshots are retained on the external target.

> **Note:** The Retain option can be set to a maximum value of 365000.

You can add additional archive schedules.

> **Note:** If you specify multiple archive schedules with the same external target that archives the same set of snapshots with different retention periods, the snapshots are retained for the longest specified time period.

f. **Lock**: Specify how many days, weeks, months, or years the retained snapshots should be locked. You cannot delete the snapshots during the specified period.

> **Note:**
> - This option is available only if you have enabled the **DataLock** option at the policy level.
> - The DataLock duration must be less than or equal to the retention period.

g. To add an additional archival schedule, repeat these steps. To archive to multiple External Targets clusters, create multiple replication schedules with different External Targets.

> **Important:** If you specify multiple archival schedules with the same External Targets that archive the same set of snapshots but with different retention periods, the snapshots are retained for the longest time period. For example if you define two archival schedules with the same External Targets that retain the same set of snapshots but one schedule specifies a 90 day retention period and other schedule specifies a 180 day retention period, the set of snapshots are retained for 180 days.

For archival, the extended retention schedule is as follows:

- **Weekly**: The first successful run that started on a Sunday or later.
- **Monthly**: The first successful run that started on the first day of the month or later.
- **Yearly**: The first successful run that started on the first day of the year or later.

19. **Add CloudSpin:** Select this option to add a CloudSpin schedule. For more information, see Clone VMs to AWS Cloud Using CloudSpin.
    a. **CloudSpin to:** Select an existing cloud source or register a new cloud source. A cloud source is a cloud service instance where the snapshots are converted and copied **to**. For details about the different fields, see the appropriate source topic:
        - Register an Azure Cloud Source at Subscription Level
        - Register or Edit AWS
    b. **After every run:** Specify how often snapshots are converted and copied to the cloud source.
    c. **Retain for xx days:** Specify how long the snapshots are retained on the cloud source.

> **Note:** The **Retain** option can be set to a maximum value of 365000.

    d. **Lock**: Specify how many days, weeks, months, or years the retained snapshots should be locked. You cannot delete the snapshots during the specified period.

> **Note:** This option is available only if you have enabled the **DataLock** option at the policy level.

e. **CloudSpin only fully successful Runs:** Check this box displays if you want converting and copying to occur only if the protection run successfully backs up all the protected objects, such as 10 out of 10 VMs. For example, if a protection run backs up only 7 out of 10 VMs, and this check box is selected, the protection run will not be converted and copied. If a protection run backs up only 7 out of 10 VMs, and this check box *is not* selected, the protection run can be converted and copied to a cloud source.

f. Specify the appropriate settings for the cloud service:

AWS

Select the Cloud Region in AWS where the VM should be copied to. For more information, see AWS Regions.

You can optionally add **Tags** to the resources by clicking **Add Tag**. These tags are key-value pairs that will be assigned to all the temporary and permanent resources created during Cloudspin. The maximum number of tags allowed is 40.

Azure

Resource Group—Select the Azure Resource Group where the converted VM image should be copied to.

Storage Account—Select an Azure Storage Account to use when copying the converted VM image.

Storage Container—Select an Azure Storage Container in the Storage Account to hold the converted VM image.

20. Click **Create**.

When editing an existing policy that is configured to capture backup snapshots on a minutes or hourly time schedule and you change the schedule to occur daily, weekly or monthly, you are prompted to specify a start time for each of the Protection Groups that use this policy.

The policy is immediately available to use in Protection Groups.

Note:

- For Cohesity 6.6 and higher versions, on NoSQL platforms, when you pause the backup job for a longer time than the retention period, the last successful snapshot will not be deleted even after the retention period expired on that snapshot.
- For Cohesity 6.6 and higher versions, the last snapshot is always retained on HDFS/HBase/Hive, MongoDB, Cassandra, and CouchBase platforms.

Considerations

- By default, all backups are incremental unless manually switched to Full. However, if no previous incremental backup exists for a specific object, a full backup will be automatically initiated, potentially generating warnings for certain applications.
- When scheduling a continuous backup policy, the protection run does not honor the start time.

  For example, specify a policy to capture the continuous backup at 9 AM every day with a periodicity of 7 hours.

Days	Description
Day 1	• First backup - 9 AM • Second backup - 4 PM (after 7 hours) • Third backup - 11 PM
Day 2	• First backup - 6 AM • Second backup - 1 PM In this example, on Day 2 the backup did not start at 9 AM.

Manage Policies

A standard policy must be associated with a Protection Group to protect objects.

Use the default standard policies or create your own custom policies. Policies save time because you do not need to repetitively enter settings. For more information, see Default Policies.

From the policy manager, you can create a new protection policy. For more information, see Create or Edit a Standard Policy.

You can also view policy settings, copy, edit and delete policies using the actions for each policy.

Column Name	Description / Action
Policy Name	Drill-down on this link to view the settings for the policy.
Organization	The name of the organization to which the policy belongs.
Backup	Snapshots are stored in the current (local) Cohesity cluster and there are one or more extended retention rules.
Replication	Snapshots are stored in the current (local) Cohesity cluster and are archived in the cloud.
Archive	Snapshots are stored in the current (local) Cohesity cluster and are archived in the cloud.
CloudSpin	Snapshots are stored in the current (local) Cohesity cluster and are archived in the cloud.
Actions	To manage an existing policy, click the actions menu (⋮) select one of the following options: • **Copy**( )–Create a new policy from an existing policy. • **Edit** ( )–Configure the settings of an existing policy. Follow the instructions provided in Create or Edit a Standard Policy. • **Delete** ( )–Delete this policy. If a Protection Group is based on a policy and the policy is deleted, the Protection Group keeps the original settings and is not deleted.

> **Note:** For jobs protecting servers, indexing is set at the Protection Group level.

Default Policies

The following default standard policies are provided:

- Gold
  - Take a snapshot every 4 hours and retain it for 7 days.
  - The first snapshot taken every day is retained for 90 days.
  - The first snapshot taken every week is retained for 365 days.
  - The first snapshot taken every month is retained for 1095 days.
  - Retry capturing snapshots 3 times every 5 minutes before reporting an error

- Silver
  - Take a snapshot every 12 hours and retain it for 14 days.
  - The first snapshot taken every day is retained for 90 days.
  - The first snapshot taken every week is retained for 365 days.
  - The first snapshot taken every month is retained for 730 days.
  - Retry capturing snapshots 3 times every 5 minutes before reporting an error
- Bronze
  - Take a snapshot every day and retain it for 30 days.
  - The first snapshot taken every week is retained for 90 days.
  - The first snapshot taken every month is retained for 365 days.
  - Retry capturing snapshots 3 times every 5 minutes before reporting an error
- Protect Once
  - Take a FULL snapshot and retain it for 14 days

> **Note:** Objects protected with the Protect Once policy are considered as unprotected if no other policy covers them. They may have valid snapshots but no ongoing protection.

> **Note:**
> Use the Protect Once policy only for a one-time activity, as Protection Groups using the Protect Once policy cannot be edited to remove or modify individual sources. This policy is intended to protect all included objects as a single unit.
>
> To make changes, consider the following options:
>
> - Wait for snapshots to expire, then delete and recreate the Protection Group.
> - Delete the Protection Group and all associated snapshots (if no longer needed).
> - Create a new Protection Group with the desired sources.

Customize Indexing

Toggle on **Indexing**, click **Add Inclusion** or **Add Exclusion** and enter a directory. Click the **X** to delete a directory from the list.

By default, the root directory (`/`) is included and entire file system is indexed. When a directory is indexed, its sub-directories are also indexed.

You can limit indexing to a specific set of files and directories to minimize the disk space used by the index. For example, you could add two additional rules to the first default rule:

- include `/`
- exclude `/var`
- include `/var/log`

In this example, all files in the `/` root directory are indexed except the files found in the `/var` directory. However, the files in the `/var/log` are indexed because there is an explicit include rule.

Include and Exclude Rules:

- Use forward slashes (/) for both Windows and Unix as the separator character.
- Do not specify the drive letter and colon in front of the directory path for Windows.
- Drive letters are allowed for Windows systems to limit the drives searched for a specified path.
- If **no drive** is specified then **all** drives for which the path is found are indexed.

Examples	Description
/C/Users	On Windows, include or exclude Users directory in C:\Users
/Users	On Windows, include or exclude Users directories in all drives, i.e. C:\Users, D:\Users.
/var/log	On Linux or Unix, include or exclude /var/log directory.

The default directory inclusion/exclusion list are as follows:

- include /
- exclude /$Recycle.Bin
- exclude /Windows
- exclude /Program Files
- exclude /Program Files (x86)
- exclude /ProgramData
- exclude /System Volume Information

- exclude /Users/*/AppData
- exclude /Recovery
- exclude /var
- exclude /usr
- exclude /sys
- exclude /proc
- exclude /lib
- exclude /grub
- exclude /grub2
- exclude /opt/splunk
- exclude /splunk

Excluding Pseudo Filesystems from Indexing

In a Linux operating system, specific filesystems, such as /sys, /proc, /dev, /run, /tmp, /dev/pts, and /dev/shm, are dynamically generated by the kernel during system boot. These filesystems are commonly referred to as 'pseudo filesystems' or 'virtual filesystems.' Their contents are unique to the system and temporary, meaning they do not persist beyond system reboots.

Due to their transient nature, Cohesity does not recommend indexing these pseudo filesystems. Attempting to index them not only lacks practical utility but may also result in inaccurate or inconsistent indexing outcomes.

Configure Pre & Post Scripts

Cohesity provides you the option to run user-defined scripts as a part of the Protection Group. When creating a Protection Group, you can configure:

- Pre Scripts: Scripts that will be executed before the protection run.
- Post Scripts: Scripts that will be executed after the protection run.
- Post Snapshot Scripts: Scripts that will be executed after the snapshot has been created on the source.

Supported Sources

You can run the Pre Script and Post Script for the following sources:

- Physical Servers
- Databases - MS SQL Server and Oracle Database
- NAS
- SAN

Depending on whether the Cohesity Agent is installed on the source, the settings required to configure the pre and post scripts in the Protection Group vary.

Configure Scripts for Sources with Agent Installed

Out of the supported sources for pre and post-scripts, physical servers and databases require Cohesity Agent to be installed on the source.

Following are the settings required for configuring the pre and post scripts for sources with Agent installed:

Pre-Script

- **Script Path**: Depending on the agents, specify the full path or name of the script as mentioned in the **Note** section below. A script matching the specified name must be available in the user_scripts directory found in the agent installation directory on the protected server. (See the table below for details about the directory path.) The script can be in any format, for example: shell script, batch file, executable (`.exe`), etc. The script must be configured with the permissions to allow execution (`+x`).

> **Note:**
> - For Physical servers, MS SQL server and Oracle, only the script name is required (provided the script is available in the `user_scripts` directory found in the agent installation directory).
> - For NAS and SAN sources, the absolute/full path to the script is required.
> - On Windows platforms, PowerShell, Python, or Perl scripts (any script that is not a .exe file) must be embedded within a .bat script before they are called.

> **Important:** Cohesity does not validate the content of the script and is not responsible for the content of the script or any actions done by the script. You must verify the script actions before a script or executable is run.

- **Script Params**: Optional field for passing one or more parameters to the script. All specified parameter values are passed as a single string to the script.
- **Timeout**: Optional field for specifying how long a script is allowed to execute before the Cohesity cluster stops the execution of the script. By default, the timeout is 15 minutes. If no value is specified, there is no timeout and the script is not stopped. This can cause a protection run to hang if the script never completes. The timeout is applied from the moment the script starts. For example, if the timeout is set to 15 minutes but the script takes 20 minutes to complete, it will be stopped after 15 minutes. Therefore, users should set the timeout based on the expected runtime of the script to avoid premature termination.
- **Continue Backup if script fails**: This toggle applies only to Pre-Scripts and is enabled by default. When this toggle is enabled, the Protection Group backs up the objects even if the Pre-Script fails. If this toggle is disabled and the Pre-Script fails, the backup is canceled but the Post-Script still runs.

Post-Script

- **Script Path**: Depending on the agents, specify the full path or name of the script as mentioned in the **Note** section below. A script matching the specified name must be available in the specified directory. The script can be in any format, such as a shell script or executable. The script must be configured with the permissions to allow execution (+x).

> **Note:**
> - For Physical servers, MS SQL server and Oracle, only the script name is required (provided the script is available in the user_scripts directory found in the agent installation directory).
> - For NAS and SAN sources, the absolute/full path to the script is required.
> - On Windows platforms, PowerShell, Python, or Perl scripts (any script that is not a .exe file) must be embedded within a .bat script before they are called.

> **Warning:** Cohesity does not validate the content of the script and is not responsible for the content of the script or any actions done by the script. You must verify the script actions before a script or executable is run.

- **Script Params**: Optional field for passing one or more parameters to the script. All specified parameter values are passed as a single string to the script.
- **Timeout**: Optional field for specifying how long a script is allowed to execute before the Cohesity cluster stops the execution of the script. By default, the timeout is 15 minutes. If no value is specified, there is no timeout and the script is not stopped. This can cause a protection run to hang if the script never completes. The timeout is applied from the moment the script starts. For example, if the timeout is set to 15 minutes but the script takes 20 minutes to complete, it will be stopped after 15 minutes. Therefore, users should set the timeout based on the expected runtime of the script to avoid premature termination.

When the post-script is enabled, it runs irrespective of the pre-script run status.

More about failures, cancellations and log files:

- If the backup is successful but the Pre-Script or Post-Script fails, a warning is issued with the script error. A script failure occurs for any of the following conditions:
  - script is not found in the expected location
  - script cannot be executed
  - script returns a failure return code of non-zero

  A script is successful when the return code is 0.
- If a Post-Script is configured, the script always executes even if the backup fails.
- If a protection run is canceled, all subsequent phases are canceled.
  - For example, if a protection run is canceled while the Pre-Script is running, the Pre-Script is canceled, then the backup task is canceled and then the Post-Script is canceled.
  - For example, if a protection run is canceled while the backup task is running, the backup task is canceled and then the Post-Script is canceled.
- Script execution log files can be found in the `user_script_logs` directory found in the agent log directory on the protected server. Both Stdout and Stderr from the script execution are recorded in the script log files. (See the table below for details about the directory path.)

Available Environment Variables

The Cohesity cluster populates the following environment variables with values before the Pre-Script or the Post-Script is run. You can use these environment variables in your script.

Variable Name	Available from Script	Description
`COHESITY_JOB_ID`	Pre-Script and Post-Script	Specifies the Cohesity id assigned to the Protection Group.
`COHESITY_JOB_NAME`	Pre-Script and Post-Script	Specifies the name of the Protection Group as displayed in the Cohesity Dashboard.
`COHESITY_BACKUP_TYPE`	Pre-Script and Post-Script	`kRegular` indicates an incremental (CBT) backup. Incremental backups utilizing CBT (if supported) are captured of the target protection objects. The first run of a kRegular schedule captures all the blocks. `kFull` indicates a full (no CBT) backup. A complete backup (all blocks) of the target protection objects are always captured and Change Block Tracking (CBT) is not utilized. `kLog` indicates a database log backup. Capture the database transaction logs to allow rolling back to a specific point in time. `kSystem` indicates a system backup. System backups are used to do bare metal recovery of the system to a specific point in time. (*Applies only to IBM*) `kStorageArraySnapshot` indicates storage array snapshot backup. This backup method creates point-in-time snapshots of data on a storage array, allowing for data protection and recovery without impacting the primary storage's performance.
`COHESITY_JOB_STATUS`	Post-Script	SUCCEEDED indicates the protection run ran successfully. ERRORED indicates the protection run reported an error.

> **Note:**
> To access an environment variable in PowerShell, use:`$env:<variable name>`.
>
> Example
>
> ```
> Write-Host "COHESITY_BACKUP_HOST_ENTITY =" $env:COHESITY_BACKUP_HOST_ENTITY
> ```

The following is a sample script for Linux

```
#!/bin/bash
echo "Hello...This is a sample script".
```

```
sleep 60
echo "Executing Script"
```

The following is a sample script for Windows

```
net start OracleOraDB12Home1TNSListener
net start OracleServiceMUC99
net start OracleServiceCATPRD
@oradim -startup -sid MUC99
@oradim -startup -sid CATPRD
```

Paths to User Scripts and Script Logs for Various Agents

Operating Systems	Directory Path Description	Path
Linux	Where Pre-Script and Post-Script must be stored to be found by the Cohesity cluster.	`INSTALL_DIR/user_scripts`
	Where the script execution logs are written by the Cohesity cluster when the script is run.	`INSTALL_DIR/logs/user_script_logs/`
	If the agent is installed using the `cohesityagent` user, the default `INSTALL_DIR` is...	`/home/cohesityagent/cohesityagent/`
	If the agent is installed using the `root` user, the default `INSTALL_DIR` is...	`/root/cohesityagent`
Windows	Where Pre-Script and Post-Script must be stored to be found by the Cohesity cluster.	`INSTALL_DIR\user_scripts`
	The default `INSTALL_DIR` is...	`C:\Program Files\Cohesity`
	Where the script execution logs are written by the Cohesity cluster when the script is run.	`C:\ProgramData\Cohesity\user_script_logs`
Solaris **AIX** **RHEL 5.x**	Where Pre-Script and Post-Script must be stored to be found by the Cohesity cluster.	`INSTALL_DIR/user_scripts`
	The default `INSTALL_DIR` is...	`/usr/local/cohesity/agent`
	Where the script execution logs are written by the Cohesity cluster when the script is run.	`LOG_DIR/user_script_logs`
	The default `LOG_DIR` is...	`/var/log/cohesity/agent`

Configure Scripts for Sources without Agent Installed

Out of the supported sources for pre and post-scripts, NAS and SAN sources do not require Cohesity Agent to be installed on the source.

Following are the settings required for configuring the pre and post scripts for sources without Agent installed:

- **Hostname or IP Address (Unix only)**: Specify either the hostname or the IP address of the remote Unix server where the script will run. (Only Unix is currently supported.)
- **Cluster SSH Public Key**: Click Copy Key to Clipboard to copy the entire string including the ssh-rsa string preceding the key, the public key and the username@cluster_name string after the key. Open the authorized_keys file (typically located in the ~/.ssh directory) and paste the entire string to the end of the file.
- **Username**: Specify the username on the remote Unix machine that will execute the script. The username must match the username specified in the authorized_keys file.

You can configure a Pre-Script, a Post-Script or both scripts.

Pre-Script

- **Script Path**: Depending on the agents, specify the full path or name of the script as mentioned in the **Note** section below. A script matching the specified name must be available in the specified directory. The script can be in any format, such as a shell script or executable. The script must be configured with the permissions to allow execution (+x).

> **Note:**
> - For Physical servers, MS SQL server and Oracle, only the script name is required (provided the script is available in the user_scripts directory found in the agent installation directory).
> - For NAS and SAN sources, the absolute/full path to the script is required.

> **Warning:** Cohesity does not validate the content of the script and is not responsible for the content of the script or any actions done by the script. You must verify the script actions before a script or executable is run.

- **Script Params**: Optional field for passing one or more parameters to the script. All specified parameter values are passed as a single string to the script.
- **Timeout**: Optional field for specifying how long a script is allowed to execute before the Cohesity cluster stops the execution of the script. By default, the timeout is 15 minutes. If no value is specified, there is no timeout and the script is not stopped. This can cause a protection run to hang if the script never completes. The timeout is applied from the moment the script starts. For example, if the timeout is set to 15 minutes but the script takes 20 minutes to complete, it will be stopped after 15 minutes. Therefore, users should set the timeout based on the expected runtime of the script to avoid premature termination.
- **Continue Backup if script fails**: This toggle applies only to Pre-Scripts and is enabled by default. When this toggle is enabled, the Protection Group backs up the objects even if the Pre-Script fails. If this toggle is disabled and the Pre-Script fails, the backup is canceled but the Post-Script still runs.

Post-Script

- **Script Path**: Depending on the agents, specify the full path or name of the script as mentioned in the **Note** section below. A script matching the specified name must be available in the specified directory. The script can be in any format, such as a shell script or executable. The script must be configured with the permissions to allow execution (+x).

  > **Note:**
  > - For Physical servers, MS SQL server and Oracle, only the script name is required (provided the script is available in the user_scripts directory found in the agent installation directory).
  > - For NAS and SAN sources, the absolute/full path to the script is required.

  > **Warning:** Cohesity does not validate the content of the script and is not responsible for the content of the script or any actions done by the script. You must verify the script actions before a script or executable is run.

- **Script Params**: Optional field for passing one or more parameters to the script. All specified parameter values are passed as a single string to the script.
- **Timeout**: Optional field for specifying how long a script is allowed to execute before the Cohesity cluster stops the execution of the script. By default, the timeout is 15 minutes. If no value is specified, there is no timeout and the script is not stopped. This can cause a protection run to hang if the script never completes.

The timeout is applied from the moment the script starts. For example, if the timeout is set to 15 minutes but the script takes 20 minutes to complete, it will be stopped after 15 minutes. Therefore, users should set the timeout based on the expected runtime of the script to avoid premature termination.

When the post-script is enabled, it runs irrespective of the pre-script run status.

More about failures, cancellations and log files:

- If the backup is successful but the Pre-Script or Post-Script fails, a warning is issued with the script error. A script failure occurs for any of the following conditions:
  - script is not found in the expected location
  - script cannot be executed
  - script returns a failure return code of non-zero

  A script is successful when the return code is 0.
- If a Post-Script is configured, the script always executes even if the backup fails.
- If a protection run is canceled, all subsequent phases are canceled.
  - For example, if a protection run is canceled while the Pre-Script is running, the Pre-Script is canceled, then the backup task is canceled and then the Post-Script is canceled.
  - For example, if a protection run is canceled while the backup task is running, the backup task is canceled and then the Post-Script is canceled.
- Script execution log files can be found in the `user_script_logs` directory found in the agent log directory on the protected server. Both Stdout and Stderr from the script execution are recorded in the script log files. (See the table below for details about the directory path.)

Available Environment Variables

The Cohesity cluster populates the following environment variables with values before the Pre-Script or the Post-Script is run. You can use these environment variables in your script.

Variable Name	Available from Script	Description
`COHESITY_JOB_ID`	Pre-Script and Post-Script	Specifies the Cohesity id assigned to the Protection Group.

Variable Name	Available from Script	Description
`COHESITY_JOB_NAME`	Pre-Script and Post-Script	Specifies the name of the Protection Group as displayed in the Cohesity Dashboard.
`COHESITY_BACKUP_TYPE`	Pre-Script and Post-Script	`kRegular` indicates an incremental (CBT) backup. Incremental backups utilizing CBT (if supported) are captured of the target protection objects. The first run of a kRegular schedule captures all the blocks. `kFull` indicates a full (no CBT) backup. A complete backup (all blocks) of the target protection objects are always captured and Change Block Tracking (CBT) is not utilized. `kLog` indicates a database log backup. Capture the database transaction logs to allow rolling back to a specific point in time. `kSystem` indicates a system backup. System backups are used to do bare metal recovery of the system to a specific point in time. (*Applies only to IBM*) `kStorageArraySnapshot` indicates storage array snapshot backup. This backup method creates point-in-time snapshots of data on a storage array, allowing for data protection and recovery without impacting the primary storage's performance.
`COHESITY_JOB_STATUS`	Post-Script	SUCCEEDED indicates the protection run ran successfully. ERRORED indicates the protection run reported an error.
`COHESITY_BACKUP_VOLUME_SNAPSHOT_SUFFIX`	Pre-Script	Contains the suffix Cohesity uses for creating the Pure Storage Volume snapshot.
COHESITY_BACKUP_GROUP_SNAPSHOT_SUFFIX	Post-Script	Contains the suffix Cohesity uses for creating the Pure Protection Group snapshot.

DataLock

DataLock is the Cohesity WORM (write once read many) feature that locks and retains files for compliance and regulatory purposes and to protect the files from accidental deletion.

DataLock enables your storage to be WORM-compliant to ensure the snapshot is not aged prematurely. When it is enabled, your data such as backups and archives cannot be deleted by anyone, including administrators, until the DataLock expires. This retention lock helps to prevent ransomware attacks on your clusters.

> **Note:** The contents of backups cannot be modified because of Cohesity SpanFS's immutable file system design.

The Cohesity Dashboard displays a lock ( ) icon to indicate DataLocked policies and snapshots.

Features of DataLock

- You can apply DataLock for each schedule (such as backup, extended retention, archive, and replication) defined in the Protection Policy.
- You can specify DataLock duration in days, weeks, months, or years for each schedule defined in the Protection Policy.

Considerations

- The DataLock expiry date is set at the time of snapshot creation and does not change afterward.
    - Extending snapshot retention—whether for runs created before or after the change—does not extend the DataLock expiry.
    - Modifying snapshot retention after creation does not alter the originally configured DataLock expiration.
- Using DataLock with longer snapshot retention periods and for auto-protected sources can result in the cluster running out of space.
- DataLock retention must be less than or equal to backup retention.

Policy

The following table provides DataLock details when creating a protection policy by a user with a Data Security role (DSO) and a Non-Data Security role.

Workflow	Data Security	Non-Data Security
Create Policy	• DataLock is enabled by default when creating a protection policy, replication target, archival target, and CloudSpin. • Disable DataLock for backup, replication, archive, and CloudSpin. • Edit both backup retention and the Lock retention period.	• DataLock is enabled by default when creating a protection policy, replication target, archival target, and CloudSpin. • Disable DataLock for backup, replication, archive, and CloudSpin. • Edit both backup retention and the Lock retention period.
Edit Policy	• Able to edit the DataLock setting for backup, replication, archival, and CloudSpin. • You can add a new replication, archival or CloudSpin target, and to add DataLock for a target, click **Add DataLock**. • Able to edit both backup retention and the Lock retention period. • Delete targets irrespective of the DataLock setting.	• Unable to edit the DataLock setting for backup, replication, archival, and CloudSpin. • You can add a new replication, archival or CloudSpin target without DataLock because **Add DataLock** is disabled. • Unable to edit the Lock retention period but the backup retention can be modified. • Targets with DataLock enabled cannot be deleted.

Protection

The following table provides DataLock details when created by a user with a Data Security role (DSO) and a Non-Data Security role.

Workflow	Data Security	Non-Data Security
Run Now	• Able to edit the backup retention period. • Able to add a new target but the DataLock setting is not available. • If you want to specify DataLock for a new target it must be enabled when creating or editing the protection policy. • Delete existing targets irrespective of the DataLock setting.	• Able to edit the backup retention period. • Able to add a new target but the DataLock setting is not available. • If you want to specify DataLock for a new target it must be enabled when creating the protection policy. • Delete existing targets irrespective of the DataLock setting.

Workflow	Data Security	Non-Data Security
Edit Run	• Able to edit the backup retention period. • Able to add a new target but the DataLock setting is not available. If you want to add DataLock to a new target it must be enabled when creating or editing the protection policy. • Unable to delete existing targets. • Delete the DataLocked snapshot only if the Lock retention period expires or DataLock is disabled on the snapshot.	• Able to edit the backup retention period. • Able to add a new target but the DataLock setting is not available. If you want to add DataLock to a new target it must be enabled when creating the protection policy. • Unable to delete existing targets. • Delete the DataLocked snapshot if the Lock retention period expires or DataLock is disabled on the snapshot.

Upgrade

Upgraded Cluster	New Cluster
• No change to the existing system and user-created policies. • DataLock is enabled by default when creating a new policy. • You can disable DataLock on specific targets (Backup, Replication, Archival, and CloudSpin) when creating a new policy irrespective of the user's role.	• DataLock is enabled by default for both system policy and new policy created by a user. • You can disable DataLock on specific targets (Backup, Replication, Archival, and CloudSpin) when creating a new policy irrespective of the user's role.

Extended Retention

Snapshots selected for extended retention are retained for a longer time than defined in the protection schedule. You can retain the first snapshot of the backup run taken on the specified hours, every run, days, weeks, months, or years for a longer period of time.

> **Note:** The retention period can be set to a maximum value of 365000 days.

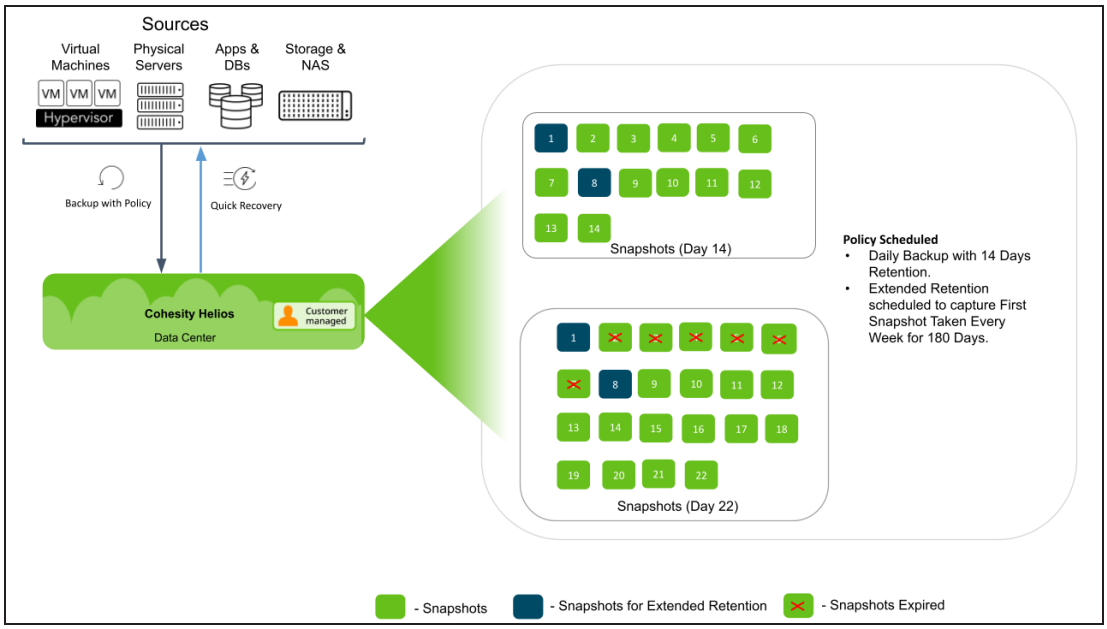

The table below lists the options available for protection policy and extended retention policy and the respective outcome.

Backup	Extended Retention	Description	Example
Every	Retain the First Snapshot taken every		

Backup	Extended Retention	Description	Example
Minutes	Run	The first snapshot captured after every run is retained for the specified number of days in the extended retention policy.	
	Hours	The first snapshot captured during the specified hour is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured during the hour is retained for the number of days specified in the protection policy.	Define a protection policy that captures snapshots every 10 minutes and retains the snapshots for 14 days. Specify an extended retention policy to capture the first snapshot taken every hour and retain the snapshot for 28 days. **6:38PM** - The first backup starts. The snapshot captured is retained for an extended period of 28 days because this is the first snapshot in the hour 6PM - 7PM. **6:48PM** - The snapshot captured is retained for 14 days. **6:58PM** - The snapshot captured is retained for 14 days. **7:08PM** - The snapshot captured is retained for 28 days as this is the first snapshot in the hour 7PM - 8PM.

Backup	Extended Retention	Description	Example
	Days	The first snapshot captured on the specified day is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the protection policy.	Define a protection policy that captures snapshots every 10 minutes and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every day and retain it for 28 days. **6:30PM (Day 1)** - The first backup starts. The snapshot captured is retained for an extended period of 28 days because this is the first snapshot for Day 1. **6:40PM (Day 1)** - The snapshot captured is retained for 14 days. **11:50PM (Day 1)** - The snapshot captured is retained for 14 days. **12:00AM (Day 2)** - The snapshot captured is retained for an extended period of 28 days because this is the first snapshot for Day 2.

Backup	Extended Retention	Description	Example
	Weeks	The first snapshot captured during the specified week is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the protection policy.	Define a protection policy that captures snapshots every 10 minutes and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every week and retain the snapshot for 28 days. Assuming the week starts on Sunday: **Week 1, Sunday 6:30PM** - The snapshot captured is retained for 28 days. **Week 1, Thursday 8:30PM** - The snapshot captured is retained for 14 days. **Week 1, Saturday 11:50PM** - The snapshot captured is retained for 14 days. **Week 2, Sunday 12:00AM** - The snapshot captured is retained for 28 days.
	Months	The first snapshot captured during the specified month is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the protection policy.	Define a protection policy that captures snapshots every 10 minutes and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every month and retain the snapshot for 28 days. **January** - The first snapshot captured is retained for 28 days. Every subsequent snapshot captured in the month is retained for 14 days. **February** - The first snapshot captured is retained for 28 days. Every subsequent snapshot captured in the month is retained for 14 days.

Backup	Extended Retention	Description	Example
Days	Run	Cohesity recommends configuring the Extended Retention policy always more than the scheduled backup policy.	

Backup	Extended Retention	Description	Example
	Hours		

Backup	Extended Retention	Description	Example
	Days	The first snapshot captured during the specified day is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the backup policy.	Define a backup policy that captures snapshots every day and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every 2 days and retain the snapshot for 28 days. **7:00PM (Day 1)** - The first backup starts. The snapshot captured is retained for an extended period of 28 days. **7:00PM (Day 2)** - The snapshot captured is retained for 14 days. **7:00PM (Day 3)** - The snapshot captured is retained for an extended period of 28 days. **7.00PM (Day 4)** - The snapshot captured is retained for 14 days.
	Weeks	The first snapshot captured during the specified week is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the backup policy.	Define a backup policy that captures every day and retains it for 14 days. Specify an extended retention policy to capture the first snapshot taken every week and retain the snapshot for 28 days. Assuming the week starts on Sunday: **Week 1, Sunday 6:30PM** - The snapshot captured is retained for 28 days. **Week 2, Sunday 6:30PM** - The snapshot captured is retained for 28 days.

Backup	Extended Retention	Description	Example
	Months	The first snapshot captured during the specified month is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the backup policy.	Define a backup policy that captures snapshots every day and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every month and retain the snapshot for 28 days. **January** - The first snapshot captured is retained for 28 days. Every subsequent snapshot captured in the month is retained for 14 days. **February** - The first snapshot captured is retained for 28 days. Every subsequent snapshot captured in the month is retained for 14 days.

Backup	Extended Retention	Description	Example
Week	Run	Cohesity recommends scheduling the Extended Retention policy always more than the scheduled backup policy.	

Backup	Extended Retention	Description	Example
	Hours		

Backup	Extended Retention	Description	Example
	Days		

Backup	Extended Retention	Description	Example
	Weeks	The first snapshot captured during the specified week is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the backup policy.	Define a backup policy that captures snapshots every week and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every two weeks and retain the snapshot for 28 days. Assuming the week starts on Sunday: **Week 1, Sunday 6:30PM** - The snapshot captured is retained for 28 days. **Week 2, Sunday 6:30PM** - The snapshot captured is retained for 14 days. **Week 3, Sunday 6:30PM** - The snapshot captured is retained for 28 days.

Backup	Extended Retention	Description	Example
	Months	The first snapshot captured during the specified week is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the backup policy.	Define a backup policy that captures snapshots every week and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every month and retain the snapshot for 28 days. Assuming the week starts on Sunday: **Week 1, Sunday 6:30PM** - The snapshot captured is retained for 28 days. **Week 2, Sunday 6:30PM** - The snapshot captured is retained for 14 days. **Week 3, Sunday 6:30PM** - The snapshot captured is retained for 14 days. **Week 4, Sunday 6:30PM** - The snapshot captured is retained for 14 days.

Backup	Extended Retention	Description	Example
Month	Run	Cohesity recommends scheduling the Extended Retention policy always more than the scheduled backup policy.	
	Hours		
	Days		
	Weeks		
	Months	The first snapshot captured during the specified month is retained for the number of days as mentioned in the extended retention policy. Every subsequent snapshot captured is retained for the number of days specified in the backup policy.	Define a backup policy that captures snapshots every month and retains the snapshot for 14 days. Specify an extended retention policy to capture the first snapshot taken every two months and retain the snapshot for 28 days. **January, February** - The first snapshot captured is retained for 28 days. Every subsequent snapshot captured in the month is retained for 14 days. **March, April** - The first snapshot captured is retained for 28 days.

You can create one or more extended schedules.
For example, you could define a protection schedule that captures snapshots every hour and retains them on the Cohesity cluster for 1 day. This protection schedule retains 24 sets of snapshots on a rotating basis, each hour a new set of snapshots are retained and the oldest set of snapshots is deleted.
If you want to retain a subset of these snapshots for a longer period of time, you could define an additional extended retention schedule to retain a single subset of snapshots daily for 90 days. This extended retention schedule retains 90 sets of snapshots on a rotating basis, each day a new set of snapshots are retained and the oldest set of snapshots is deleted.

> **Note:** If you specify multiple extended retention rules for the same set of snapshots with different retention periods, the snapshots are retained for the longest specified time period. For example if you define two extended retention rules that retain the same set of snapshots but one rule specifies a retention period of 90 days and another specifies a retention period of 180 days, the set of snapshots are retained for 180 days.

The extended retention schedule for backups is based on calendar-based scheduling for weekly, monthly, and yearly. For archive, replication, and CloudSpin; the extended retention schedule is as follows:

- **Weekly**: The first successful run that started on a Sunday or later.
- **Monthly**: The first successful run that started on the first day of the month or later.
- **Yearly**: The first successful run that started on the first day of the year or later.

For backup schedule, Cohesity allows you to set the day for weekly backup and any day of the week for monthly backup.

> **Note:** If a backup fails, the extended retention is applied to the next scheduled snapshot taken by the Cohesity cluster. However, if a manual backup was run in place of the failed backup after selecting "all objects" when selecting the VMs, extended retention is applied to the manual snapshot. For example, you've set the extended retention to be applied on the first snapshot on the first of every month. The scheduled backup on Feb 1st fails. After selecting all the objects (VMs), you run a manual backup successfully. The extended retention is applied to the snapshot of the manual backup.

Chapter 9: SmartFiles (File and Object Services)

This chapter includes the following topics:

- SmartFiles (File and Object Services) Overview 1264
- Ideal Use Cases for SmartFiles 1267
- Fundamentals of SmartFiles 1269
- Cohesity Views 1281
- File Shares (SMB) 1343
- Network File System (NFS) 1372
- Objects (S3) 1393
- SmartFiles Dashboard 1450
- OpenStack Swift Storage Protocol 1467
- NAS Tiering 1482
- File Services Audit Logs 1505
- Antivirus Software to Protect Cluster Data 1514
- Access Management 1518

SmartFiles (File and Object Services) Overview

Cohesity SmartFiles is an enterprise-class, software-defined, data-centric, multiprotocol file and object solution for the enterprise that transcends traditional offerings in terms of manageability, scale, security, efficiency and multi-tiered data management. SmartFiles modernizes and simplifies data and application management by providing one platform for multiple workloads. It is a modern converged target for consolidating data silos and securely managing unstructured content and application data, including digital libraries, archives, rich media, video surveillance, big data, and backup data sets.

You can use Cohesity file and object services for the following use cases:

- **Backup targets**—Store backups from other products, such as Veeam. Reduce storage requirements with automatic deduplication and compression.
- **Home directories**—Create personal folders for Windows and Linux users. Use ID mapping to ensure consistent access control. Establish view user quotas and block certain file types. Create share aliases to simplify access.
- **Corporate file shares**— Store, access and share company files and documents. Lock and retain files for compliance and regulatory purposes.
- **Video and image stores**—Stream videos and store images of virtually any file size.
- **Archiving and tiering**—Archive infrequently used data to public clouds for long-term retention. Use cloud tiering for seamless capacity expansion into the cloud. Replicate into the cloud for disaster recovery and test/dev operations.
- **PACS images**—Store digital medical images for compliance and disaster recovery purposes. Use built-in antivirus software to protect against any viruses and ransomware. Lock and retain files for HIPAA compliance.
- **AWS S3 integration**—Use Cohesity cluster instead of S3 buckets.

Key Components

SmartFiles uses the following Cohesity cluster components:

- **Storage domains**—A named storage location on a cluster. When a cluster is created, it contains a default storage domain. You can create additional storage domains and customize them for the intended use case.

  When you create a storage domain, you define its options, including deduplication, compression, encryption, quotas, fault tolerance and erasure coding. Storage domains contain views.

- **Views**—A named storage location in a storage domain used for storing user data such as files, backup snapshots, objects and cloned VMs.

When you create a view, it inherits some options, like deduplication and compression, from its storage domain, but you can override them. You also define attributes such as access protocols, a Quality of Service policy, security settings, quotas, WORM (DataLock), file filtering and more.

Views are automatically assigned protocol-specific mount paths so they can be used as NFS exports and SMB shares.

Before a client system can mount a view, the client system's subnet must be added to the allowlist in the Cohesity cluster.

- **Allowlists**—Lists of subnets. Systems with IP addresses in those subnets can mount and access views. The Cohesity cluster supports global, view-level and share-level allowlists.

Supported Access Protocols

SmartFiles supports the following protocols for accessing views from client systems:

- NFS 3.0 and NFS 4.1—Mount a View Using NFS
- SMB 3.0 and SMB 2.x—Mount a View Using SMB
- AWS S3—Objects (S3)

Optimize for the Use Case

You can and should configure storage domains and views for the expected use case. For example, the data stored in a view can be written to faster SSD storage or slower, less expensive HDD storage, depending on the use case. The view's Quality of Service (QoS) policy determines the type of storage used and the latency factor in writing data to the view. The following use cases demonstrate why you might choose one QoS over another.

- **Videos**—Let's say you plan to store surveillance videos from security cameras that are continually streaming videos to a view. This workload entails sequential writes (recording the video) and sequential reads (video playback). Because this workload is not latency-sensitive, the data can be written to slower HDD storage. You can accomplish this by configuring the view with the appropriate QoS policy. In this case, you would use one of the Backup Target QoS policies.
- **Home directories**—Let's say you plan to set up home directories for university students to store their work in. This workload entails random reads and writes. Because this workload is latency-sensitive (students are actively accessing their work), the data should be written to faster SSD storage. In this case, you would use one of the TestAndDev QoS policies.

View and Storage Domain Options at a Glance

The following table lists configuration options and indicates whether the option is set at the view level or storage domain level.

Category	Option	Description	Setting
Storage efficiency	Deduplication	Eliminates redundant copies of data to reduce storage required on the Cohesity cluster.	Set for storage domain. Inherited by view and can be modified.
	Compression	Further reduces storage requirements with compression algorithms.	
Resiliency	Fault tolerance	Ensures organizations can continue conducting business in the event of a resource outage on the Cohesity cluster.	Storage domain
	Erasure coding	Ensures high availability by distributing data chunks across disks and nodes in the Cohesity cluster.	Storage domain
Security and compliance	Allowlists	Add subnets to allowlists so devices in those subnets can mount and access views using the NFS, SMB, or S3 protocol.	View
	Security styles (Native, Unified or NTFS)	Keep Linux and Windows permissions separate, or keep them in sync so when one changes, the other automatically changes. Or control NFS and SMB access by using the file's SMB ACL.	View
	Encryption	Encrypt data-at-rest and data-in-flight using the industry standard 256-bit Advanced Encryption Standard (AES) in Galois/Counter Mode. Cohesity cluster is also FIPS 140-3 compliant.	View
	Antivirus scanning	Use your choice of antivirus software to protect data on the Cohesity cluster.	View
	DataLock (WORM)	Enable long-term retention of files and prevent modification for compliance and regulatory purposes.	View
	File extension blocking	Allow or disallow certain file types in the view.	View

Category	Option	Description	Setting
Authentication and authorization	Active Directory	Authenticate SMB shares using Active Directory credentials.	View
	NTLM	Authenticate SMB shares based on NTMLv2, for clients that are not part of Active Directory. The clients may be part of Active Directory, but not able to use Kerberos, so they would fall back to using NTLM. Also, if a share is accessed, it would always use NTLM, as Kerberos requires using a Fully Qualified Domain Name (FQDN).	View
	LDAP	Authenticate NFS exports and SMB shares using LDAP credentials.	View
	Multi-protocol access	Automatically map permissions for multi-protocol views if the Cohesity cluster is joined to an Active Directory domain.	View
	RBAC	Create and manage custom Cohesity cluster administration roles for domain users and groups.	View
Quotas	Physical	Limit the physical data capacity for a storage domain.	Set for storage domain. Inherited by view and can be modified.
	Logical	Limit the logical data capacity of a view for users.	
Cold data	Cloud Tier	Automatically move cold, rarely accessed data to a cloud storage external target.	Storage domain
Performance	QoS policy	Optimize latency, throughput and priority of view I/O based on the expected workload.	View

Ideal Use Cases for SmartFiles

Cohesity SmartFiles is a software-defined, data-centric, multiprotocol file and object solution for the enterprise. The following image illustrates a few capacity-centric workloads which are ideal for SmartFiles:

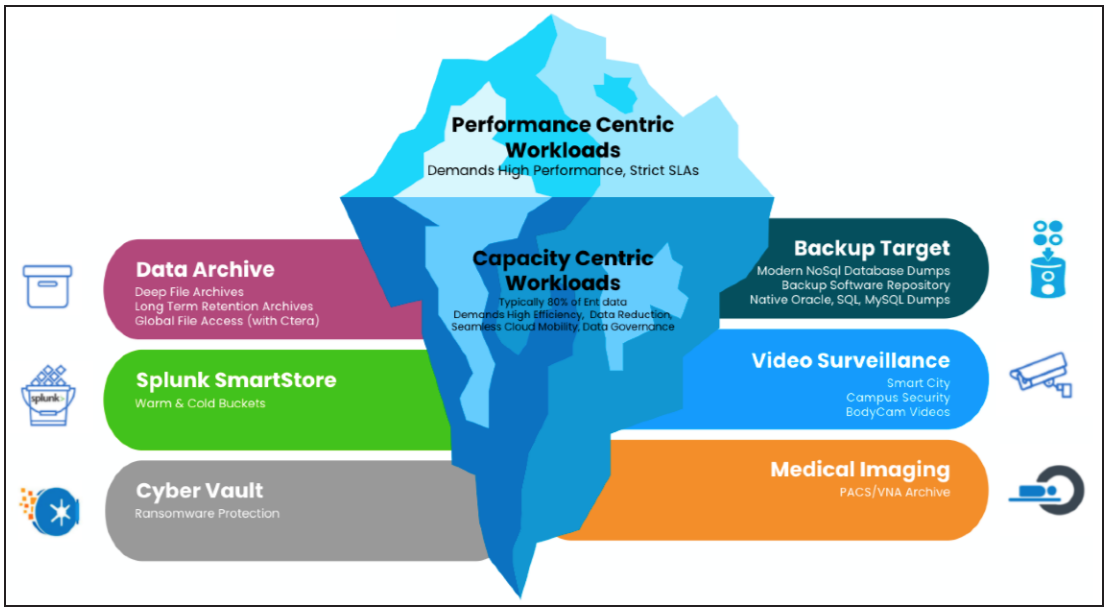

There are dozens of capacity-centric workloads and use cases in which SmartFiles provides benefits. Amongst these use cases are:

- **Backup Target**—Includes the following workloads:
  - Backup software repository for applications such as Veeam Backup & Replication, Commvault Backup & Recovery, IBM Spectrum Protect, and NetBackup AdvancedDisk. For more information, see the following documents:
    - Use Cohesity for Your Veeam Backup Repository
    - Use Cohesity for Your Commvault Storage Libraries
    - Use Cohesity for Your IBM Spectrum Protect Storage Pools
    - Use Cohesity for Your NetBackup AdvancedDisk Storage
  - Database dumps for modern NoSQL applications such as Redis, CockroachDB, Vertica, and MongoDB.
  - Database dumps for traditional applications such as native MySQL, MS SQL, and Oracle SQL. For more information, see the following documents:
    - Streamline Microsoft SQL Server Native Backups with Cohesity
    - Best Practices for Oracle RMAN Backups on Cohesity Views
- **Data Archive**—Long-term retention for digital documents and audio and video files.
- **Splunk SmartStore**—Splunk SmartStore reduces overall cost of ownership by plugging into SmartFiles S3 views.
- **Video Surveillance**—Repository for surveillance videos captured using applications such as Milestone. Smart cities and campuses can use SmartFiles for storing videos.
- **Medical Imaging**—PACS archives and VNA archives.

- **Cyber Vault**—Long-term retention of data in an immutable format using Cohesity DataLock or S3 Object Lock and Legal Hold.

Fundamentals of SmartFiles

The following topics introduce you to key file and object services concepts:

About Deduplication

Deduplication eliminates redundant copies of data to reduce storage usage on the Cohesity cluster. Redundant data blocks are replaced with a pointer to the unique data copy, which ensures that only unique instances of data are transferred over the network and retained on storage media.

When you create a storage domain, you choose one of the following options:

- **Inline deduplication:** Data is deduplicated as it is written to the Cohesity cluster. Generally, Cohesity recommends inline deduplication. It is the default when creating a storage domain.
- **Post-process deduplication:** Data is deduplicated after it is written to the Cohesity cluster, which can improve the speed of I/O operations.
- **No deduplication:** Data is not deduplicated at all. You may want to disable deduplication if access speed is more important than conserving disk space, or if you are storing data that has already been deduplicated by another product.

Storage domains contain user-created views. The views inherit the deduplication setting from the storage domain. If inline deduplication is enabled for the storage domain, you can override it and always use post-process deduplication for the view.

> **Tip:** If the expected workload for a view is latency sensitive, Cohesity recommends post-process deduplication because it improves the speed of I/O operations to the view.

For more information about deduplication, see the white paper Cohesity Fault Tolerance — Data Integrity for Modern Web Scale Environments.

QoS Policies for Cohesity Views

You can assign a QoS policy to a view. A Cohesity view provides a storage location with NFS, SMB, and S3 mount paths in a storage domain on the Cohesity cluster. A QoS policy optimizes the view's performance in terms of throughput, IOPS, and latency for different workloads. The policies also let you prioritize I/O operations for certain views above others.

You need to assign the QoS policy to the view when creating the view. You can create a view using a predefined template or custom template. For more information, see:

- Create Views Using a Predefined Template
- Create Views Using a Custom Template

Predefined templates improve the performance and can also be used for different use cases and scenarios. Cohesity recommends creating a new view using a predefined template.

While creating a view, you can select a single category and the QoS policy defaults to the best optimization mode.

The table below lists the default QoS policies:

Category	Default QoS Policy
File Shares	Test And Dev High
Backup Target	Backup target High
Object Services	Test And Dev High

The following table lists the default properties for view templates. When you create a view using templates, these optimal settings for use cases are selected by default. For more information on view templates, see Cohesity Views.

View Category	View Template	Properties		
		Deduplication	Compression	Performance (QoS)

View Category	View Template	Properties		
File Shares	PACS Archive (Picture archiving and communication system)	No	No	Test And Dev High
	Multimedia	No	No	Test And Dev High
	Video	No	No	Backup Target High
	General Purpose Files	Yes	Yes	Test And Dev High
	Digital Archive	Yes	Yes	Test And Dev High
Backup Target	Commvault	Yes	Yes	Backup Target Commvault
	Veeam	Yes	Yes	Backup Target SSD
	General	Yes	Yes	Test And Dev High
	TSM	Yes	Yes	Test And Dev High
	Application Dumps (Oracle, SQL, SAP Hana)	Yes	Yes	Test And Dev High
Object Services	Containers	Yes	Yes	Test And Dev High
	General	Yes	Yes	Test And Dev High
	Splunk Smartstore	Yes	Yes	Test And Dev High
	Hadoop	Yes	Yes	Test And Dev High

There are two basic QoS policies, Test And Dev and Backup Target, and each has variants by storage media and priority.

QoS Policy	Optimized for I/O Workload Type	Priority ↓
Test And Dev	Mixed I/O workloads such as file shares. Best for writes and sequential single stream reads and writes. Recommended for Cohesity SmartFiles.	High
		Low
Backup Target	Best for sequential, multi-stream reads and writes. Recommended for Cohesity DataProtect and Cohesity Archive. Recommended for Cohesity SmartFiles when used as a backup target.	SSD
		High
		Low

Be sure to select an appropriate QoS policy for the view. By specifying a priority and the type of workload you expect, the Cohesity cluster can optimize how it processes data in regard to latency, throughput, and priority.

The following table lists the QoS policies in the decreasing order of QoS priority.

QoS Policy	Description
Journaled Sequential Dump	Use this policy for workloads that write large amounts of data sequentially to a very small number of files and do not keep many I/Os outstanding. By default data is written to the SSD and has the highest priority and low latency.
TestAndDev High	Use this policy for workloads that require lower I/O latency or do not keep many I/Os outstanding, as the I/Os are given higher priority compared to other QoS policies. Data is written to the SSD.
TestAndDev Low	The same as TestAndDev High, except that the I/Os with this QoS policy are given lower priority compared to I/Os with TestAndDev High when there is contention.
Backup Target CommVault	Use this policy to intelligently detect and exclude application-specific markers to achieve better deduplication when CommVault backup application is writing to a Cohesity View. Data is written to SSD and has the same priority and latency as TestAndDev High.
Backup Target Auto	*(Applicable only for Cohesity C6000 series)* Use this policy for workloads such as backups, which keep many I/Os outstanding. This policy distributes I/Os across SSDs and HDDs to achieve maximum performance based on the current usage. The priority for processing workload with this policy is lower than Journaled Sequential Dump, TestAndDev High, and Backup Target CommVault.

QoS Policy	Description
Backup Target SSD	Use this policy for workloads such as backups, which keep many I/Os outstanding, but in this case, Cohesity sends both sequential and random I/Os to SSD. The latency is lower than other Backup Target policies.The priority for processing workload with this policy is the same as Backup Target Auto.
Backup Target High	Use this policy for non-latency sensitive workloads such as backups, which keep many I/Os outstanding. Data is written to HDD and has higher latency compared to other QoS policies writing to a SSD The priority for processing workload with this policy is the same as Backup Target Auto.
Backup Target Low	The same as Backup Target High, except that the priority for processing workloads with this policy is lower than workloads with Backup Target Auto / High /SSD when there is contention.

Cohesity allows you to pin views to the SSD tier. For more information, see Pin Views to SSD.

In hybrid nodes that have both SSDs and HDDs, QoS policies use the storage media differently to optimize performance and life of the media. Many factors such as I/O block size and I/O pattern are considered. In general, Test And Dev uses SSDs for accelerating overall write performance. Backup Target often reads and writes directly from HDDs. For sequential reads, all QoS policies prefetch data intelligently into RAM and/or SSDs.

ID Mapping for Multi-Protocol Views

ID mapping maps Unix and Windows identifiers to provide consistent access control to files on shared storage, such as Cohesity views. Specifically, the Unix user and group IDs (UIDs/GIDs) are mapped to Windows user and group Security Identifiers (SIDs).

> **Tip:** ID mapping is used only for views that support both the NFS (Unix) and SMB (Windows) protocols *and* that have the Security Style set to either Unified or NTFS, as described later in this topic. ID mapping is irrelevant to views that support just one protocol.

You can use your organization's existing ID mapping in an external server or configure the Cohesity cluster to maintain internal ID mapping, as described in the following sections.

External Server ID Mapping

If the Cohesity cluster is joined to an Active Directory (AD) domain controller and your organization has existing ID mapping on any of the following servers, you can configure the Cohesity cluster to use that mapping for access to multi-protocol views. The administrator of the external server maintains the ID mapping.

- RFC2307 within AD
- SFU within AD
- Centrify
- LDAP

You can also configure fallback mapping in case the external server is not available.

Internal ID Mapping

If the Cohesity cluster is joined to AD but your organization does not have an existing ID map, you can still configure the Cohesity cluster to use one of the following basic mapping types. The ID mapping will be stored internally on the Cohesity cluster.

- Fixed - A user-defined UID and GID are applied to all files created on the view.
- RID - The relative ID of the SID is used as the UID or GID.

Required Configuration for ID Mapping

To use ID mapping, do following in the Cohesity Dashboard:

- Configure a mapping type for the AD domain controller(s) on the Active Directory page. For more information, see Join Active Directory.

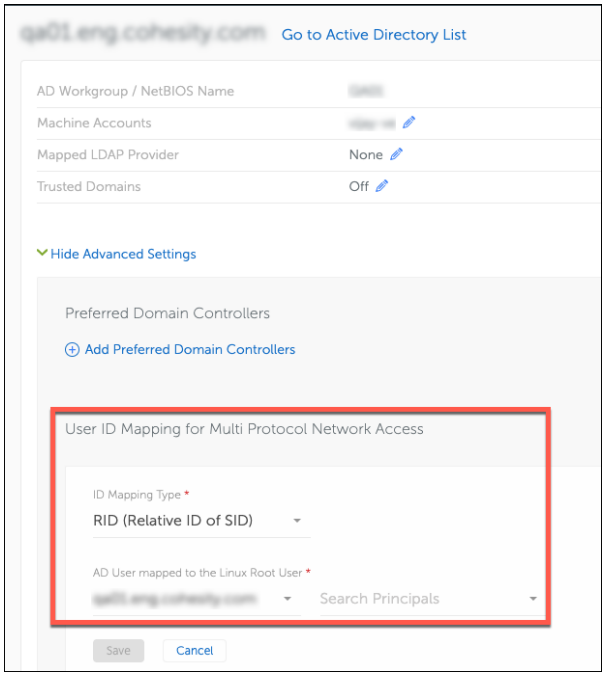

- Configure the following options for multi-protocol views. For information, see Create Views Using a Predefined Template.
  - **Read/Write Protocol** must be set to the required protocol (NFS, NFS v4.1 or SMB).

Optionally S3 access is selected. It is the default Read-Only Protocol. You can disable the access if you need it.

ID mapping is irrelevant if Read/Write Protocol or Read-Only Protocol is set to NFS only, SMB only or S3 only.

- **Security Style** must be set to Unified or NTFS.

  Unified means that Unix and Windows permissions are kept consistent; when one is changed, the other is automatically changed.

  NTFS means that access via SMB and NFS is controlled by the Windows ACL on the file/folder.

  ID mapping is irrelevant if Security Style is set to Native.

User and Directory Quotas for Views

In addition to the storage domain quotas, you can enable user quotas and/or directory quotas at the view level.

User Quotas

User quotas limit how much view capacity a user can use. After a view is created, you can set a default quota for all users and override the default for individual users.

> **Note:** If the view is used through both NFS and SMB, Cohesity strongly recommends using ID mapping to ensure accurate quota accounting. For more information, see ID Mapping for Multi-Protocol Views.

When a user reaches 85% of the quota, an alert (`UserExceededQuotaAlertLimit`) is triggered in the Cohesity cluster. If email notifications are configured for alerts, emails are also sent. To know more information, you can use the User Quotas.

Individual users will also receive an email if the Cohesity cluster joins an Active Directory domain and the user's email address is defined in the AD. NFS users will receive an email if they are mapped to an AD user whose email address is defined in the AD. Alert emails are sent to the addresses specified in the notification rule. AD users who wish to receive these emails must configure their email addresses in the notification rule. For more information about configuring alert notifications, see Configure Alert Notification Settings.

When the user reaches 100% of the quota, the (`UsersExceededQuota`) alert is triggered and the user cannot use additional capacity in the view until they reduce their usage or the quota is increased.

To enable user quotas:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on a view name to navigate to the corresponding view details page.
3. Click the **Quotas** tab.
4. In the **User Quotas** section, turn on the **Enable** toggle. The default logical quota for all users and the alert threshold are displayed.
5. To change the default quota for all users, click the edit icon, select different values and click **Save**.
6. To override the default quota for an individual user:
    a. Click **Add**.
    b. In the **Add User Quota Override** dialog, select a local user, AD domain, or Unix user ID from the **Principal** drop-down.

       If the view's storage domain is mapped to an AD domain, you can select only that AD domain and its trusted domains.
    c. Enter an AD user or a Unix ID.
    d. Select different values for the **Quota Override**.
    e. Click **Add**.
    f. Repeat these steps to add additional user quota overrides.

You can also use the `iris_cli user-quotas` commands to manage user quotas. For example, you can change the default 85% alert threshold. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

Directory Quotas

Directory quotas limit the capacity of directories in a view. You can set different quotas for different directories. For example, you can set a quota on the parent directory /home/sales, and set different quotas on the sub-directories /home/sales/kane and /home/sales/reynolds.

If a quota is set on a parent directory, the data usage in any of its nested sub-directories contributes to the quota for the parent. So if the data usage for /kane is 20 GiB and /reynolds is 10 GiB, they contribute 30 GiB to the quota set for /home/sales.

If the view is replicated to another Cohesity cluster, the directory quotas are available on that Cohesity cluster.

When a directory reaches the alert threshold (by default, 90% of its quota), the `DirectoriesExceededQuota` alert is triggered in the Cohesity cluster but data can still be written to the directory.

When the directory reaches 100% of its quota, the `DirectoryExceededQuotaAlertLimit` alert is triggered and data cannot be written to the directory until files are deleted or the quota is increased.

If email notifications are configured for alerts, emails are also sent.

Best Practices

- Enable directory quota at the time of directory creation. By doing so, you can ensure that the quota settings are in place from the beginning.
- After enabling the quota for a directory experiencing high IOPS, it is important to allow the current storage usage to be calculated before accessing the directory again. This allows the quota settings to be accurately applied and ensures proper management of storage usage within the directory.

To enable directory quotas:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on a view name to navigate to the corresponding view details page.
3. Click the **Quotas** tab.
4. In the **Directory Quotas** section, click **Add**. The directories in the view are displayed.
5. Click a directory folder and any sub-folders for which you want to set a quota.

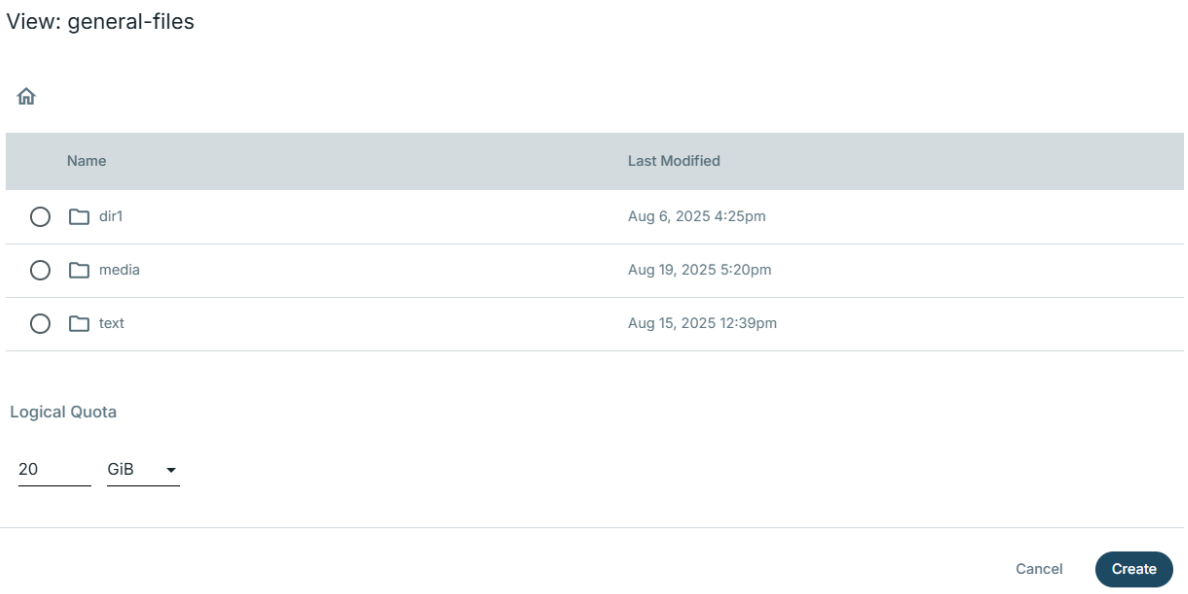

6. Set the **Logical Quota** for the selected directory.

7. Click **Create**.
   The folder and its quota information is displayed.

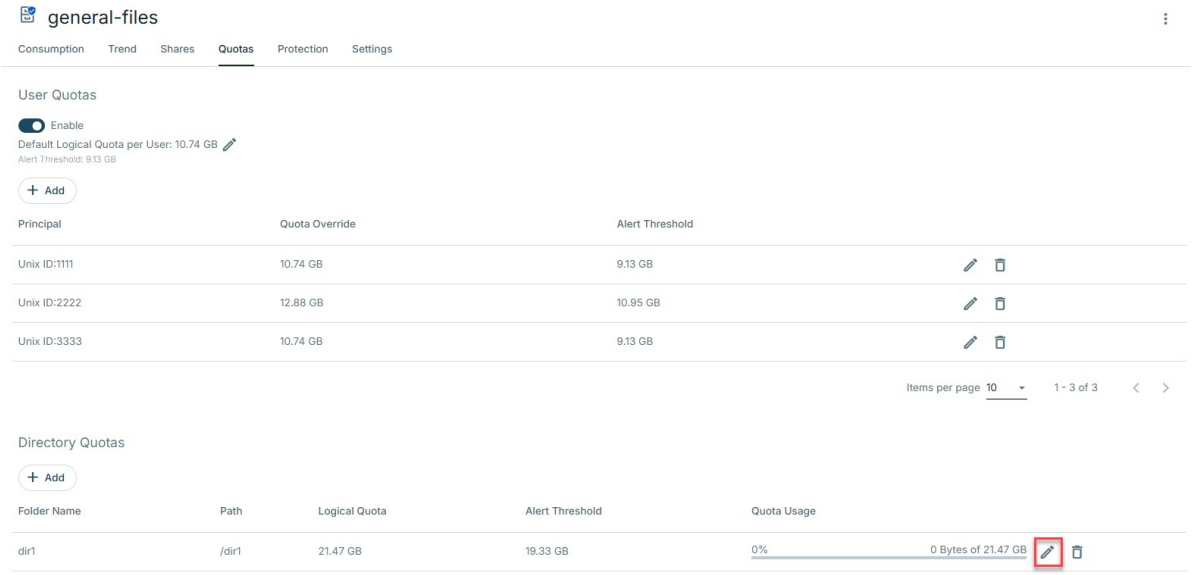

8. By default, the alert threshold is 90% of the quota. To change it, click the edit icon.
9. In the **Alert Threshold** column, change the quota and/or unit of measure and click **Save**.
10. Repeat these steps for additional directories as needed.

To delete an existing directory quota, click the action icon ⋮ for the quota and select **Delete**.

File-level DataLock for NFS and SMB Views

DataLock is the Cohesity WORM (write once read many) feature that locks and retains files in a view for compliance and regulatory purposes.

This topic describes the file-level DataLock setting, which differs from the view-level DataLock setting available when cloning a view.

When you create a view, you can enable the **File DataLock** setting to make the files in the view immutable based on the following settings:

- whether files are automatically locked after a specified period of I/O inactivity, based on the last modified time stamp
- whether files are manually locked when users or applications change the file permissions to read only or change the last accessed time, for example, by using the following Linux and Windows commands.

  Using NFS:

```
chmod -w <filename>
```

```
touch -at YYYYMMDDhhmm.ss <filename>
```

Using SMB, from PowerShell:

```
# (dir filename).Attributes = "ReadOnly"
```

```
# (dir filename).LastAccessTime = "YYYY-MM-DD hh:mm:ss"
```

In addition to file system locking, the Cohesity REST API can be used to lock files, update time stamps and get file lock status. Contact your Cohesity account team for details on the REST API.

You can also set:

- a default retention period for locked files
- an overriding lock expiration date for files

For more information about the File DataLock settings, see File DataLock.

If File DataLock is enabled for a view, the files in the view are in one of the following states: unlocked, locked or lock expired. The following table lists the actions permitted for each state.

Permitted Action	Unlocked	Locked	Lock Expired
Change the file	yes	no	no
Delete the file	yes	no	yes

> **Note:** The File DataLock option is only applicable for NFS and SMB View types.

Snapshot Self-Service

The Snapshots Self Service option enables the users to retrieve files and folders from the backup snapshots without depending on the administrator for retrieval.

You can enable the snapshot self-service for Views supporting SMB/NFS protocols when creating or editing a View. For more information, see Create Views Using a Predefined Template.

When you enable the Self-service option on a View, a hidden snapshot directory is created and the backup snapshots are listed in the hidden directory.

Using the snapshots from the hidden snapshot directory, users can copy/restore the files or folders that may have been deleted accidentally or corrupted.

> **Note:** The Self-service option applies to the Views that are protected. You can set the Snapshot Self-service option on Views without protection job enabled. However, the snapshots directory will be empty without the backup snapshots.

For NFS and SMB Views, you need to enable the Enable Snapshot Directory option.

When you enable the option with default settings, a .snapshot directory is created for NFS Views, and ~snapshot directory is created for SMB Views. You can customize the Snapshot directory name for NFS and SMB.

NFS users/clients can access the .snapshot directory and SMB users/clients can access the ~snapshot directory and copy the required files or folders.

For NFS Views, the .snapshot directory may not be visible on the Linux server after being mounted. To navigate to the .snapshot directory, run the `cd .snapshot` command inside the mount path.

The snapshot retention period in the protected views depends on the protection policy.

> **Note:** When you are using a name of your choice for the snapshot directory name, ensure that it is a unique name. If the name of the snapshot directory is the same as one of the folders in your View, you will not be able to see the backup snapshots.

Show Previous Versions

The **Show Previous Versions** option is available for SMB Views. This is another option to restore files and folders.

> **Note:** When you upgrade from any previous versions to Cohesity 6.6 or later, by default, the **Show Previous versions** option is enabled with access to **Everyone**.

Restoring Files or Folders from Snapshots via NFS/SMB

To restore a file or folder:

1. Access your View.
2. Open the **.snapshot** (for NFS) or **~snapshot**(for SMB).
3. Select a Timestamp from which files or folders need to be restored and copy the required files or folders.

Restoring Files or Folders from Previous Versions via SMB

To restore a file or folder:

1. Mount the View and right-click a View and select **Properties**.
2. In the Properties window, click the **Previous Versions** tab and click the snapshot list.
3. From the list of Snapshots, select a timestamp.
4. The Snapshot opens in a new window listing all files.
5. You can copy the required files.

Cohesity Views

A Cohesity view provides a storage location with NFS, SMB, and S3 mount paths in a storage domain on the Cohesity cluster.

You can use the NFS protocol to mount the NFS mount path onto systems. You can mount the SMB mount path onto systems as an SMB share.

You can use the view to store data such as files, backup snapshots, and cloned VMs. Views can only be mounted on any machines specified in the allowlist subnets. A view can also act as a datastore during the VM cloning process and during VM recovery. You can select the SWIFT protocol and a create SWIFT-only view for OpenStack environments.

You can create a view for the following three categories:

- File Shares
- Backup Target
- Object Services

You can select the **Read/Write** protocol and, optionally, select the **Read-only** protocol.

You can create a Cohesity view only by using a template. The template can be a predefined template or a custom template. For more information, see:

- Create Views Using a Predefined Template
- Create Views Using a Custom Template

> **Note:** If you are using Kerberos (AD or MIT), you must create a view using the same storage domain that is mapped to AD/MIT-Kerberos + LDAP.

Navigate to the **SmartFiles > Views** page to perform the following actions:

- Create views using a predefined template or custom template.
- Create custom view templates.
- Manage share aliases—View and manage Cohesity-created shares and any user-created shares on the **Shares** tab. For more information, see Manage Share Aliases.
- Configure global settings—Enable Active Directory authentication for SMB shares.

Specify subnets in the **Allowlist** that define ranges of IP addresses for systems that have privileges to access all views in the Cohesity cluster. For more information, see Add Subnets to Allowlists.

You can also clone a view, edit a view's settings, or delete a view.

The **Views** page also provides an overview of the views that have already been created. Views are grouped by storage domains. You can narrow the list by selecting a storage domain and/or typing part of a view name. To see a view's mount paths, click the view name.

Glance Bar

The glance bar provides a summary of the views:

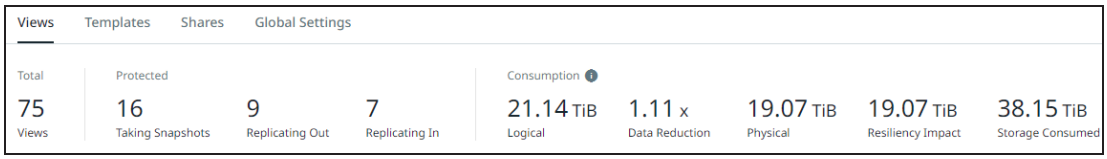

The following details are displayed:

- **Total**—The total number of views.
- **Protected**
  - The number of views taking snapshots.
  - The number of views replicating data out of the cluster.
  - The number of views replicating data into the cluster.
- **Consumption**
  - **Logical**—Size of total unique logical data of active file system as residing on Cohesity. This does not count the space consumed by the Data Protection runs (i.e. snapshots) that may be protecting the active file system.
  - **Data Reduction**—The ratio of Data In to Data Written. This ratio reflects the data deduplication and data compression operations, and demonstrates the true power and value of the Cohesity.
  - **Physical**—This number reflects unique data written, before resiliency operations. This is the unique logical data has been reduced by data deduplication and data compression by Cohesity.
  - **Resiliency Impact**—The space required to honor resiliency settings (Erasure Coding / Replication Factor) in addition to the data written. For example, if the Data Written is 100 GB and resiliency setting is RF3, the Resiliency Impact is 200 GB. That is, to honor the resiliency setting, you will need an additional 200 GB on top of the 100 GB for the Data Written.

- **Storage Consumed**—The total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for. This data can be stored either all locally or as a combination of local storage and in the Cloud Tier (if Cloud Tier is enabled). For example, if the Data Written is 100 GB and resiliency is set to RF3, leading to the need for an additional 200 GB of space, the resulting Storage Consumed is 300 GB. If the Data Written is 100 GB and resiliency is set to EC4:2, leading to the need for an additional 50 GB (as for every 4 data fragments, 2 parity fragments are added), the resulting storage consumed is 150 GB.

Views Table

The following table describes the columns on the **Views** page:

Column Name	Description / Action
View	The name of the view. Click on the view name to drill down and view the statistics and additional details.
Organization	The name of the organization that is associated with the view.
Protection	Displays the protection status.
Storage Domain	The name of the storage domain. Click on the storage domain name to drill-down and view the storage and performance statistics.
QoS Policy	The quality of service (QoS) specified for the view. By specifying a priority and the type of workload you expect, the Cohesity cluster can optimize how it processes data in regard to latency, throughput, and priority. For more information, see QoS Policies for Cohesity Views.

Column Name	Description / Action
Actions	To manage an existing view, select one of the following options: For a multitenant Cohesity cluster, you can only modify the objects that belong to your organization. • **Protect**–Create a Protection Group for the view. This action is available only for unprotected views. • **Clone** –Make a copy of the view. After cloning, the two versions of the data are unique. Changes made to the data in the new view are not replicated in the original view and changes made to the original view are not replicated in the new view. Data within a DataLock view is immutable. Follow the instructions provided in Clone a View. • **Mount Paths**–Copy the NFS, SMB, or S3 path to the clipboard. This path has to be used while mounting this view onto a local file system. For more information, see: • Mount a View Using NFS • Mount a View Using SMB **Note:** Multiple SMB paths are available if Active Directory has multiple account names; each machine account has its own path. Clicking this action copies the first SMB mount path in the list displayed on the View Details page. • **Edit**–Update the settings of a view. For more information, see Edit a View . • **Delete**–Delete the view. Deleting a view permanently deletes all the data stored on the view such as files or snapshots. If the view is used as a datastore for Cohesity cloned VMs, the VM files such as VMDK and VMX are also deleted. These cloned VMs can no longer run because these VMs can no longer access their VM files.

Filter Views Data

The **Views** page supports the following filters to pare down the data that you want to view in the table:

- **Organization**—Click the filter and select the organization from the list.
- **Category**—Click the filter and select the category: Backup Target, File Shares, or Object Services.
- **Protection Type**—Click the filter and select the protection type(s): Taking Snapshots, Replicating In, Replicating Out, or Archiving.
- **Protection Status**—Click the filter and select the protection status: Success, Warning, or Error.

Click **More Filters** to view the following filters:

- **Storage Domains**—Select the storage domain. You can also enter the name of the storage domain to filter by the storage domain name.
- **Protocol**—Select the protocol: NFS v3, NFS v4.1, SMB, S3, or SWIFT.
- **QoS Policy**—Select the QoS policy.

Hover over the search icon to view the **Search Views** search bar. Enter the name of the view to filter the views based on the name.

About Templates Tab

The **SmartFiles > Views** page includes a **Templates** tab. You can view the following details from the **Templates** tab:

- Number of total templates, predefined templates, and custom templates
- List of all predefined and custom templates

You can perform the following tasks from the **Templates** tab:

- Create a new custom template and modify existing custom templates
- Create a new Cohesity view
- Edit or delete custom templates

Cohesity supports creating customized view templates. While creating a view, you can configure settings according to your workloads and save these settings as a view template. You can use this customized template to create views. You can also edit or delete customized templates.

You can create a custom view template for the following categories:

- File Shares
- Backup Target
- Object Services
- Archive Services

Create a Custom View Template

You can create a custom view template from the **SmartFiles > Views > Templates** tab or the **SmartFiles > Views > View Templates** side sheet.

To create a custom view template from the **Templates** tab:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, go to the **Templates** tab.
3. Click **Create Template:**

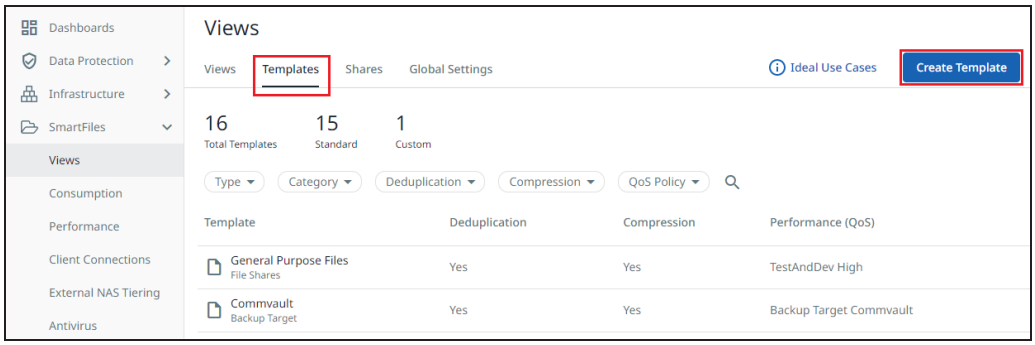

4. On the **Create View Template** dialog, complete the following:
    1. Enter a name for the template.
    2. Select a category:
        - File Shares
        - Backup Target
        - Object Services
        - Archive Services
    3. Create a new storage domain or select a storage domain from the list.

        > **Note:** Appropriate deduplication and compression settings are prepopulated while creating the storage domain.

    4. Based on the category that you selected, the custom template prepopulates a few default settings. You can update the settings based on your workload. Click **More Options** to update the default settings. For more information about the available options, see About Cohesity View Options.
5. Click **Create**.

To create a custom view template from the **View Templates** side sheet:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click **Create View**.
3. On the **View Templates** side sheet, click the **+** icon:

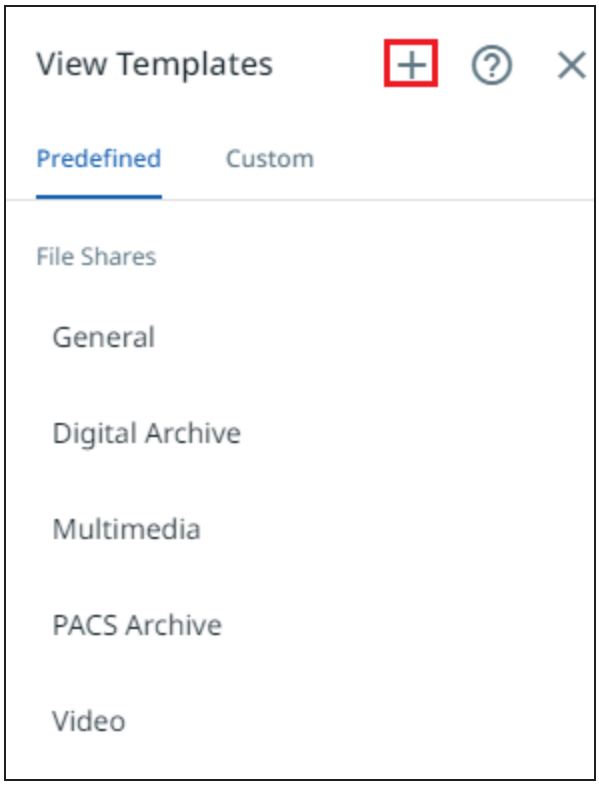

> **Note:** If no customized templates are created, you can also navigate to the **Custom** tab and click **Create Custom Template**.

4. On the **Create View Template** dialog, complete the following:
    1. Enter a name for the template.
    2. Select a category:
        - File Shares
        - Backup Target
        - Object Services
        - Archive Services
    3. Create a new storage domain or select a storage domain from the list.

        > **Note:** Appropriate deduplication and compression settings are prepopulated while creating the storage domain.

    4. Based on the category that you selected, the custom template prepopulates a few default settings. You can update the settings based on your workloads.

Click **More Options** to update the default settings. For more information about the available options, see About Cohesity View Options.

5. Click **Create**.

Create a View Template While Creating a View

To create a view template while creating a view:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click **Create View**.
3. On the **View Templates** side sheet, select a template.
4. Enter a name for the view.
5. Based on the template you selected, preconfigured settings are populated. Click **More Options** to update the settings. For more information about the available options, see About Cohesity View Options.
6. Click **Create View & Save as Template** to create the view and save the settings as a template:

7. On the **Template Name** dialog, enter the template name and click **Done**.

   The template is created.

View Custom Templates

To view custom templates from the **View Templates** side sheet:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click **Create View**.
3. On the **View Templates** side sheet, click the **Custom** tab.

   All custom view templates are displayed:

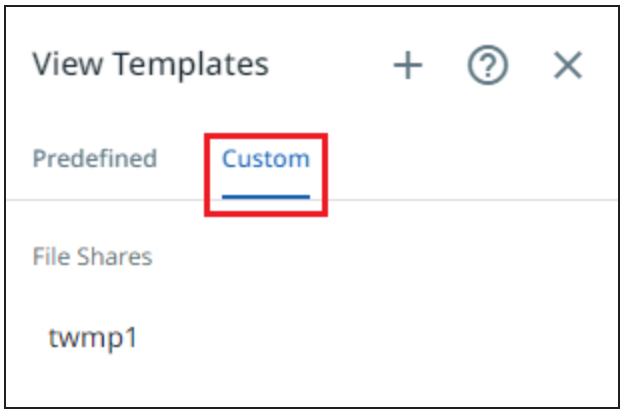

To view custom templates from the **Templates** tab:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, go to the **Templates** tab.
3. Click the **Type** filter and select **Custom**.

   All custom view templates are displayed:

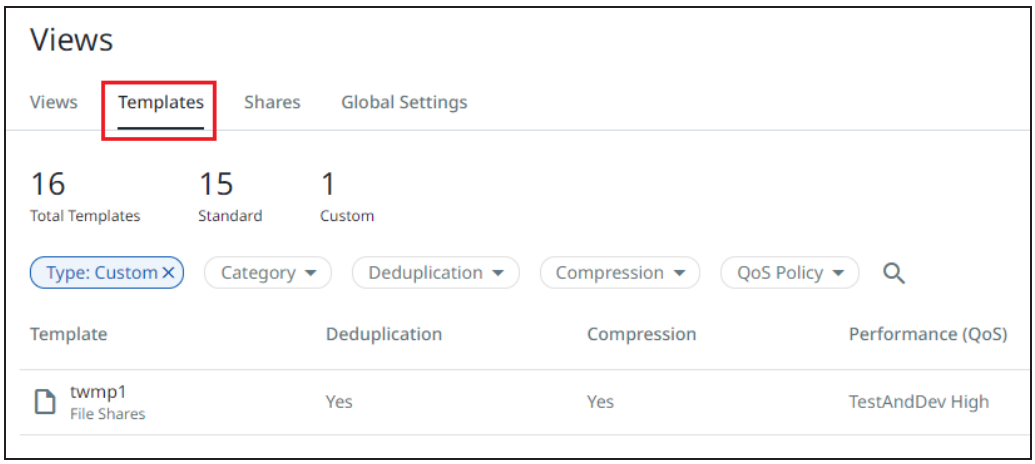

Edit a Custom View Template

To edit a custom view template:

1. Navigate to **SmartFiles > Views** and do one of the following:
   - On the **Views** page, click **Create View**. On the **View Templates** side sheet, navigate to the **Custom** tab and hover over the custom template that you want to edit. Click ⋮ and select **Edit**.
   - On the **Views** page, go to the **Templates** tab. Click the **Type** filter and select **Custom**. Hover over the custom view template that you want to edit. Click ⋮ and select **Edit**.
2. On the **Edit View Template** page, update the settings as required.
3. Click **Update** to save the changes.

   The updated settings are applied only in views that are subsequently created from this template.

Delete a Custom View Template

To delete a custom view template:

1. Navigate to **SmartFiles > Views** and do one of the following:
   - On the **Views** page, click **Create View**. On the **View Templates** side sheet, navigate to the **Custom** tab and hover over the custom template that you want to delete. Click ⋮ and select **Delete**.

- On the **Views** page, go to the **Templates** tab. Click the **Type** filter and select **Custom**. Hover over the custom view template that you want to delete. Click and select **Delete**.

2. On the **Delete Custom View Template?** dialog, click **Delete** to confirm.

Create Views Using a Predefined Template

You can create a view using predefined templates for different use cases and scenarios. Predefined templates are provided by Cohesity. Each category supports multiple templates. Templates have optimum preconfigured settings for a use case or scenario. You can select templates based on your requirements and create views faster. You can create views using a predefined template for the following categories:

- File Shares
- Backup Target
- Object Services
- Archive Services

Predefined templates and default properties for each predefined template are listed below.

Default Properties for Predefined Templates

The following table lists the default properties for predefined templates. When you create a view using predefined templates, these optimal settings for use cases are selected by default.

> **Note:** The view workflow suggests you create a new storage domain if there is no storage domain that exists with the necessary deduplication and compression settings based on the predefined template.

View Category	View Template	Properties		
		Deduplication	Compression	Performance (QoS)
File Shares	Multimedia	No	No	TestAndDev High
	Video	No	No	Backup Target High
	General Purpose Files	Yes	Yes	TestAndDev High

View Category	View Template	Properties		
Backup Target	Commvault	Yes	Yes	Backup Target Commvault
	Veeam	Yes	Yes	Backup Target SSD
	General	Yes	Yes	Backup Target High
	TSM	Yes	Yes	TestAndDev High
	Application Dumps (Oracle, SQL, SAP Hana)	Yes	Yes	TestAndDev High
Object Services	Containers	Yes	Yes	TestAndDev High
	General	Yes	Yes	TestAndDev High
	Splunk Smartstore	Yes	Yes	TestAndDev High
	Hadoop	Yes	Yes	TestAndDev High
Archive Services	General Archive	No	No	TestAndDev High

To create a view using a predefined view template:

1. Navigate to **SmartFiles > Views** and do one of the following:
   - On the **Views** page, click **Create View**. On the **View Templates** side sheet, navigate to the **Predefined** tab and select a predefined template.
   - On the **Views** page, go to the **Templates** tab. Click the **Type** filter and select **Predefined**. Hover over the predefined view template that you want to use, click ⋮, and select **Create View**.

   > **Note:** After you select the predefined template, you cannot modify the category.

2. Enter a name for the view.
3. Based on the template you selected, the predefined template populates optimum preconfigured settings. However, if you want to update the settings, click **More Options** and update the settings. For more information about the available options,

see About Cohesity View Options.

4. Click **Create** to create the view.

> **Note:** Click **Create View & Save as Template** to create the view and save the settings as a custom view template. On the **Template Name** dialog, enter the template name and click **Done**. The customized template is created.

Create Views Using a Custom Template

Cohesity supports creating customized view templates and using a customized template to create views. For more information, see About Templates Tab.

You can create a view using a custom template for the following categories:

- File Shares
- Backup Target
- Object Services
- Archive Services

To create a view using a custom template:

1. Navigate to **SmartFiles > Views** and do one of the following:
   - On the **Views** page, click **Create View**. On the **View Templates** side sheet, navigate to the **Custom** tab and select a custom template.
   - On the **Views** page, go to the **Templates** tab. Click the **Type** filter and select **Custom**. Hover over the custom view template that you want to use, click ⋮, and select **Create View**.

> **Note:** After you select the custom template, you cannot modify the category.

2. Enter a name for the view.
3. Based on the template you selected, the customized template populates preconfigured settings. However, if you want to update the settings, click **More Options** and update the settings. For more information about the available options, see About Cohesity View Options.
4. Click **Create** to create the view.

> **Note:** Click **Create View & Save as Template** to create the view and save the settings as a custom view template. On the **Template Name** dialog, enter the template name and click **Done**. The customized template is created.

About Cohesity View Options

Protection

Turn on the **Protection** toggle to specify the protection group and protection policy. If enabled, all views created using this template would use this protection group and adhere to the selected protection policy. In the policy, you can also define how the objects should be replicated.

- **Protection Group**—Use this option to protect objects as a group, or individually for granular level control. You can either create a new protection group or select an existing protection group from the drop-down list.
  - **New Group**—Enter a new group name. This will create a new protection group on the first use of the template.

    > **Note:** Ensure that you specify a group name that is not in use.

  - **Existing Group**—Select a protection group from the drop-down list.
- **Storage Domain**—Select a storage domain to be used during protection. You can also create a new storage domain.
- **Protection Policy**—Select a policy from the drop-down list. The policy is read-only. If you want to update the policy, go to the **Data Protection > Policies** page to edit the policy.

Storage Lifecycle Management Rules

You can reduce storage space utilization by managing the file timestamp attributes. You can define criteria for identifying aging files and specify actions to take upon meeting these conditions. By automating the deletion of aged files, storage lifecycle management ensures optimal performance and resource utilization, enabling you to reclaim valuable storage space.

> **Important:** Cohesity strongly recommends enabling **File Services** audit logs on Views on which storage lifecycle management is enabled.

Click **No Lifecycle rules have been added yet** and click **Add Rule**:

In the **Add Storage Lifecycle Management Rule** dialog, do the following:

- **Name**—Enter a name for the rule.
- **Type**—Select one of the following:
  - **Allow Rule**—Files matching this rule are deleted.
  - **Deny Rule**—Files matching this rule are excluded from storage lifecycle management, even if the files match other rules.
- **Scope**—Select one of the following:
  - Select the **Apply to all files in the View** option to apply this storage lifecycle management rule on every file within the View.
  - Select the **Limit the scope of this rule using filters** option to limit the scope by using the following filters:
    - **Filter by File Extension**—Enter the file extension. For example, you can enter .pdf to restrict the scope to only PDFs stored in the View.
    - **Filter by Minimum File Size**—Specify the minimum file size limit.
    - **Filter by Maximum File Size**—Specify the maximum file size limit.
- **Aging Policy**—Select one of the following:
  - **Number of Days**—Specify the number of days.
  - **Date**—Select a date.
  - **None**—If you have selected the type as **Deny Rule**, choosing **None** ensures that the aging policy is not applied to the deny rule.

  To calculate aging, you can select Last Modification Time or Creation Time.
  - **Creation Time**—Refers to the timestamp when a file is initially created and saved to the storage medium.
  - **Last Modification Time**—Refers to the most recent timestamp when a file's data was modified. This timestamp updates whenever you modify the data in the file.

All the rules added are listed in the **Storage Lifecycle Management Rules** section.

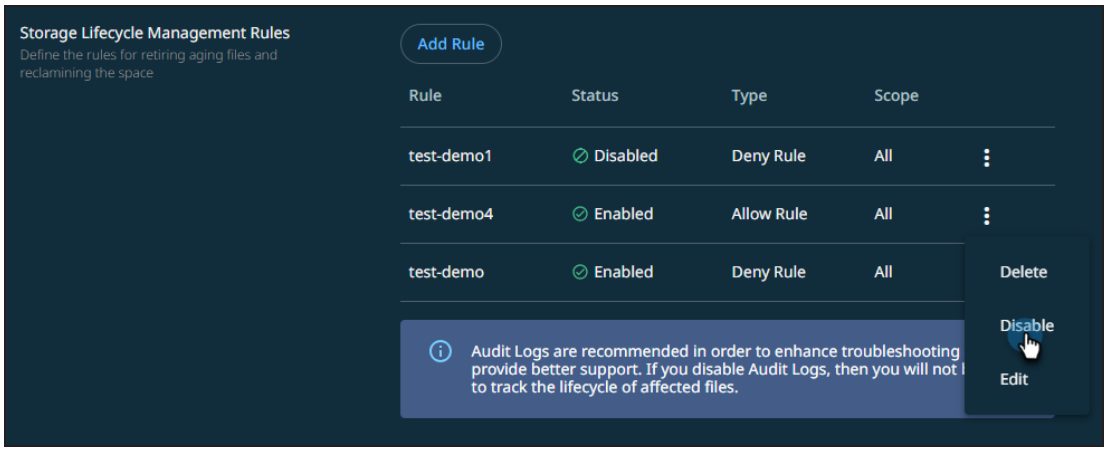

The following details are displayed:

- **Name**
- **Status**: Enabled or disabled
- **Type**: Allow rule or deny rule
- **Scope**

Hover over the rule and click the vertical ellipsis to perform the following actions:

- **Delete**: Allows you to delete the rule
- **Enable** or **Disable**: Allows you to enable or disable the rule

  > **Note:** Only the rules that are enabled are considered during the execution of lifecycle management.

- **Edit**: Allows you to edit the rule

The tables below provide several examples of the aging policy:

- Allow Rule Examples
- Deny Rule Examples

Allow Rule Examples

Aging Policy	Description
	Files that have not been modified in the last **15** days are deleted.
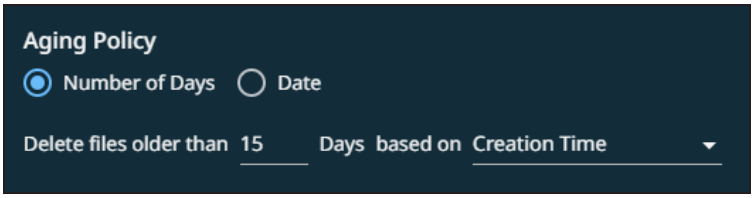	Files created more than **15** days ago are deleted.
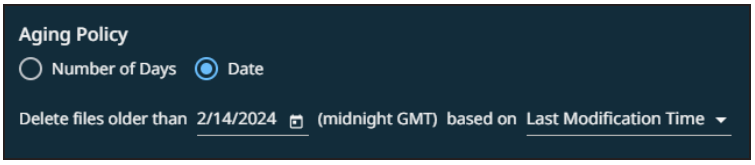	Files modified before **February 14, 2024**, are deleted.
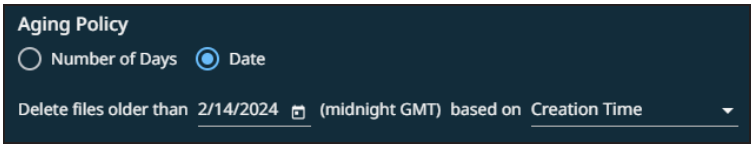	Files created before **February 14, 2024**, are deleted.

Deny Rule Examples

Aging Policy	Description
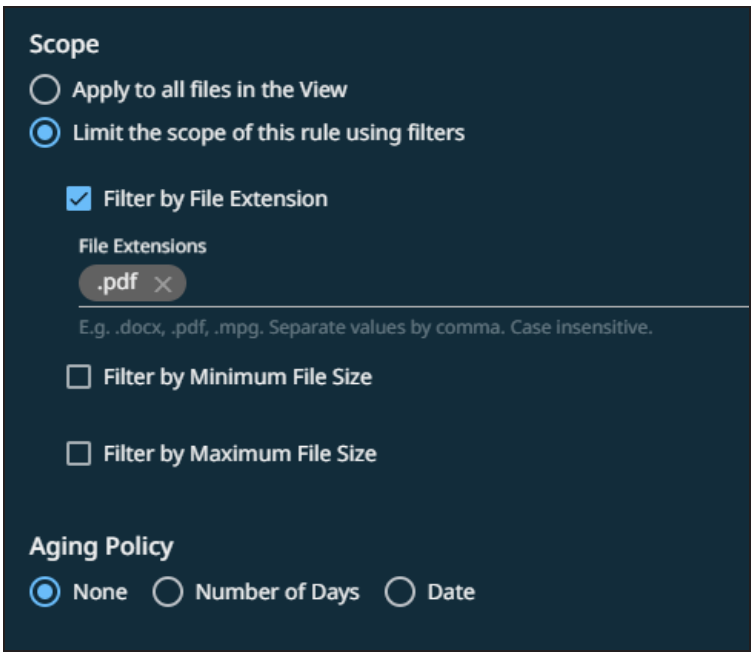	PDF files are not deleted.
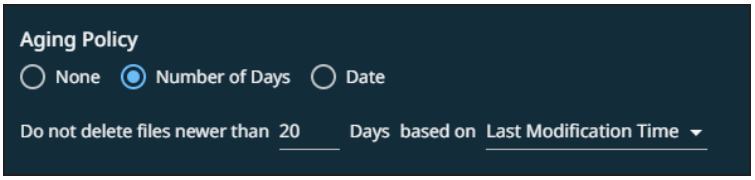	Files modified within the last **20** days are not deleted.
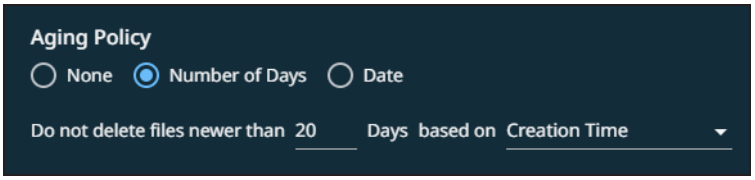	Files created within the last **20** days are not deleted.
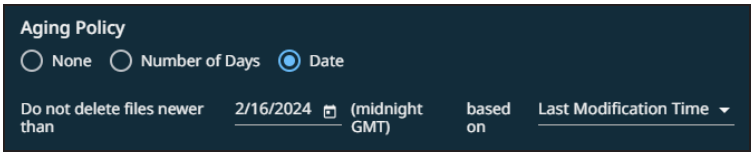	Files modified after **February 16, 2024**, are not deleted.
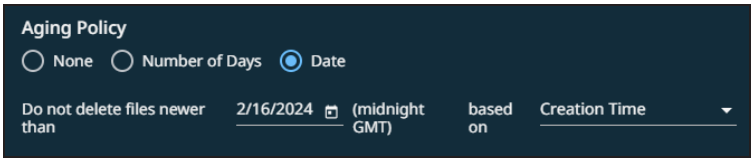	Files created after **February 16, 2024**, are not deleted.

Important Points to Note

- When specifying file extensions, make sure to include the period (.) before the extension. For example, use ".pdf" instead of just "pdf".
- Regardless of the sensitivity of the View, file extensions are treated as case-insensitive and are handled the same way. For example, if a case-sensitive View contains files named f1.pdf, f2.PDF, and f3.Pdf, and a rule specifies to preserve only ".pDf" files, then all three files (f1.pdf, f2.PDF, and f3.Pdf) will be preserved.
- If a file meets the criteria for deletion based on a storage lifecycle management rule while it is open, the system still proceeds with deleting the file.
- Cohesity strongly recommends backing up Views on which storage lifecycle management is enabled. If any files are unintentionally deleted, you can recover them from the backup.
- The system honors **File DataLock** settings and storage lifecycle management rules are only applied to files that are not locked.
- If you replicate a View that includes storage lifecycle management rules, the rules are disabled since the View is read-only on the remote cluster. If you intend to use this View by cloning it, note that the storage lifecycle management rules will be present, but you must manually enable them on the View.

  > **Note:** If the remote cluster does not support Storage Lifecycle Management, the storage lifecycle management rules are not replicated to the remote cluster.

- When you archive and subsequently restore a View, the storage lifecycle management rules are not restored along with the View.
- When you clone a View, the storage lifecycle management rules associated with that View are also cloned and lifecycle actions are executed on this cloned View according to the rules.
- Users with the privilege to edit a View can add or update **Storage Lifecycle Management Rules** on that View.
- Storage lifecycle management rules applied to a View are executed on all eligible files within that View, regardless of the individual file's access permission settings.
- If you intend to enable **Storage Lifecycle Management Rules** on a View designated as a down-tiering target, Cohesity recommends not to enable lifecycle management on this target View if the down-tiering task uses the **Tier with symbolic links** option. Enabling lifecycle management under such circumstances may result in dangling symlinks on the primary source when files are deleted by lifecycle management on this target View. For more information about Symbolic Links, see Configure a Down-Tiering Task.

- By default, lifecycle job execution occurs across the cluster every seven days. During this process, the job identifies all Views governed by lifecycle policies and identifies files eligible for deletion according to the policy of each View. If a significant number of files are eligible for deletion, it may take multiple runs of the lifecycle job to remove them all. Cohesity focuses on prioritizing other critical workloads on the server; as a result, the job refrains from being more aggressive in deleting all files in a single run. Therefore, if a lifecycle job does not delete all files in one go, note that the subsequent job runs will eventually complete the task. Although Cohesity can schedule the lifecycle job to run daily, doing so would consume resources needed by active and higher-priority workloads on the Cohesity server. Hence, the default periodicity of once every seven days is maintained.

Lifecycle Management Details in Audit Logs

The **Cluster** audit logs feature captures all events that occur at the cluster level. To review storage lifecycle management details, you can download the CSV file from the **System > Audit Logs > Cluster** tab.

The following is an example of an audit log record:

```
filerLifecycleManagement:
  rules:
  - agingPolicy:
      agingCriteria: LastModificationTime
      days: 180
    fileFilter:
      fileExtensions:
      - .docx
      - .pdf
      fileSize:
        maxBytes: 1073741824
        minBytes: 1048576
    name: rule1
    status: Enabled
    type: Allow
```

The **File Services** audit logs feature allows you to capture and monitor events happening on your SMB and NFS Views. These logs document actions such as file deletions, last modification times, creation times, file sizes, and other pertinent details. Timestamps are included in these logs, enabling you to cross-reference with cluster audit logs and determine the most recently active storage lifecycle management rules.

The following is an example of a **File Services** audit log record:

```
{"Timestamp" : "2023-12-12T17:35:33.918Z", "RecordID" : "2cda:6b", "Protocol" : "SnapFS", "RequestType" :
"Delete ClientIP" : "10.14.52.77", "ViewName" : "test_view2", "Fsid" : 8817 ", "Result" : "Ok", "EntityType"
: "File", "EntityPath": "f_100kb_90day.pdf", "EntityInode" : 8035, "EntityModifyTimestamp" : "2023-09-12T01:
30:09.000Z", "InodeSizeBytes" : 102400, "SourceParentEntityInode" : 8815, "LCMRuleName" : "r1",
"EntityCreationTimestamp" : "2023-12-12T17:29:51.860Z"}
```

- `EntityPath` displays the name of the file that is being deleted.
- `LCMRuleName` displays the name of the storage lifecycle management rule.

Case Sensitive File or Folder Names

Turn on the **Case Sensitive File or Folder Names** toggle if you want the uppercase and lowercase version of the same letter in a file or folder name to be treated differently on the view. For example: `A` is different from `a`.

> **Note:** You cannot edit this setting after you create the view.

Performance

QoS Policy—Select the QOS policy. A QoS policy determines to which storage media data is written and the priority of I/O when contention occurs. For more information, see QoS Policies for Cohesity Views.

Pin View to SSD—Turn on the toggle to pin the view to the SSD tier so that critical data and data that has to be retrieved quickly are not downtiered to HDD or cloud. For more information, see Pin Views to SSD.

Security

Configure the security settings for the view. You can add and manage **Global IP Allowlist**, **Netgroup Allowlist**, and select **Security Style** for the view.

IP Allowlist

1. Select **Override Global IP Allowlist** to ignore the global IP allowlist and honor only specific entries provided in the view. Select **Extend Global IP Allowlist** to honor the global IP allowlist and the specific entries provided in the view.
2. Click **Add**.
3. In the **Add Allowlist** dialog, enter the subnet in CIDR format.
4. Select **NFS Permissions**, **SMB Permissions**, and **S3 Permissions**.
5. Select the **NFS Squash** option: **None**, **Root**, or **All**.
   For more information, see NFS Squashing.
6. Optionally, you can enter a description for the IP allowlist of the view.

Netgroup Allowlist

1. Select **Override Global Netgroup Allowlist** to ignore the global netgroup allowlist and honor only specific entries provided in the view. Select **Extend Global Netgroup Allowlist** to honor the global netgroup allowlist and the specific entries provided in the view.
2. Click **Add**.
3. In the **Netgroup Name** field, enter a name.
4. Select **NFS Permissions**.
5. Select the **NFS Squash** option: **None**, **Root**, or **All**.

For more information, see NFS Squashing.

6. Click **Add**.

Security Style

Select the **Security Style** for the view. This setting is available only if you select both NFS and SMB protocols.

- **Native**—Windows and Linux permissions are stored separately.
- **Unified**—Windows and Linux permissions are kept consistent such that when one changes, the other is modified automatically.
- **NTFS**—Access via SMB and NFS is controlled by an SMB ACL present on the file.

> **Note:** Permission mapping occurs only if the cluster has joined an Active Directory domain that has **User ID Mapping for Multi Protocol Network Access** enabled. For more information, see ID Mapping for Multi-Protocol Views.

Dedupe & Compression

Turn on the **Dedupe & Compression** toggle to always use post-processing, irrespective of the storage domain policy. This option is available only if deduplication and compression are enabled on the storage domain.

Logical Quota

Select one of the options to set the **Logical Quota** for the view:

- **Inherit from Storage Domain**
- Select **Override**, turn on the following toggles, and set the limit—**Logical Quota** or **Alert Threshold**

File DataLock

Turn on the **File DataLock** toggle to make the files in this view immutable based on the settings described below. File-level DataLock is typically used for compliance and regulatory purposes.

> **Important:** You can enable **File DataLock** only while creating the View. You cannot enable **File DataLock** on an existing View.

> **Note:** The **File DataLock** option is only applicable for NFS and SMB view types.

- **File DataLock Mode**—Select one of the following modes:
  - **Compliance**—Once File DataLock is enabled for the view, it cannot be disabled by anyone. However, a Cohesity user with the Data Security privilege can edit other DataLock settings. A user with the Data Security privilege can extend the override date into the future, but no one can shorten the date. Compliance mode means the view complies with SEC Rule 17a-4.
  - **Enterprise**—A Cohesity user with Admin privileges or Data Security privileges can rename the view, delete the view, extend or shorten the override date, switch to Compliance mode, and edit the File DataLock settings. This mode is useful for giving storage administrators additional control over the view.

    > **Note:**
    > - Apart from the local admins, Domain Admins also can delete files in a view with File DataLock enabled and operating in Enterprise Mode.
    > - Enterprise mode is supported on 6.4.1 and later clusters. If the view is replicated to a pre-6.4.1 cluster, Enterprise mode is automatically changed to Compliance mode. If the view is then replicated back to a 6.4.1 or later cluster, it retains the Compliance mode setting.

- **Default Lock Period**—Indicates the default time period that locked files remain locked if a user or application does not specify a default retention period. When the time period ends, files can be deleted but not changed.
  If you set the **Default Lock Period** to **Forever**, you cannot override it by enabling the **Override (Lock all files until a specific date)** option.
- **Autolock Files**—Indicates whether files are automatically locked after the specified period of I/O inactivity (based on the last modified timestamp).
  Override is supported only if **Autolock Files** is enabled. If **Autolock Files** is disabled, override cannot be used as legal hold.
- **Manually Lock Using**—By default, manual file locking occurs when a user or application sets the file's read-only attribute. Select one of the following:
  - **Read-Only**—Locks files when the file permission is set to read-only. For example, set the Windows file properties attribute or use the Linux `chmod` command.
  - **Future ATime**—Lock files when the last access time (**atime**) is set. For example, use Windows PowerShell commands or the Linux `touch` command.

  Auto lock and manual lock are not mutually exclusive. You can have both the options active at the same time.

- **Minimum Lock Period**: Optional. Indicate a minimum time for which the files remain locked if the user or application does not specify a default retention period.

  If an application or user locks the files for a default period and the application locks the files for a minimum period that is less than the default period, the request is silently rejected and the files remain locked as per the default retention period.
- **Maximum Lock Period**: Optional. Indicate a maximum time to ensure that a file's last access time cannot be set outside this period. A request to do so results in a permission denied error if the manual locking mode is **FutureATime**. Otherwise, the request will be silently rejected and the file will be locked for the default retention duration.

  This is useful for storage administrators to reclaim the storage space by specifying a maximum lock period after which the files can be either deleted or archived.

  > **Note:** By default, the **Maximum Lock Period** is set to **None** and cannot be edited if the **Default Lock Period** is **Forever**.

  The **Minimum Lock Period** and **Maximum Lock Period** fields are valid for both **FutureATime** and **Read-Only** manual lock.
- **Override (Lock all files until a specific date)**—Indicates whether the lock on files expires on a specific date. This setting is typically used for legal hold purposes. Files in a lock expired state cannot be deleted until this date.
  If the **Default Lock Period** is set to **Forever**, overriding only extends the period for the defined amount of time and does not actually override the files.

File Filtering

Turn on the **File Filtering** toggle if you want to block or allow files with specified file extensions from being written to this view. Select one of the following options:

- **Denied File Types**—If you want to specify a list of file extensions for files that are blocked from being written to this view. For example, if you specify a list with the `.exe` and `.dll` extensions, files with these extensions cannot be written to this view.
- **Allowed File Types**—If you want to specify a list of file extensions for files that are allowed to be written to this view. For example, if you specify a list with the `.cvs` and `.txt` extensions, only files with these extensions can be written to this view.

> **Note:** Special characters are not permitted in file extension names.

Enter file extensions separated with commas, spaces, tabs, or newlines and click **Add**.

NFS Options

You can control NFS access at the host IP, subnet, or netgroup level. The most restrictive permissions are applied.

- **Discoverable Shares**—When creating a Cohesity view, NFS shares are set as not discoverable by default. To change this default behavior, enable the **Discoverable Shares** toggle. Subsequently, you can see the discoverable shares in the `showmount` command output.
- **Root Permissions**—When creating a Cohesity view, the initial permissions are as follows:
  - Read, write, and execute permissions for the root user (UID=0)
  - Read and execute permissions for the group root (GID=0)
  - Read and execute permissions for others

  The Cohesity cluster administrator can set UID/GID for the root of the Cohesity view:

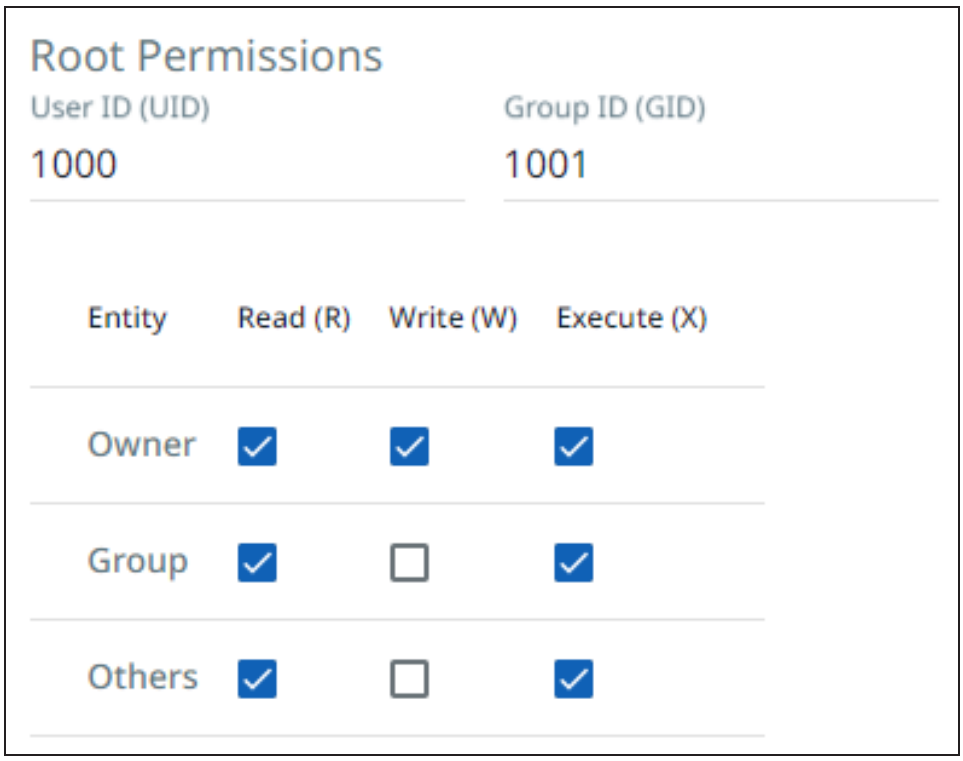

- Root Squash
- All Squash

Supported NFS Security Modes

Select at least one network authentication protocol:

- Unix Authentication
- Kerberos Authentication - Client needs to mount with sec=krb5.
- Kerberos Integrity - Client needs to mount with sec=krb5i.
- Kerberos Privacy - Client needs to mount with sec=krb5p.
  For more information, see NFS Security Modes for Kerberos Environments.

> **Note:** The authentication provider assigned to the storage domain is used for Kerberos authentication. For more information, see Kerberos.

SMB Options

Click the **SMB Options** to enable or disable the following SMB options:

- **Browsable Shares**—Turn on the toggle if you want shares of this view to be discoverable by the Windows `net view` command.

  Windows behavior prevents Cohesity SMB shares from being automatically discoverable. Use the `net view` command to probe the cluster explicitly using \\<Cluster-machine-account-name> or \\<Cluster-vip-FQDN> or \\<Cluster-VIP>.
- **Access Based Enumeration**—If this setting is enabled, then when a user browses an SMB share, only the files and folders that a user has privileges to access are displayed.
- **SMB Encryption**—Turn on the toggle to enable encryption. Turn on the **Enforce Encryption** toggle to disallow any SMB2 or SMB3 connections that do not have encryption enabled. Turn off the toggle to allow all connections.
- **Fast Durable Handles**—Turn on this toggle to improve performance, but continuous availability is lost. For more information, see Fast Durable Handles.
- **SMB Oplocks**—By default, the **SMB Oplocks** toggle is turned on when you create a view. If SMB Oplocks is enabled:
  - Client-side caching of read-ahead, write-behind, and lock information improves SMB share performance.
  - File and directory leasing also gets enabled on newer clients such as Windows 2012 or later or Windows Server 2012 R2 or later. SMB leasing enables the client to cache file and directory metadata, which improves the overall file sharing performance.

    For more information, see SMB Oplocks.
- **Offline File Caching**—The SMB offline feature allows the SMB client to use files offline when the SMB clients are disconnected from the network. For more information, see SMB Offline.
- **NTFS Root Permissions**—For an SMB share, by default, everyone has full control permissions for the view. You can optionally define custom permissions for the view. Search for and select the principal (user or group) that you want to define permissions for. Select **Allow** or **Deny**. Then select the permissions and the folder/file level to apply them to. Click **Add**. Permissions adhere to the prescribed behavior of ACEs in a DACL. For more information, see Manage Access Using NTFS Permissions.

> **Note:** While editing a view, changes to the permissions apply only to new files and folders that are created after the change. Changes to permissions are not applied recursively to the view's existing files and folders.

If the view's storage domain is mapped to an AD domain, you can select only that AD domain and its trusted domains.

- **Share Level Permissions**—For an SMB share, you can optionally define custom permissions for the share. The share-level permissions are stored on the cluster as view metadata and are independent of the NTFS permissions.

  Search for and select the principal (user or group) that you want to define permissions for. Select **Allow** or **Deny**. Click **Add**. Permissions adhere to the prescribed behavior of ACEs in a DACL.

- **Super Users**—An SMB Super User is a privilege for AD users and groups, and local SMB users and groups. Select the AD domain and principal (user or group). Click **+** to add additional super users.

  A user with this privilege has complete access to view or share content overriding any assigned ACLs. However, the privilege cannot override the global allowlist and file-level data lock permissions. The privilege is set at the share-level for a view and not as NTFS root permissions. When the Super User privilege is assigned to a user for a specific view or share, the user has complete access to all the contents of the view or share. For more information, see SMB Super User.

> **Note:** The most restrictive permissions are applied. The permissions can be **Full Control**, **Modify**, **Read Only**, or **Super User**. For example, if NTFS permissions are Read and Write, but the share permission is Read, the Read permission is applied.

Snapshot Self-Service

Click **Snapshot Self-Service** and define the following to enable the users to retrieve files and folders from the backup snapshots without depending on the administrator for retrieval. For more information, see Snapshot Self-Service.

- In the **NFS Snapshot Options** section, turn on the **Enable snapshot directory** button. Optionally, you can specify a name for the snapshot directory.

  > **Note:** When you use a custom name for the snapshot directory name, ensure that the name is unique and is not present inside the view.

- In the **SMB Snapshot Options** section, you can turn on the **Show previous versions** and **Enable snapshot directory** buttons. Optionally, you can specify a name for the snapshot directory. Click **Add** and set the SMB snapshot access permissions.

  > **Note:** When you use a custom name for the snapshot directory name, ensure that the name is unique and is not present inside the view.

Antivirus

If antivirus software is installed on the Cohesity cluster, you can enable antivirus scanning on the view. For more information, see Antivirus Software to Protect Cluster Data.

Audit Logs

Turn on the **Audit Logs** toggle to enable audit logs. For more information, see File Services Audit Logs.

Description

Click **Description** and provide a description for the view.

Manage Access Using NTFS Permissions

You can use NTFS permissions to manage access to the files and folders that are stored in NTFS file systems. For an SMB share, by default, everyone has the **Full Control** permission level for the view. You can optionally define custom permissions for the view.

To define custom permissions for an SMB share, complete the following steps:

1. Navigate to **File Services > Views > All Views**.
2. On the **All Views** page, select **Create View > New View**.
3. Enter a name for the view.
4. Select one of the following categories:
   - **File Shares**
   - **Backup Target**
5. Click the **Read/Write Protocol** drop-down list and select **SMB**.
6. Click **More Options**.
7. Click **SMB Options**.
8. In the **NTFS Root Permissions** section, click **Add**.
9. In the **Add NTFS Root Permissions** pop-up window, do the following:

1. Select the AD domain from the drop-down list.
2. Select the Principal from the drop-down list. The list of users is displayed under the **Users** category. The list of groups is displayed under the **Groups** category.

   Cohesity supports the following well-known local groups:

   - Administrators
   - Users
   - Guests
   - Power Users
   - Backup Operators

     For more information, see Supported Windows Well-Known SIDs.
3. Select the type of access:
   - Allow
   - Deny
4. Click the **Applies To** drop-down list and select one of the following:
   - This folder, subfolders, and files.
   - This folder and subfolders
   - This folder and files
   - This folder
   - Subfolders only
   - Files only
5. Select the permission level:
   - Full Control
   - Read & Write
   - Modify
   - Read Only
6. Click **Add**.

The permissions are added and the details are displayed in the Permissions table.

For more information about creating a view, see Create Views Using a Predefined Template.

To update NTFS permissions in an existing view, complete the following steps:

1. Navigate to **File Services > Views > All Views**.
2. On the **All Views** page, hover over the view that you want to edit, click the **More Options** icon, and select **Edit**.

3. On the **Edit View** page, click **SMB Options**, and update the permissions.
4. Click **Update**.

NFS Squashing

You can configure NFS root squash or all squash at the subnet, netgroup, or view level. Subnet/netgroup level squashing supersedes view-level squashing.

Root Squash

By default, root squashing is enabled on a Cohesity view and root user requests are processed as though it is made by the anonymous user (UID=65534). If you disable squashing, the root user on the client gets the same level of access to files as the root user on the Cohesity view. Root squashing allows the root user (UID=0) to be trusted by clients defined in a given allowlist.

All Squash

All squash maps all users, including root, to a defined UID/GID. The default UID used for squashing is 65534.

Related Topics

- About Cohesity View Options
- Global Settings for NFS
- Update the Global Allowlist

Edit a View

You can edit a View and modify the View its settings. Some of the View settings cannot be modified or updated.

To edit a View, complete the following steps:

1. Navigate to **SmartFiles > Views**.
2. In the **All Views** page, for the required View click the ⋮ icon and select **Edit**.
3. In the **Edit View** screen, modify the required View properties. For more information, see Cohesity Views.
4. Edit the required properties and click **Update**.

> **Note:** You cannot edit most of the properties of a view if the View is created using View templates and some of the View properties cannot be edited for any View type.

> **Important:** You cannot edit the **File DataLock** option. You can enable **File DataLock** only while creating the View.

Pin Views to SSD

Cohesity allows you to pin views to the SSD tier. This feature enables you to pin views to the SSD tier so that critical data and data that has to be retrieved quickly are not downtiered to HDD or cloud. You can do the following:

- Pin a view to an SSD when you create the view.
- Edit an existing view and pin the view. In this case, only the new data written after the view is pinned is retained on the SSD tier.
- Pin a view for an indefinite period of time. Select **Forever** if you want to set this option. The data resides on the SSD forever.
- Pin a view for a definite period from the time of its last access. You can select **Hours**, **Days**, **Months**, or **Years** and specify an appropriate number in the field. The view is pinned for the specified duration from its last access time. After the specified period, the data is automatically downtiered to HDD.
- Only the following QoS policies support pinning views:
  - Test And Dev High
  - Backup Target SSD

  For more information, see QoS Policies for Cohesity Views.
- Unpin a view from an SSD.

> **Note:** If the system has used the reserved SSD capacity, view pinning is not honored and the data is downtiered to HDD.

Pin a View to the SSD

To pin a view when creating a view, complete the following steps:

1. Navigate to **SmartFiles > Views > All Views**.
2. Click **Create** and select **View**.

3. In the **Create View** pop-up window, click **More Options**.
4. Click **Performance**:
5. From the **QoS Policy** drop-down list, select one of the following QoS policies:
    - **Test And Dev High**
    - **Backup Target SSD**
6. Enable the **Pin View to SSD** toggle button.
7. Select one of the following options to set the duration:
    - **Forever**—Pins the view forever. The data resides on the SSD forever.
    - **Hours**—Enter the number of hours for which the view should be pinned.
    - **Days**—Enter the number of days for which the view should be pinned.
    - **Months**—Enter the number of months for which the view should be pinned.
    - **Years**—Enter the number of years for which the view should be pinned.

    > **Note:** The view is pinned for the specified duration from its last access time.
8. Complete the other steps required to create a view. For more information, see Create Views Using a Predefined Template.
9. Click **Create**.

To pin an existing view, complete the following steps:

1. Navigate to **SmartFiles > Views > All Views**.
2. Hover over the view that has to be pinned to the SSD, click the **More Options** icon, and select **Edit**.
3. In the **Edit View** page, click **Performance**.
4. From the **QoS Policy** drop-down list, select one of the following QoS policies:
    - **Test And Dev High**
    - **Backup Target SSD**
5. Enable the **Pin View to SSD** toggle button.
6. Select one of the following options to set the duration:
    - **Forever**—Pins the view forever. The data resides on the SSD forever.
    - **Hours**—Enter the number of hours for which the view should be pinned.
    - **Days**—Enter the number of days for which the view should be pinned.
    - **Months**—Enter the number of months for which the view should be pinned.

- **Years**—Enter the number of years for which the view should be pinned.

> **Note:** The view is pinned for the specified duration from its last access time.

7. Click **Update**.

> **Note:** An algorithm runs in the background every 12 hours. When the pinning properties are updated on the view, only after the algorithm runs, the data is downtiered to HDD if it is cold (rarely accessed) for more than the specified duration.

Unpin a View from the SSD

You can unpin a view manually if the data is no longer required to be retrieved quickly. After you unpin the view, the data stored on the SSD is downtiered to HDD.

> **Note:** An algorithm runs in the background every 12 hours. After the view is unpinned, the next run of the algorithm clears the pinning properties, and the data is downtiered to HDD.

To unpin a view from the SSD, complete the following steps:

1. Navigate to **SmartFiles > Views > All Views**.
2. Hover over the view that has to be unpinned from the SSD, click the **More Options** icon, and select **Edit**.
3. In the **Edit View** page, click **Performance**.
4. Disable the **Pin View to SSD** toggle button.
5. Click **Update**.

Clone a View

You can make a clone of an existing view. After cloning, the two versions of the data are unique. Changes made to the data in the new view are not replicated in the original view and changes made to the original view are not replicated in the new view. Data within a DataLock view is immutable. You can also clone a view from a view Snapshot. For more information, see Clone a View From a Snapshot.

To clone a View:

1. Navigate to **SmartFiles > Views**.
2. Hover over the view that you want to clone, click ⋮ and select **Clone**.
3. Select the clone's **Purpose**.

   Select **Clone as DataLock View** if you want to clone a view as write once read many. Data within a DataLock view is immutable, and most view properties cannot be edited. (Note that view-level DataLock is different than file-level DataLock, as described in File-level DataLock for NFS and SMB Views.)
4. In the **Task Name** field, you can optionally change the default name for the clone task.
1. **View Name:** If you are creating or cloning a new view, specify a name for the view. Forward slashes (/) are not allowed in the view name.

   > **Important:** Renaming a view changes its default share name. Clients won't be able to mount the view unless they use the new share name. Clients using a share alias are not affected by the rename.

   For a DataLock view, select a date for lock expiration. During the lock period, data within the view cannot be edited by any user.

   During the lock period, users with Data Security privileges can delete the view (unless File DataLock is enabled for the view, in which case the view cannot be deleted), extend the lock period (it cannot be reduced) , add IPs in the allowlist and change the description. After the lock period expires, the view can be deleted by any user that has permission to delete it.
5. The Storage Domain is already selected. The new view is always created in the same storage domain as the original view.

   If the selected storage domain has Inline Deduplication enabled, an option appears that lets you override the default storage domain Inline Deduplication setting for the current view. Select this option if desired. Cohesity recommends enabling Inline Deduplication.

   If Inline Deduplication is on, deduplication occurs when the Cohesity cluster writes the data to the Cohesity View. If Inline Deduplication is off, deduplication occurs after the Cohesity cluster writes the data to the Cohesity View.

   For more information, see About Deduplication.

   1. **QoS Policy:** A QoS Policy determines to which storage media data is written and the priority of I/O when contention occurs. Select an appropriate QoS policy for the view. By specifying a priority and the type of workload you expect, the Cohesity cluster can optimize how it processes data in regard to latency,

throughput, and priority.

For details and recommendations, see QoS Policies for Cohesity Views.

2. To have privileges to mount and access a view, a system must be specified in the allowlist subnet. For more information, see Add Subnets to Allowlists.

6. Click **Continue**.

You can view the cloned Cohesity view on the **Test & Dev** page.

The view is ready to store data or act like a datastore during the VM recovery and VM cloning processes.

After cloning, the two versions of the data are unique. Changes made to the data in the new view are not replicated in the original view and changes made to the original view are not replicated in the new view. Data within a DataLock view is immutable.

Clone a View From a Snapshot

You can clone a view directly from a snapshot of the type view. This clone action creates a new view with a copy of the data captured in the snapshot. After cloning, the two versions of the data are unique. Changes made to the data in the new view are not replicated in the original view and changes made to the original view are not replicated in the new view. Data within a DataLock view is immutable.

You can clone a view from a snapshot created using any of the following actions:

- A snapshot captured directly from a view on the current Cohesity cluster.
- A snapshot located on a remote Cohesity cluster that was replicated (copied) from the Cohesity cluster that captured the snapshot.
  When the first Snapshot is replicated to the remote Cohesity cluster, a DR view is created on this remote Cohesity cluster. A DR view is a placeholder view that does not store data and cannot be cloned or mounted. You can use this DR view to access the replicated Snapshots using the Cohesity UI on the receiving (remote) Cohesity cluster. You can create a new view on the remote Cohesity cluster by cloning a replicated Snapshot associated with the DR view. In addition, when the first Snapshot is replicated to the remote Cohesity cluster, an Inactive Job is also created on the remote Cohesity cluster. You can use an Inactive Protection Group of type View to access the replicated Snapshots using the Cohesity UI on the remote Cohesity cluster that received the Snapshots. An Inactive Protection Group of type View is a placeholder Job and you cannot failover this type of Inactive Protection Group.

  While creating a Protection Group for a view, if you have enabled the option to create a view on the remote Cohesity cluster, then along with the snapshot, a view will be created on the remote Cohesity cluster from the latest snapshot during the first Protection Group run. The same view will be reused and will always point to the latest snapshot after every run

To learn more, see Protect Cohesity Views.

You can also clone a view directly from an existing view. For more information, see Clone a View.

To clone a view from a snapshot:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on the view name that has been backed up into a snapshot that you want to clone.

   > **Note:** If you are on a Cohesity cluster with replicated snapshots that were captured on another Cohesity cluster, you can access the replicated snapshots from a special view created on the remote Cohesity cluster called a **DR View**. A DR view is a placeholder view that does not store data and cannot be cloned or mounted. You can use this DR view to access the replicated Snapshots using the Cohesity UI on the receiving (remote) Cohesity cluster. You can create a new view on the remote Cohesity cluster by cloning a replicated Snapshot associated with the DR view.

3. On the view details page, click the **Protection** tab.
4. Hover over the snapshot that you want to clone into a new view, click ⋮, and select **Clone from Local Target**

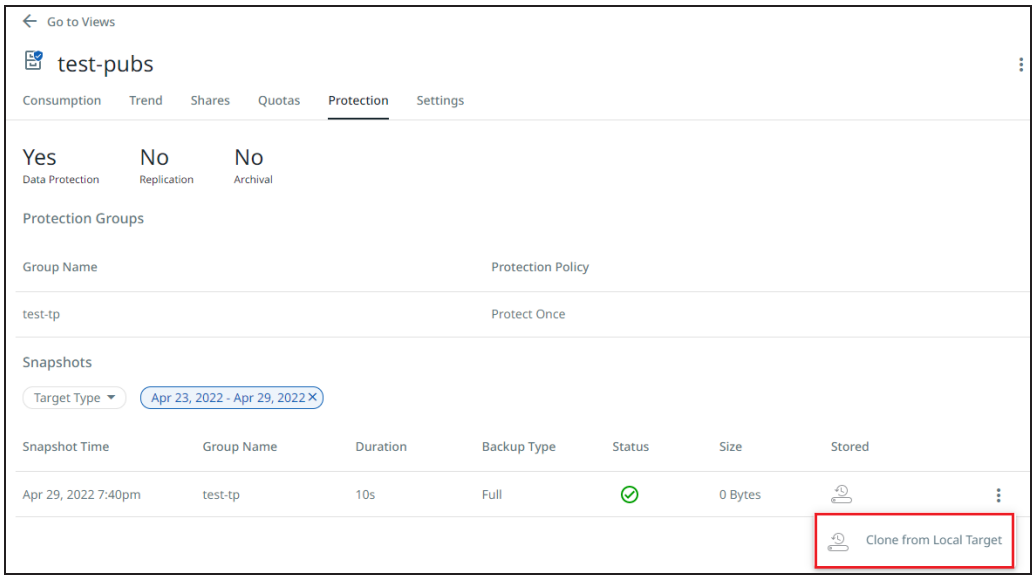

   1. Select the clone's **Purpose**.

      Select **Clone as DataLock View** if you want to clone a view as write once read many. Data within a DataLock view is immutable, and most view properties

cannot be edited. (Note that view-level DataLock is different than file-level DataLock, as described in File-level DataLock for NFS and SMB Views.)

2. In the **Task Name** field, you can optionally change the default name for the clone task.
3. **View Name:** If you are creating or cloning a new view, specify a name for the view. Forward slashes (/) are not allowed in the view name.

   > **Important:** Renaming a view changes its default share name. Clients won't be able to mount the view unless they use the new share name. Clients using a share alias are not affected by the rename.

4. For a DataLock view, select a date for lock expiration. During the lock period, data within the view cannot be edited by any user.

   During the lock period, users with Data Security privileges can delete the view (unless File DataLock is enabled for the view, in which case the view cannot be deleted), extend the lock period (it cannot be reduced) , add IPs in the allowlist and change the description. After the lock period expires, the view can be deleted by any user that has permission to delete it.
5. The Storage Domain is preselected. The new view is created in the same Storage Domain as the cloned snapshot.

   If the selected storage domain has Inline Deduplication enabled, an option appears that lets you override the default storage domain Inline Deduplication setting for the current view. Select this option if desired. Cohesity recommends enabling Inline Deduplication.

   If Inline Deduplication is on, deduplication occurs when the Cohesity cluster writes the data to the Cohesity View. If Inline Deduplication is off, deduplication occurs after the Cohesity cluster writes the data to the Cohesity View.

   For more information, see About Deduplication.
6. **QoS Policy:** A QoS Policy determines to which storage media data is written and the priority of I/O when contention occurs. Select an appropriate QoS policy for the view. By specifying a priority and the type of workload you expect, the Cohesity cluster can optimize how it processes data in regard to latency, throughput, and priority.

   For details and recommendations, see QoS Policies for Cohesity Views.
7. To have privileges to mount and access a view, a system must be specified in the allowlist subnet. For more information, see Add Subnets to Allowlists.

5. Click **Continue**.

You can view the cloned Cohesity view on the **Test & Dev** page.

The view is ready to store data or act like a datastore during the VM recovery and VM cloning processes.

Clone the Contents of a Directory in a View

You can recursively clone (copy) the contents of a directory in a View to a new directory in the same View or a different View using the `view clone-directory` Cohesity CLI command.

Limitation

The default value for the maximum number of files in the source hierarchy that can be atomically cloned together is 8192. Contact Cohesity Support to increase the maximum number of files in the source hierarchy that can be atomically cloned together.

To copy the contents of a directory in a View:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Clone the contents of a directory using the Cohesity CLI `view clone-directory` command:

   ```
   admin>view clone-directory destination-directory-name=<DestDirName>
   destination-directory-path=<DestPath> source-directory-
   path=<SourcePath>
   ```

   > **Tip:** You can use the CLI Help feature to view the `view clone-directory` options by entering the following command: `view clone-directory help`.

CLI Option	Placeholder Name	Example	Description
`destination-directory-name=`	**<DestDirName>**	`Dest`	Specify the name of the new directory where the source directory will be recursively copied, under the specified destination path. The new directory will be a child under the destination directory and must not already exist.
`destination-directory-path=`	**<DestPath>**	`/ViewDest`	Specify the directory path of the destination directory where the cloned directory will be created. You must start the path with a slash (`/`) and include the View name in the directory path. You can optionally preface the path with the Storage Domain name.

CLI Option	Placeholder Name	Example	Description
`source-directory-path=`	**<SourcePath>**	`/ViewSrc`	Specify the full path of the source directory to clone (or recursively copy) the contents from. You must start the path with a slash (`/`) and include the View name in the directory path. You can optionally preface the path with the Storage Domain name.

For example, the following directory structure and files exist prior to running the example CLI command:

- `/ViewSrc/test/file.txt`
- `/ViewSrc/test/src/text.txt`
- `/ViewDest`

The following example command is run:

```
view clone-directory destination-directory-name=Dest destination-directory-path=/ViewDest source-directory-path=/ViewSrc/test
```

After the above command is run, the following directories and files exist in the Views.

- `/ViewDest/Dest/file.txt`
- `/ViewDest/Dest/src/text.txt`
- `/ViewSrc/test/file.txt`
- `/ViewSrc/test/src/text.txt`

Protect Cohesity Views

You can protect the data in a view by capturing snapshots of the view. A view Protection Group captures snapshots according to the protection schedule defined by the associated protection policy. Snapshots can be used as follows:

- The snapshot can be cloned into a new view that is stored in the same storage domain as the original snapshot. The new view contains the version of the files and directories as they were captured when the original snapshot was created.
- The snapshot can be replicated (copied) to another storage domain on a remote Cohesity cluster. The replication schedule is defined by the protection policy associated with the Protection Group.

Other backup products, such as Veeam, can store data in a view.

The steps below provide a sample scenario of how view snapshots are captured, cloned and replicated:

1. A view Protection Group captures a snapshot of all files and directories stored in a view. In this example, all the files and directories currently stored in the `myView` view are captured in the `3/21/24 01:00 AM` snapshot.

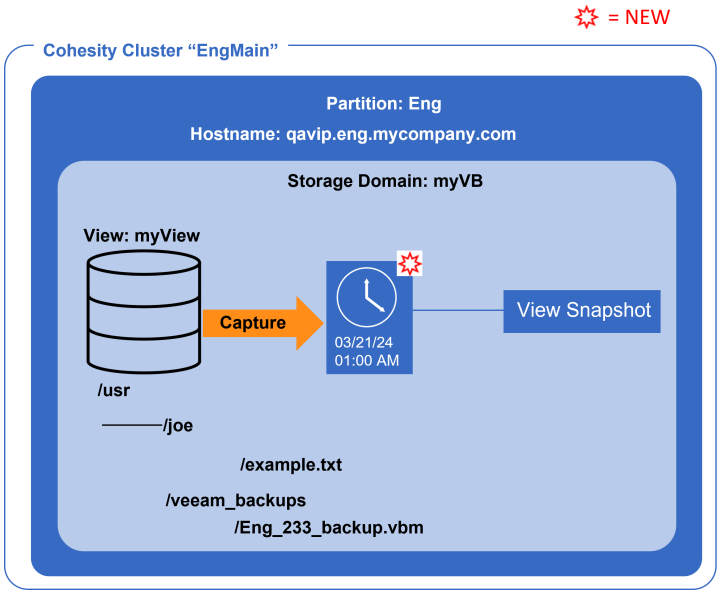

2. If a replication schedule is defined for protection policy associated with Protection Group that created the snapshot, then the following would be created on the remote Cohesity cluster ("EngReserve, below):
    - A copy of the snapshot is replicated from the Cohesity cluster capturing the snapshots ("EngMain", below)
    - A DR view is created when the first snapshot is replicated to the remote Cohesity cluster. A DR view is a placeholder view that does not store data and cannot be cloned or mounted. You can use this DR view to access the replicated

snapshots using the Cohesity UI on the receiving (remote) Cohesity cluster. You can create a new view on the remote Cohesity cluster by cloning a replicated snapshot associated with the DR view. In addition, when the first snapshot is replicated to the remote Cohesity cluster, an Inactive Job is also created on the remote Cohesity cluster.

- A new view with Read-Only permission is created from the latest replicated snapshot (if you have enabled the option to create view on the remote Cohesity cluster). This view can have any user-specified name provided while creating Protection Group for a view and will always point to the latest replicated snapshot. If a disaster strikes, you can clone this view instead of the latest snapshot. However, if you want to restore a data set that is not the latest, then you must clone the view from the respective snapshot.

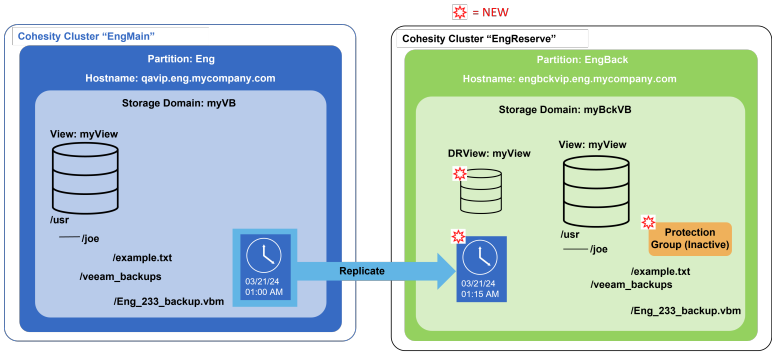

3. In the original `myView` view, a new file called `Eng_400_backup.vbm` is created at 1:30 AM.

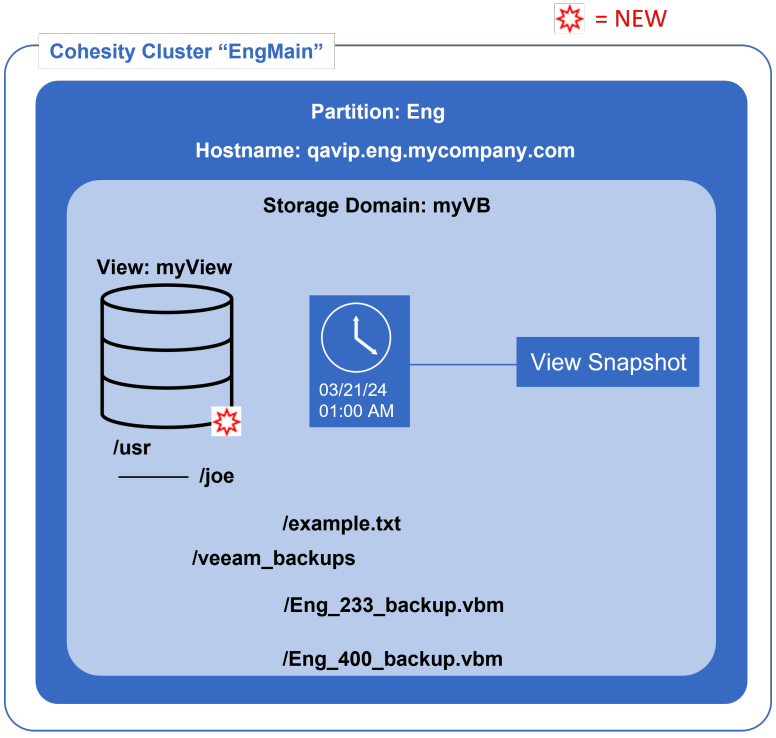

4. At 2:00 AM, a new view snapshot is captured. The 2:00 AM snapshot has captured the new file called `Eng_400_backup.vbm` but the original 1:00 AM snapshot does not have the file.

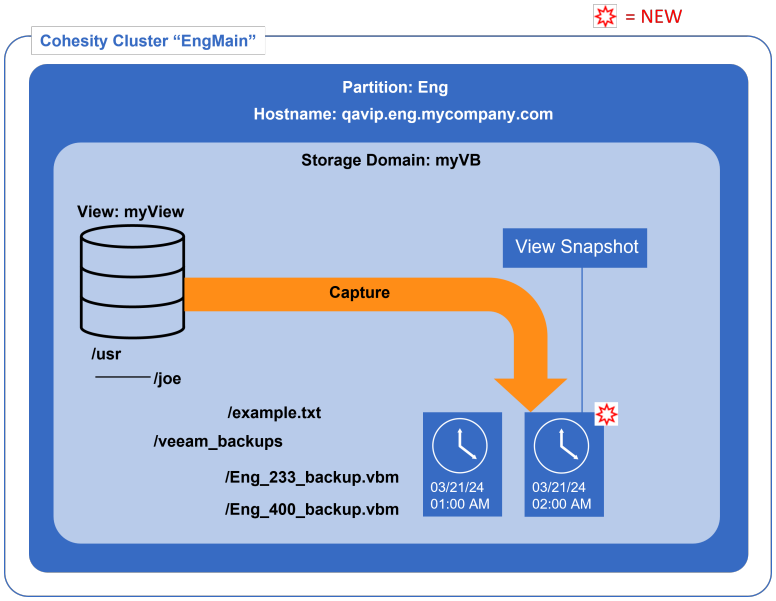

5. You clone the 1:00 AM snapshot in the capturing Cohesity cluster. A new view (called myViewClone) is created with the versions of the files and directories as originally captured by the 1:00 AM snapshot. The new view is always created in the same

Storage Domain as the original view. The new view (called myViewClone) does not contain the Eng_400_backup.vbm file.

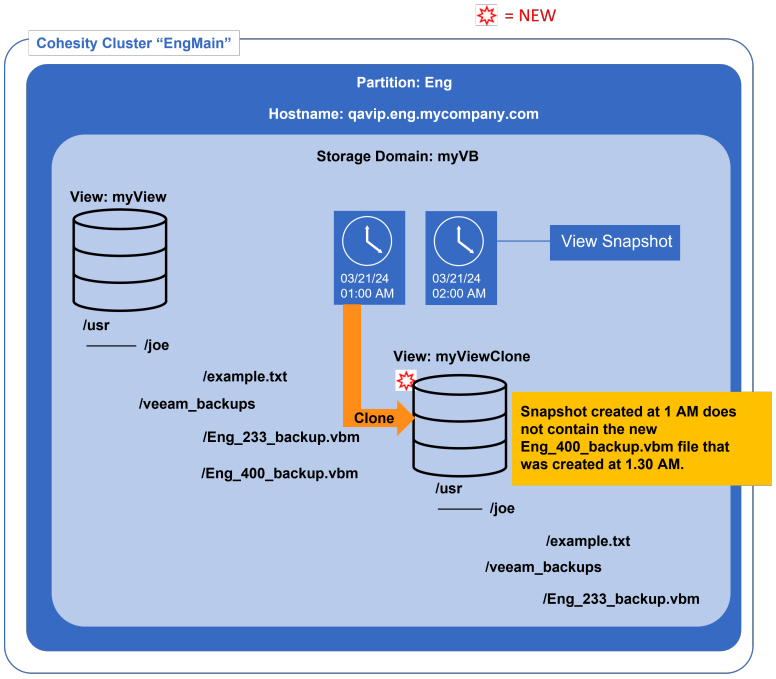

6. As defined by the replication schedule of the Job, along with the latest view a second copy of the snapshot is replicated to a remote Cohesity cluster at 2:15 AM. In this example, every snapshot captured on the Cohesity cluster ("EngMain", below) is replicated to the remote Cohesity cluster ("EngReserve", below).

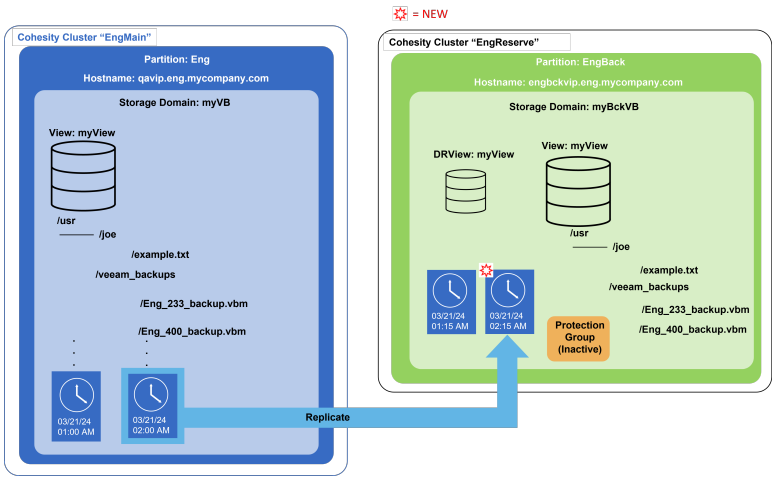

7. If disaster strikes and the original view is no longer available, you can recover from the disaster by cloning the snapshot stored on the remote Cohesity cluster. When the View Snapshot is cloned, a new View is created that contains the version of the files and directories as they were captured when the Snapshot was created. In this example, since the new view created on the remote Cohesity cluster ("EngReserve",

below) was replicated from the second snapshot, the new view (called `myBckVM`) contains the `Eng_400_backup.vbm` file.

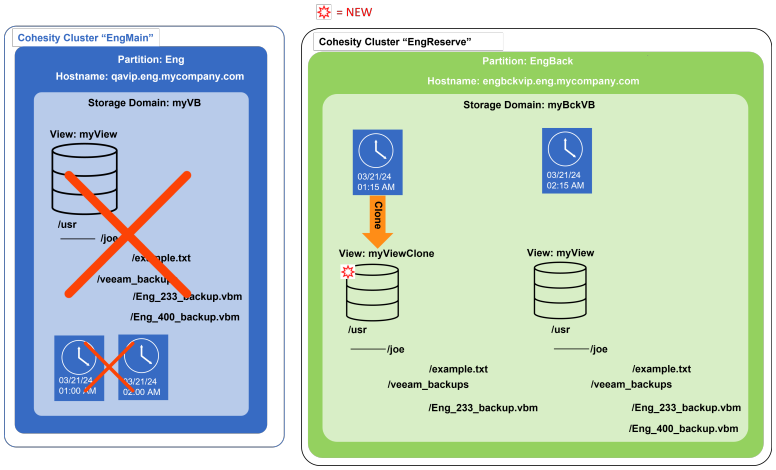

Add or Edit a Protection Group for a View

The Protection Group captures snapshots of views and snapshots of views are required for the following tasks:

- creating a new view from a view snapshot
- replicating snapshots
- archiving snapshots

The Cohesity cluster supports backing up the following entities stored in the view:

- ordinary files
- directories
- symbolic links
- hard links

These entities can be captured in a view snapshot and restored to a new view. The backup and recovery of the following special files is supported: pipes, block devices, sockets and character devices.

To create a Protection Group for a view:

1. Navigate to **Data Protection > Protection** and do one of the following:
   - To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
   - To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. If creating a new job, select **Protect > Cohesity View**.

3. **Views:** Select the view that you want to protect and then click **Save Selection**. The list shows the Storage Domain that contains the view. View Snapshots are always stored in the same Storage Domain as the original view.
4. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
5. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.
6. **Remote View**: This option is available only if the policy you selected has replication enabled. In the **Remote cluster** fields, enter a name for each view that will be created on the remote Cohesity cluster corresponding to the views on the primary Cohesity cluster. Ensure that the name provided for the views are unique and are not the same as any other view names on the remote Cohesity cluster. If the name provided is not unique, then on the remote Cohesity cluster, a new view will be create with"-1" appended to the provided name. In the next protection run, the views in the remote Cohesity cluster will be overwritten by the latest replicated snapshot of the view on the primary Cohesity cluster.
1. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
7. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

   The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
8. **Indexing** :Indexing is enabled by default and is required for file recovery. Optionally customize by adding indexing rules. Type index text. Click **Add** under **Include** and **Exclude** fields to add more index terms.
9. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
10. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
11. **Alerts:** Optional. Select one or more of the following settings if you want Alerts to be created for the following triggers:

- **Success:** Create an Informational Alert when a Protection Group completes successfully. Emails are not sent when Informational Alerts are created.
- **Failure:** Create an Critical Alert if the Protection Group fails to complete. Emails are sent when Critical Alerts are created.

For more information about Alerts, see Alerts.

12. **Description:** Optional. Click the plus sign and enter a description for the Protection Group.
13. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Recover Files or Folders

You can now recover the files and folders of a view that has been protected. You can either recover the files and folder of a view to an alternate view or to the same view. This section describes how to perform the recovery of a view.

Considerations

- You can select only those views as targets that have the **Case Sensitive File or Folder Names** option same as the source view. For example, if the file of a view you selected for recovery has the **Case Sensitive File or Folder Names** option disabled, then from the **Select View** drop-down, you can select only those views having the **Case Sensitive File or Folder Names** option disabled.
- You can recover files or folders from a view having the **File DataLock** option enabled. However, you cannot recover files or folders to a view having the **File DataLock** option enabled.
- The target view and the source view must belong to the same storage domain.
- You can recover files or folders from an NFS view to an SMB view and vice-versa. However, symlinks will not be preserved for cross protocol recoveries.
- You can recover files or folders from an archive and replicated view.
- You cannot recover files or folders to a Read-only view.
- You cannot recover files or folders from an S3 view or to an S3 view.

Recover to the Original View

To recover files and folders to the original view:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Cohesity View > Files or Folders**.
3. Select any one of the following options to search for the file or folder:

- **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
- **Browse**: Allows you to search and recover files using the name of the view or Protection Group from any backup available on the Cohesity cluster.

If you select..	Then..
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. Select a date range and the snapshot point for recovery. **Note:** • If you have selected multiple files for recovery, then any files not present in the newly selected snapshot will not be recovered. • By default, the snapshot is recovered from the local cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the Location option. 3. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial view or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. By default, only the files and folders that are indexed are displayed. To display all the available files and folders, toggle off the **Browse on Indexed Data** option. For S3 Views, both indexed and non-indexed browse are supported. By default, the **Browse on Indexed Data** option is disabled. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Once you select the files and folder to recover, select one of the following options: • **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task completes, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file. • **Next**: If you select this option, then continue with next step to configure the file recovery options. 6. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps:

If you select..	Then..
	1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. • For S3 Views, only indexed browsing is supported for archived snapshots; non-indexed browsing is not supported. • Recovery from archived snapshots of an S3 View is not supported. 7. Click **Select Recovery Point** and go to step 4.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     > **Note:** The download option is disabled for S3 Views.

Character replacement explanation for : and \ in file and directory names in zip files

If the original file or directory name contains a ‘:’ , it is replaced with ‘_COLON_’ in the archive.
Example:

- The original name is "a:b" then the archived name is "a_COLON_b"
- The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

If the original file or directory name contains a ‘\’, it is replaced with '_BACK_SLASH_' in the archive.
Example:

- The original name is “x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **Original view** to recover the files or folders to the same view.
6. To recover to a different location in the original view, enable the **Recover to alternate path** option, and then provide the location to which the files or folders are to be recovered in the **Alternate Path** field. By default, the files and folders will be recovered to the original location.
7. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is disabled to create the files and folders in the specified location. Any duplicate files are skipped. Enable this option to overwrite the existing files and folders.

     Consider this example, you are recovering a file, **1.txt** from the source view V1 to the root directory of the target view V2. If a file, **1.txt** already exists in the root directory of the target view V2, then the **Overwrite** behavior is as follows:

     - If the **Overwrite** option is enabled during recovery, the **1.txt** in the target view is overwritten with the data in the **1.txt** file of the source view V1.
     - If the **Overwrite** option is disabled during recovery, the **1.txt** in the target view V2 will not be overwritten with the data in the **1.txt** file of the source view V1.

     If **Overwrite** is enabled, recovering a file to the source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
   - **Preserve File/Folder Attributes**: This option is not applicable for the file or folder recovery of a view.

- **Continue on Error**: By default, this option is enabled to continue the recovery even if one of the objects encounters an error. Disable this option to discontinue the recovery task if one of the objects encounters an error.
- **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
- **Task Name**: Change the default name of the recovery task.
- **Save Success Files**: Enable this option to download the success logs for a recovery run.

8. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover to an Alternate View

To recover files and folders to the alternate view:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** located at the upper-right corner of the page and select **Cohesity View > Files or Folders**.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse**: Allows you to search and recover files using the name of the view or Protection Group from any backup available on the Cohesity cluster.

If you select..	Then..
Files and Folders	1. Enter characters of the file name, path, or specify the wildcard character *. A list of items is displayed. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. Select a date range and the snapshot point for recovery. **Note:** • If you have selected multiple files for recovery, then any files not present in the newly selected snapshot will not be recovered. • By default, the snapshot is recovered from the local cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the Location option. 3. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial view or Protection Group name. You can specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, or Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. By default, only the files and folders that are indexed are displayed. To display all the available files and folders, toggle off the B**rowse on Indexed Data** option. For S3 Views, both indexed and non-indexed browse are supported. By default, the **Browse on Indexed Data** option is disabled. **Note:** • Changing the snapshot after selecting the items (files or folders) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folders) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Once you select the files and folders to recover, select one of the following options: • **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task completes, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file. • **Next**: If you select this option, continue with the next step to configure the file recovery options. 6. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps:

If you select..	Then..
	1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. • For S3 Views, only indexed browsing is supported for archived snapshots; non-indexed browsing is not supported. • Recovery from archived snapshots of an S3 View is not supported. 7. Click **Select Recovery Point** and go to step 4.

4. Once you select the files and folders to recover, select one of the following options:
   - **Next Recover Option**: If you select this option, continue to step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task completes, from the **Recoveries** page, click the task name and then click **Download Files** to download the generated zip file.

     > **Note:** The download option is disabled for S3 Views.

Character replacement explanation for : and \ in file and directory names in zip files

If the original file or directory name contains a ‘:’ , it is replaced with ‘_COLON_’ in the archive.
Example:

- The original name is "a:b" then the archived name is "a_COLON_b"
- The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

If the original file or directory name contains a ‘\’, it is replaced with '_BACK_SLASH_' in the archive.
Example:

- The original name is “x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **Alternate view** to recover the files or folders to a different view.
6. From the **Select View** drop-down list, select an existing view. You can also create a new view by clicking **Create View** from the **Select View** drop-down.

> **Note:**
> You can select only those views as targets that have the:
>
> - **Case Sensitive File or Folder Names** option same as the source view. For example, if the file of a view you selected for recovery has the **Case Sensitive File or Folder Names** option disabled, then from the **Select** View drop-down, you can select only those views having the C**ase Sensitive File or Folder Names** option disabled.
> - **File DataLock** option disabled.
> - Same storage domain as the source view.

7. In the **Path** field, enter the path where you want to recover the file or folder in the target view. If the path you specified does not exist in the target, then Cohesity will create the path you specified and recover the file and folder in that path.
8. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is disabled to create the files and folders in the specified location. Any duplicate files are skipped. Enable this option to overwrite the existing files and folders.

For example, you are recovering a file, **1.txt**, from the source view V1 to the root directory of the target view V2. If a file, **1.txt**, already exists in the root directory of the target view V2, then the **Overwrite** behavior is as follows:

- If the **Overwrite** option is enabled during recovery, the **1.txt** in the target view is overwritten with the data in the **1.txt** file of the source view V1.
- If the **Overwrite** option is disabled during recovery, the **1.txt** in the target view V2 will not be overwritten with the data in the **1.txt** file of the source view V1.

If **Overwrite** is enabled, recovering a file to the source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

- **Preserve File/Folder Attributes**: This option is not applicable for the file or folder recovery of a view.
- **Continue on Error**: By default, this option is enabled to continue the recovery even if one of the objects encounters an error. Disable this option to discontinue the recovery task if one of the objects encounters an error.
- **Cluster Interface**: By default, the **Auto Select** option is enabled, and the recovery task automatically selects the correct VLAN. If you disable this option, select a configured interface group from the **Interface Group** drop-down.
- **Task Name**: Change the default name of the recovery task.
- **Save Success Files**: Enable this option to download the success logs for a recovery run.

9. Click **Recover**.

   The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Restore Previous Versions of SMB Files and Folders

You can restore previous versions of files and folders that are in Cohesity-protected SMB Views. This feature enables you to perform the restore from the system where the Views are mounted and provides the same functionality as the Windows "Restore previous versions" function.

The snapshot to be restored is generated only when the backup task is performed on the SMB view.

> **Note:**
> - The file or folder you want to restore must have already been backed up by a previous View Snapshot on the Cohesity cluster.
> - The versions available for restore are determined by the snapshots stored in the Cohesity cluster. For example, if the retention period specified in the Protection Policy retains snapshots for only two days, then only the versions from the past two days are available for restore.

To restore a previous version:

1. On the system, right-click the file or folder that is under a backed up view and select **Restore previous versions**.
2. Click the **Previous Versions** tab.
3. Select the version you want to restore to and click **Restore**.

Lock or Unlock Cohesity View Configuration

You can now lock or unlock the configuration of a Cohesity view. The lock or unlock feature provides additional security checks on the view's configuration. Only users with the **Data Security** privilege can lock or unlock a view.

You can lock or unlock a Cohesity view using the Cohesity CLI only. For more information, see view for view-related CLI operations:

- view lock
- view unlock

When a view is locked, no users can perform configuration-related operations such as:

- Access the properties of the view
- Create, edit, or delete aliases of the view
- Enable or disable user or directory quotas

Overwrite a View

You can use the Cohesity CLI `view overwrite` command to overwrite a View with the contents of another View in the same Cohesity cluster. The target View becomes a clone of the source View. This is useful if you want to maintain a View on a remote Cohesity cluster (used for replication) that has information from the latest replication of the source View. For example, you might manually create a View on the remote Cohesity cluster and then overwrite that View with a View snapshot after each replication task runs.

To overwrite a View:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Overwrite the target View with the source View contents by using the following Cohesity CLI command:

   ```
   admin>view overwrite source-view-name=<view1> target-view-name=<view2>
   ```

Delete a View

You can delete a View if you do not need the data in a View. When you delete a View permanently all data stored in the View such as files, Snapshots and Cloned VM files.

> **Note:** You cannot revert the View delete operation.

To delete a View, complete the following steps:

1. Navigate to **SmartFiles > Views**.
2. In the **All Views** page, for the required View click the ⋮ icon and select **Delete**.
3. In the Delete View screen, type **Yes** to confirm and click **Delete**.

Remote Views

When a View is replicated from a source cluster to the target/remote cluster, a read-only View is automatically created on the target/remote cluster, which is also known as the Remote View. Remote Views are always read-only. By default, Remote Views are visible to clients and visible in the Cohesity GUI.

To disable Remote Views visibility on clients for all protocols:

1. Navigate to **SmartFiles > Views** and click the **Global Settings** tab.
2. Turn off the **Visible to Clients** toggle.

To disable Remote Views visibility on the Cohesity cluster GUI:

1. Navigate to **SmartFiles > Views** and click the **Global Settings** tab.
2. Turn off the **Visible in Cohesity UI** toggle.

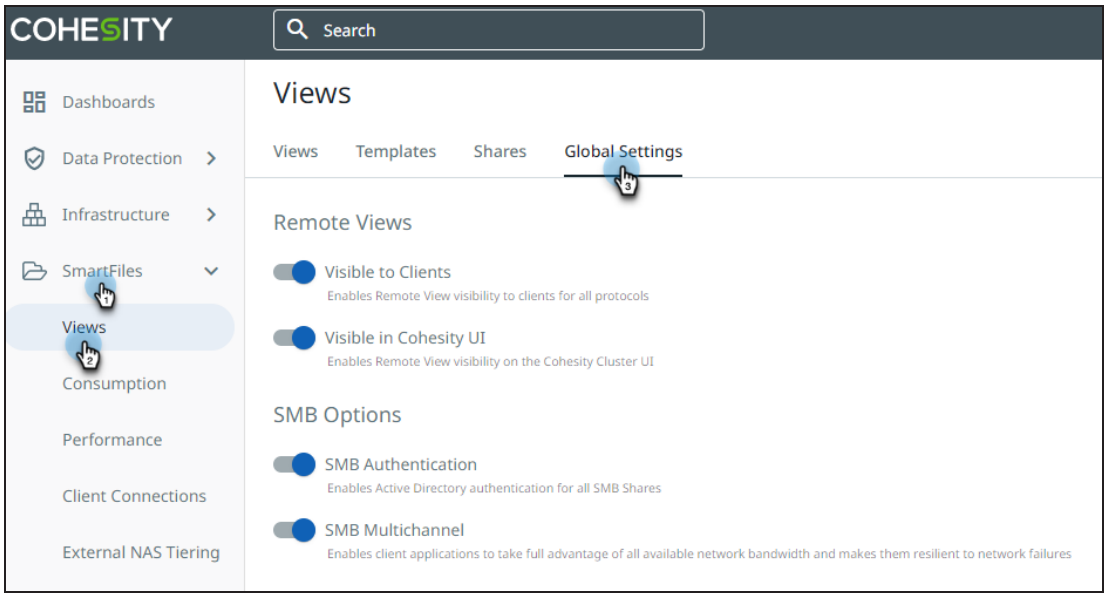

If the **Visible to Clients** option is disabled, Remote Views are not visible to clients through Windows Explorer or the Linux/Unix `showmount` command. However, clients can still access the Remote Views through the direct mount path.

If the **Visible in Cohesity UI** option is enabled, the Remote View is displayed on the Views page. The **Remote View** label is displayed next to the view name:

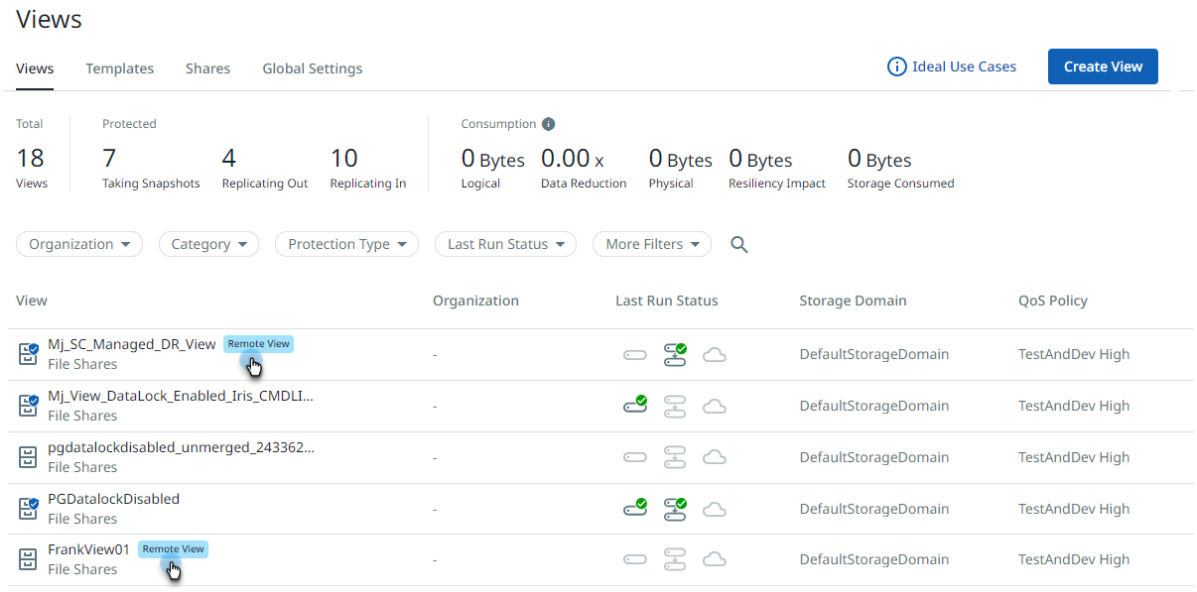

Consumption Details of a View

The view details page includes the **Consumption** tab, which displays the storage consumption and storage savings details for a specific Cohesity view.

To view the consumption details of a specific Cohesity view:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on a view name to navigate to the corresponding view details page.

   The **Consumption** tab is displayed by default:

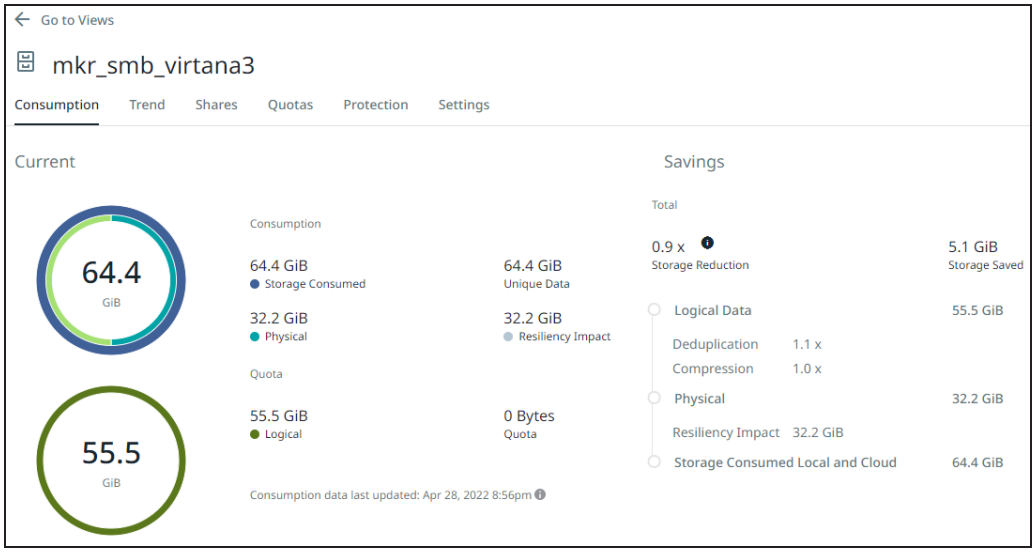

Related Topics

- Storage Statistics
- User and Directory Quotas for Views

Data Growth Trend of a View

The view details page includes the **Trend** tab, which shows the data growth trend. The data is presented in a chart. You can customize the time period by selecting **Last Day**, **Last Week**, or **Last Month**. Turn on or off the **Storage Saved** or **Resiliency Impact** toggle to add or remove data from the chart.

To view the data growth trend of a specific Cohesity view:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on a view name to navigate to the corresponding view details page.
3. Click the **Trend** tab and do one of the following:

- Select **Data** from the drop-down to view the data growth trend:

- Select **Files** from the drop-down to view the growth trend of directories and files:

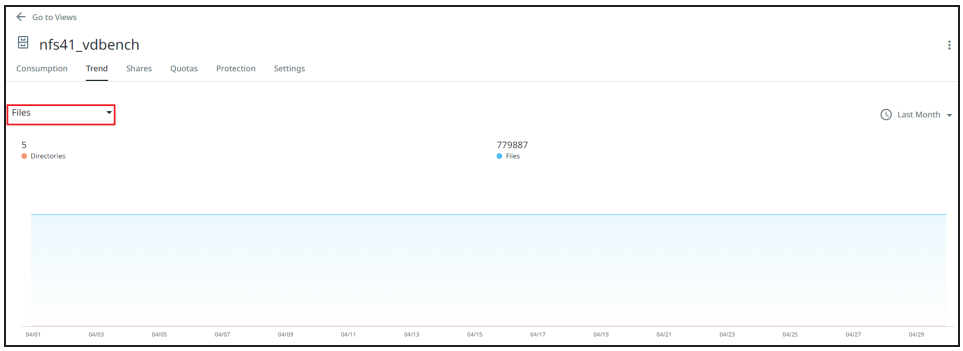

Protection Details of a View

The view details page includes the **Protection** tab, which displays the protection details of the view. The following details are displayed:

- Data Protection status
- Replication status
- Archival status
- Protection group and protection policy details
- List of snapshots and snapshot details

To view the protection details of a specific Cohesity view:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on a view name to navigate to the corresponding view details page.

3. Click the **Consumption** tab:

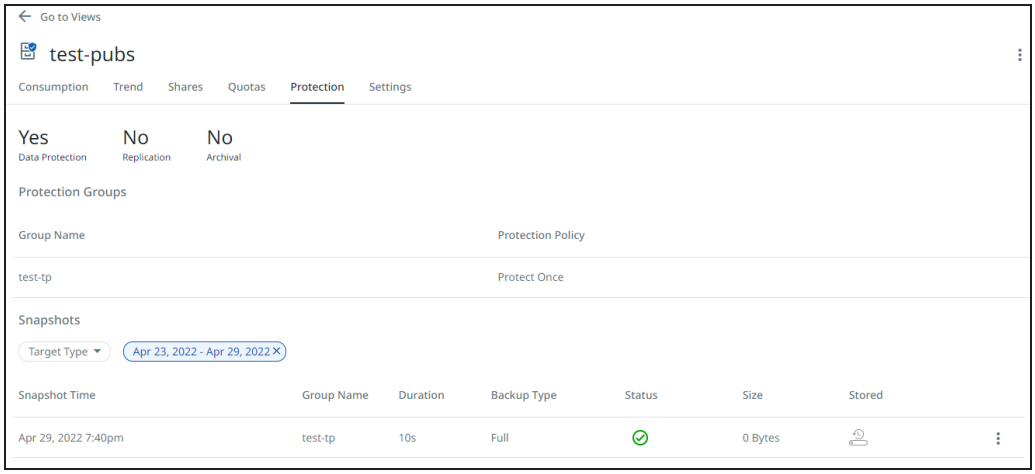

> **Note:** If the Cohesity view is not protected, a message is displayed. You can click **Create Protection Group** to protect the view. For more information, see Add or Edit a Protection Group for a View.

Cohesity View Details Page

The view details page includes the **Settings** tab, which displays all the details of the Cohesity view.

To view the details of a specific Cohesity view:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on a view name to navigate to the corresponding view details page.
3. Click the **Settings** tab.

   The **View Details** section is displayed. All the details specific to the view are displayed. You can also navigate to the **File Services Audit Logs** page by clicking the **Manage Audit Logs** option.

File Shares (SMB)

The following topics provide guidance about using and managing file shares:

SMB Shares

Mounting and accessing views as SMB shares requires the following high level steps.

1. Review the information about optimizing storage domains and views for your use case.

   See SmartFiles (File and Object Services) Overview.

2. Create storage domains and views for your use case.

   See Create or Edit Storage Domains and Create Views Using a Predefined Template.

3. Add the client systems that will mount and access the views to the allowlist.

   See Add Subnets to Allowlists.

4. Mount the views to client systems so users and applications can store data in them.

   See Mount a View Using SMB.

Once the above steps are completed, you can use a view to store data such as files and directories. Other backup products, such as Veeam, can also store data in a view, as shown in the following figure.

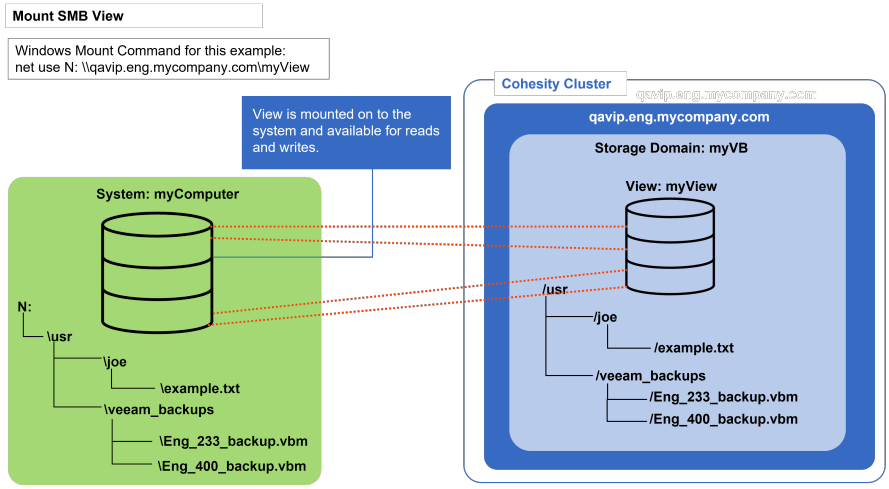

SMB Features

The following SMB features are supported on the Cohesity cluster:

- SMB Oplocks
- Continuous Availability
- Fast Durable Handles
- SMB Offline
- SMB Multichannel

SMB Oplocks

The SMB Oplocks feature provides the capability for an SMB client to lock files and cache content locally. By caching the content on the client side, it improves the performance of read and write operations on the file and reduces the amount of network traffic required.

Enabling SMB Oplocks on a View also activates the file and directory leasing feature on newer clients, such as Windows 2012 or later, or Windows Server 2012 R2 or later. This feature enhances performance by allowing the client to cache file and directory metadata.

The following table outlines the features utilized by the SMB client when the SMB Oplocks feature is enabled:

SMB Version	Clients	Feature Used
SMB 2.0	• Windows Vista or later • Windows Server 2012 • Mac OS X 10.9	The client uses the SMB2 oplocks feature
SMB 2.1	• Windows 12 or later • Windows Server 2012	The client uses the SMB2 file leasing feature
SMB 3.0 (or later)	• Windows 8 or later • Windows Server 2012 or later	The client uses the SMB2 file and directory leasing feature

By default, the SMB Oplocks feature is enabled for a View, allowing the client to lock files and cache content locally for enhanced performance. However, it's important to note that enabling SMB Oplocks may introduce data inconsistency, as the data is cached locally on the client. If you prefer to ensure that the client commits writes to the SMB server synchronously, without caching the content locally, you have the option to disable the SMB Oplocks feature. This can be done by editing the View settings and modifying the **SMB Oplocks** option accordingly. For more information, see SMB Options.

Continuous Availability

Continuous Availability (CA) is a collection of features in SMB 3.0 protocol that aims to facilitate smooth failover in a clustered environment. Cohesity does not support all CA features.

Cohesity supports the Persistent Handles feature. By default, Persistent Handles is enabled on the Cohesity View, and the cluster node stores file handle operations on the disk. This allows SMB clients to reconnect to a different node in case of a node failure and seamlessly continue their operations with the same file handle. The Persistent Handles feature primarily pertains to the SMB 3.0 protocol and this can have an impact on SMB 2 clients.

By leveraging the Persistent Handles feature, Cohesity provides enhanced availability and seamless failover capabilities within a clustered environment, ensuring continuous access to data even in the event of a node failure.

> **Note:** If you enable the **Fast Durable Handles** option on a View, the Continuous Availability feature gets disabled.

Fast Durable Handles

The Fast Durable Handles option is designed to enhance performance for SMB clients by improving I/O operations, metadata operations, file listing, and more. It offers improved resiliency against service restarts of the Cohesity data service and short network failures. However, it's important to note that Fast Durable Handles might not be able to handle more significant events such as node failures or crashes of the SMB server service.

> **Note:** By default, when creating a View, the Fast Durable Handles feature is disabled. Enabling this feature for a View results in the Continuous Availability feature being disabled for that View.

You can enable the Fast Durable Handles feature for applications that can tolerate I/O failures and would benefit from faster performance on the client side. This feature is particularly useful for applications like Microsoft Office App or scenarios involving bulk changes of permissions.

> **Note:** If you enable or disable the **Fast Durable Handles** option for a View, it is necessary to unmap and then remap the SMB share associated with that View.

SMB Offline

The SMB Offline feature enables SMB clients to access and work with files even when they are disconnected from the network. By default, this feature is disabled, which means no caching is performed. However, you have the option to enable it on a per-View basis during the View creation process, and you can also modify it later for existing Views.

When the **Offline File Caching** option is enabled during View creation, the manual caching option is also enabled. If the **Offline File Caching** option is not enabled, manual caching is not enabled, ensuring that no offline access or caching is performed for that View.

SMB Multichannel

The SMB multichannel feature, available with the SMB 3.0 protocol, is supported at the cluster level. It enables SMB clients to utilize multiple network connections concurrently for read and write operations between the client and server. This capability enhances the throughput and performance of SMB communication by leveraging multiple network paths.

By leveraging SMB multichannel, users can experience improved data transfer speeds and optimized network utilization, resulting in enhanced performance during file transfers and data access. For information on enabling or disabling SMB Multichannel, see Enable or Disable SMB Multichannel.

Prerequisites for using SMB Multichannel

- Windows 8 or a later version, Windows Server 2012, or Windows Server 2012 R2 is required on both server and client.
- Both the server and client must be running on the SMB 3.0 protocol.
- The Windows Server must have a minimum of two network interfaces or the Receive Side Scaling (RSS) feature must be enabled.
- One or more network adapters should be configured using NIC teaming

Limitation

The option to advertise multiple IP addresses on the cluster is not supported.

Client Authentication Workflow with SMB Encryption Enabled

The following is the workflow for clients to authenticate to Cohesity when SMB encryption is enabled:

1. SMB Client reaches the Domain Controller and self-identifies using credentials.
2. SMB Client asks for ticket to service (Cohesity).
3. Domain Controller checks if SMB Client with the above credentials has access to the requested service. If accessible, the Domain Controller sends the ticket to SMB Client which only the service provider (Cohesity) can decode.
4. SMB Client sends this ticket in the SMB SessionSetup request.
5. Cohesity decodes this ticket. If successful, Cohesity sends *Success* in response to SessionSetup.
6. SMB Client sends TreeConnect request to connect to the specific View.
    - If SMB encryption is forced on the SMB Client, TreeConnect and all the following SMB commands will be encrypted by the SMB Client. The SMB2 specification SMB Server replies with the encrypted message to the encrypted request.
    - If SMB encryption is set to be forced from Cohesity (on View settings) above TreeConnect request as well as TreeConnect response will not be signed but in the TreeConnect response Cohesity will notify the SMB Client that the following messages for this View must be encrypted. Hence, the SMB Client must start encryption of all messages following TreeConnect.

> **Note:** The above workflow is not specific to Cohesity, except forced encryption setting on View. Cohesity follows the standard Kerberos/SMB/Encryption pattern.

Considerations

- For SMB, the maximum path length is tailored to the SMB clients that connect to Cohesity since the client machine controls the session to the SMB share. For more details, see Maximum Path Length Limitation.

SMB Super User

An SMB Super User is a privilege for AD users and groups, and local SMB users and groups. A user with this privilege has complete access to view or share content overriding any assigned ACLs. However, the privilege cannot override the global allowlist and file-level data lock permissions.

The SMB Super User privilege is helpful when you have migrated data from a different source storage to the Cohesity cluster. The privilege allows a user to access and set appropriate permissions on the migrated data.

If the SMB identity of an NFS user is given the Super User permissions, then the NFS user gets full access to the view when the view security style is set to unified or NTFS.

You can assign the SMB Super User privilege to a user when creating a view or share. You can also edit a view or share and assign the privilege. The privilege is set at the share level for a view and not as NTFS root permissions. When the Super User privilege is assigned to a user for a specific view or share, the user has complete access to all the contents of the view or share.

> **Note:** You can assign this privilege only from the Cohesity Dashboard.

A user with the Super User privilege cannot access files that are quarantined by an antivirus feature enabled on a view.

The SMB Super User privilege is replicated or archived when you replicate or archive the view or share.

Limitation:

The SMB Super User privileges that are set for shares (view-aliases) will not be effective for the mapped NFS identities for unified and NTFS security style views.

Best Practices for SMB Shares

Consider the following best practices when using views as SMB shares.

- For SMB client 2.x (Windows 7 and Windows Vista), set the HKLM\System\CurrentControlSet\Services\LanmanWorkstation\Parameters\DisableLargeMtu registry key value to 0 (zero) to enable 1 MB writes. Leaving the registry value at 1 (one) limits the size to 64 KB.
- If the SQL Server service runs as a non-regular user account (system or service account), it fails to write to an SMB share. To allow these accounts to write to an SMB

share, give the SQL server's domain machine account at least "Modify" access to the directory's ACL.

- The client's IP addresses must be within the global allowlist or View-level allowlist and use the correct subnet mask for allowlist entry.
- Ensure the LDAP server is reachable.
- Transitive trusts should be enabled if a user is part of an AD domain that requires transitive trusts.
- If you are using AD, add the cluster to the AD domain before creating a View.
- To improve performance when copying a lot of small files from a Cohesity SMB share, enable prefetching on the cluster. To enable prefetching on the cluster, contact Cohesity Support.

Add Subnets to Allowlists

Before you can mount and access a view, ensure the IP address of the system that is mounting the view is specified in a subnet that has been added to an allowlist.

Cohesity supports the following types of allowlists:

- **Global Allowlist**—The global allowlist applies to all views on the Cohesity cluster, but is overridden by a view allowlist or view share allowlist, if they exist.
- **View Allowlist**—Only the systems with IP addresses in the subnets listed in the view's allowlist have privileges to access and mount this view.

Update the Global Allowlist

To add a subnet to the Cohesity cluster's global allowlist:

1. Navigate to **SmartFiles > Views**.
2. Click the **Global Settings** tab.
3. In the **Global Subnet Allowlist** section, click **Add**.
4. In the **Add Allowlist** dialog, enter a Subnet IP in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).

   > **Note:** CIDR prefix values are allowed from 0 to 32.

5. Select any or all permissions as required.

   > **Note:** By default, the ability to set **S3 Permissions** for restricting access to specific subnets is enabled.

6. Select **NFS Root Squash** as applicable.
7. In the **Description** field, enter a description for the allowlist.

8. Click **Add**:

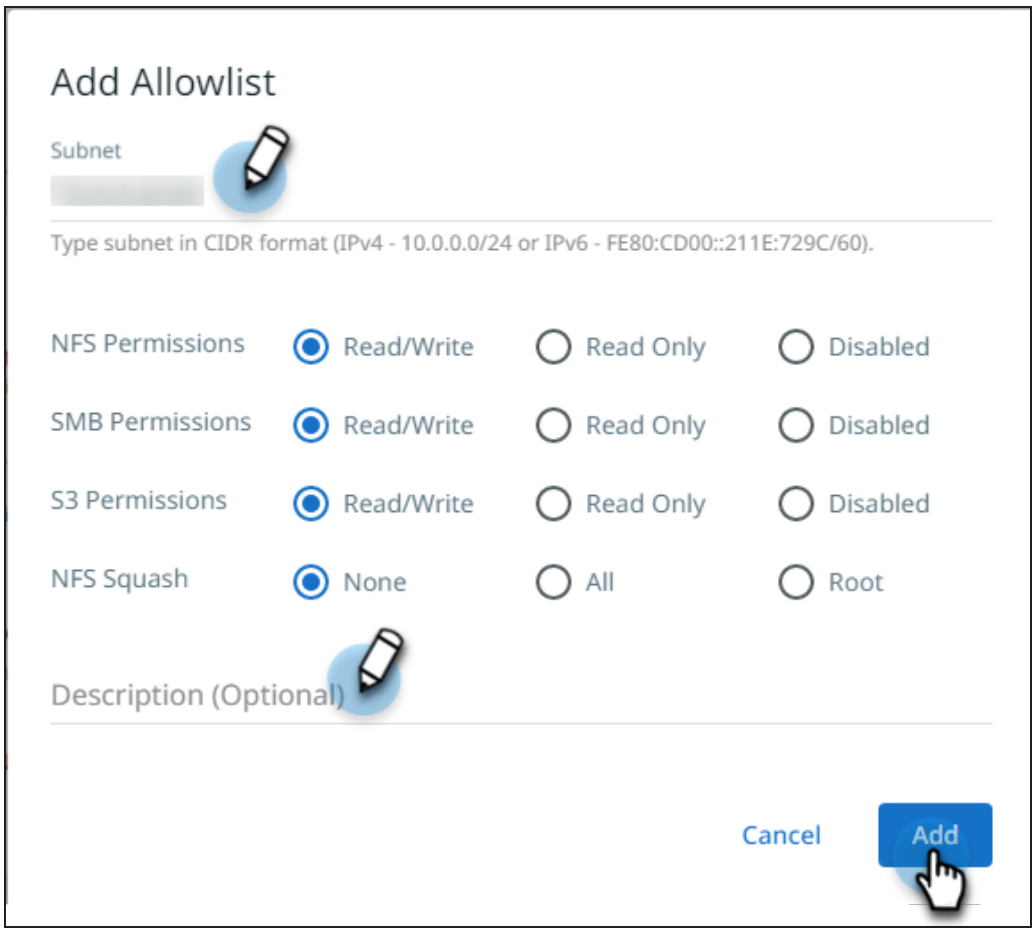

9. Repeat the preceding steps for specifying additional subnets in the global allowlist.

Update a View Allowlist

To add a subnet to an existing view's allowlist:

1. Navigate to **SmartFiles > Views**.
2. Locate the required View in the list, click the ⋮, and select **Edit** ( ).
3. In Advanced settings, click **Edit** for the Security category.
4. **Override Global IP Allowlist:** Turn on **Override Global IP Allowlist** if you want to specify a custom list for this view. Only the systems with IP addresses in the subnets listed in the view's allowlist have privileges to access and mount this view. If you enable the **Override Global IP Allowlist** toggle, the machines listed in the subnets of the global Cohesity cluster allowlist cannot access or mount this view unless the IP address of the machine is also listed in the View's Allowlist.

   Extend Global IP Allowlist: Allows you to extend the Global IP Allowlist.

   To add subnets to the allowlist of this view:

a. Under **IP Allowlist**, select: **Override Global IP Allowlist** or **Extend Global IP Allowlist**.
b. Click **Add**.
c. In the **Add Allowlist** screen, type subnet in CIDR format.
d. Select any or all permissions as required.
e. Select **NFS Squash** options.

The **NFS all squash** option is like NFS Root Squash. However, for given IP addresses, all users are mapped to the All Squash User/Group configured in the view. The default All Squash User for all views is "nobody".

> **Note:** If the **Override Global Allowlist** toggle is off, this view can only be accessed and mounted by systems specified by the global Cohesity cluster allowlist.

5. Click **Update View**.

For more information on global settings for each protocol, see the following topics:

- Global Settings for NFS
- Global Settings for SMB

Global Settings for SMB

In addition to the settings provided in Add Subnets to Allowlists, the following settings are specific to SMB.

- Update a Share Allowlist
- Disable SMB Authentication
- Enable or Disable SMB Multichannel

Cohesity supports the following types of allowlists:

Share Allowlist—If a view has a share allowlist, only the systems with IP addresses in the subnets listed in the share allowlist have privileges to mount and access the folders and subfolders in the share.

You can add a share allowlist only to views that support the SMB protocol. The share allowlist is only in effect during SMB access. If the view supports both the SMB and other protocols, then access to the other protocols is subject to the view allowlist, if present, or the global allowlist.

> **Note:** When adding a single host IP address to an allowlist, the netmask must be 255.255.255.255.

After the Cohesity cluster is upgraded to 6.6.0c, note the following:

Existing NFS or SMB view—IP restriction is retained. If you access an existing NFS or SMB view using the S3 protocol, the request fails if the existing allowlist does not include the IP address of S3 clients. Before or after the upgrade, ensure to add the IP address or subnet at the view level if the NFS or SMB view is accessed using the S3 protocol.

Update a Share Allowlist

You can add an SMB share allowlist when creating or editing a share. For more information, see Manage Share Aliases.

Disable SMB Authentication

You can disable Active Directory authentication for SMB shares if you want users' AD credentials to be used only for logging on to the Cohesity cluster.

1. Navigate to **SmartFiles > Views**.
2. Click the **Global Settings** tab.
3. Turn off the **SMB Authentication** toggle.
4. In the **Disable SMB Authentication** screen, click **Disable**.

A message appears at the bottom of the screen confirming the change.

Enable or Disable SMB Multichannel

The SMB Multichannel feature is enabled by default and you may not need to disable it for any performance or networking problem. If you disable the feature, the performance and the reliability of the SMB3 protocol might reduce.

For more information about SMB Multichannel, see SMB Multichannel.

Prerequisites for using SMB multichannel:

- Windows 8 or later, Windows Server 2012 R2, Windows Server 2012 or higher on both server and client.
- Server and client must be on SMB 3.0.
- The Windows Server must have a minimum of two network interfaces or Receive side scaling (RSS) feature is enabled.
- One or more network adapters are configured using NIC teaming.
- Assign additional VIPs to each node for improved performance of SMB multichannel.

To disable or enable the SMB Multichannel feature:

1. Navigate to **File Services**> **Views**.
2. Click the **Global Settings** tab.
3. Under the **SMB Options** section, complete the following steps:
    1. Turn off the **SMB Multichannel** toggle.
    2. In the **Disable SMB Multichannel** screen, click **Disable**.

> **Note:** Turn on the **SMB Multichannel** toggle to enable SMB Multichannel. The confirmation message is displayed at the bottom of the screen.

Mount a View Using SMB

Before you can store data on a view, you have to mount it as an SMB share on your Windows system.

Prerequisites

- The view to be mounted must be enabled for the SMB protocol.
- The system that will mount the view must be in a subnet in a Cohesity allowlist. Otherwise, the system will not have privileges to mount and access the view.

Considerations

- Keeping with the industry standard of change notification for SMB shares, recursive change notifications are not sent due to their effect on process load and network traffic.
- Filenames that contain UTF-16 character codes ranging from U+D800 to U+DFFF are not allowed in Cohesity SMB shares.
- For Linux clients that are members of AD, using "client max protocol = SMB2" in the [global] section of /etc/samba/smb.conf is not supported. Use "client max protocol = SMB3".
- Cohesity SMB shares do not support alternate data streams.
- You can add Cohesity SMB shares as a Microsoft Distributed File System (DFS) target, but note that SmartFiles does not support any additional features or functionalities provided by Microsoft DFS.
- Windows behavior prevents Cohesity SMB shares from being automatically discoverable. Use the `net view` command to probe the cluster explicitly using \\<Cluster-machine-account-name> or \\<Cluster-vip-FQDN> or \\<Cluster-VIP>.

Mount a View

1. In the Cohesity Dashboard, copy the view mount path:
    a. Select **SmartFiles > Views**.
    b. Browse for the view in the list.
    c. Click the **Mount Paths** icon of SMB.
    d. In the **Mount Paths** screen, click the copy icon for SMB mount path and click **Close**.

2. On the Windows machines, mount the view as an SMB share using the mount path copied previously. You can use OS commands or utilities.

   For example, you can use the Windows command prompt to run the following command:

   ```
   net use <Drive Letter>: \\<View Mount Path>
   ```

   Example:

   ```
   net use N: \\qavip.eng.mycompany.com\myView
   ```

   If using Windows PowerShell, you do not have to mount the view. You can change the directory directly to a UNC path as long as your Windows login has privileges for the remote share. For example, you can do the following:

   ```
   cd \\<View Mount Path>
   ```

   The mounted file system is available to store data.

   If you want to use a Cohesity SMB share as a target for backup applications such as Veeam or CommVault, ensure you provide the UNC path to the target, for example:

   ```
   \\Hostname_or_VIP\myView
   ```

3. **Important:** At this point, the Cohesity root share does not have any default permissions or Access Control Entries (ACEs). The Windows administrator must apply these permissions before the share is ready for use.
   a. On the Windows machines, right-click the share root and go to **Properties > Security**.
   b. ACEs are applied only if the Administrator specifically adds ACEs to the share. A warning is displayed asking the logged-on user to assign ACEs to the share root.

Share Based Backup Using Remote VSS

Cohesity supports shadow copy creation for SMB views through Remote VSS (Volume Shadow Copy Service). With Remote VSS, you can take the application consistent shadow copies for applications that store data and configuration files on shares.

A shadow copy of the data represents a copy of the data in the share taken at a particular instant in time. A set of one or more of these shadow copies is collectively called as a Shadow copy set. Each shadow copy and shadow copy set is represented with a unique GUID.

Create Backup Using Remote VSS

The generic steps to take a backup using Remote VSS are as follows:

1. Prepare for Shadow Copy.
    - The VSS client running on the application first collects information about what shares must be included while taking a shadow copy set. A shadow copy refers to one share location.
2. Create Shadow Copy.
    - Cohesity creates a clone of every view share in the shadow copy set.
    - Once the shadow copy is created, the clone of the shadow copy set is exposed.
3. Recovery of Shadow Copy set.
    - During the creation of shadow copy, if there are any modifications being done to the files that are being cloned, the VSS application must update these copies so that the shadow copy set is in a consistent state.
    - If auto recovery is enabled, then read/write permissions to that share must be set.
    - After recovery, all created clones are updated with read-only permissions.
4. Optional: Delete Shadow Copy.
    - The shadow copy can be deleted in the future if no more is required.

List Active SMB Connections and File Opens

You can use the Cohesity CLI to list active SMB connections and file opens for views.

List Active SMB Connections

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. List the active SMB connections for Views by using the Cohesity CLI `view list-smb-connections` command:

   ```
   admin>view list-smb-connections
   ```

   > **Tip:** You can also use the CLI Help feature to view the `view list-smb-connections` options by entering the following command: `view list-smb-connections help`.

CLI Option	Example	Description
`names=`	`ViewName1,ViewName2`	Optional. Specify a comma separated list of View names. The active SMB connections for these Views are displayed.

List Active SMB File Opens

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. List the active SMB file opens by using the Cohesity CLI `view list-smb-file-opens` command:

   ```
   admin>view list-smb-file-opens
   ```

   > **Tip:** You can also use the CLI Help feature to view the `view list-smb-file-opens` options by entering the following command: `view list-smb-file-opens help`.

CLI Option	Example	Description
`cookie=`	`string`	Optional. Copy and paste the cookie returned from a previous command's output to get the next set of entries in the output. A cookie is included at the end of the output if there are additional entries to display. For example, if you specify `page-count=10` and there are still more entries to display, include the cookie in the command and run it again to display the next set of entries. If this option is not specified, the command always outputs the first set of entries.
`file-path=`	`FilePath`	Optional. Specify the file path to an SMB file relative to the SMB mount path. The active SMB file opens for this file are displayed. If this option is not specified, the command outputs all file paths that have file opens.
`name=`	`ViewName`	Optional. Specify the name of the View. The active SMB file opens for this View are displayed. If this option is not specified, the command outputs all open file paths in all Views in the Cohesity cluster.

CLI Option	Example	Description
`page-count=`	`10`	Optional. Specify the number of entries you want the output to display. This limits the number of entries in the command output. If this option is not specified, the command outputs a maximum of 1000 entries or the number of file opens, whichever is smaller.

Supported Windows Well-Known SIDs

Starting with Cohesity 6.5.1, Cohesity added support for the following Windows well-known security identifier (SIDs) groups. For more information on SIDs, Well-known-security-identifiers-in-windows-operating-systems.

SID	Name
S-1-5-32-544	Administrators
S-1-5-32-545	Users
S-1-5-32-546	Guests
S-1-5-32-547	Power Users
S-1-5-32-548	Account Operators
S-1-5-32-549	Server Operators
S-1-5-32-550	Print Operators
S-1-5-32-551	Backup Operators
S-1-5-32-552	Replicators
S-1-5-32-582	Storage Replica Administrators

> **Note:** These supported SID groups are Windows SMB-specific groups and these are different from Cohesity inbuilt roles.

A Windows Administrator can assign any of the supported SID groups to a Cohesity share mounted on Windows Server or use the Microsoft Management Console (MMC) console to grant or deny access for the assigned SID group.

Once the supported SID group is assigned to the Cohesity share by the Windows Administrator, you can see the SID group name in the View properties under the NTFS root permissions of the Cohesity Dashboard.

For example, if you grant Read/Write permissions for the Backup Operators group to a View, the View Settings in the Cohesity Dashboard displays the group name (**File Services**>**Views**>(click the View name)**Settings**> **SMB Options**> **NTFS Root permissions**).

Limitation

You cannot assign these SID group roles to a View from the Cohesity Dashboard.

Microsoft Management Console (MMC) Integration

You can use the Shared Folders Microsoft Management Console (MMC) snap-in on Microsoft Windows clients to perform basic configuration and management actions for SMB shares on a Cohesity cluster.

Set up a Windows Client to Manage SMB Shares on Cohesity

Before you can manage Cohesity SMB shares, you must map a drive on the Windows client where MMC will be used to a Cohesity share.

When mapping, ensure you do the following:

- Add the Cohesity cluster to an Active Directory domain.
- Use the FQDN of the Cohesity cluster. For example: \\MyCluster.cohesity.com\

  Where, *MyCluster* is the machine account used to add the Cohesity cluster to the Active Directory domain. For more information, see Active Directory.

  The VIP of the Cohesity cluster is also supported but using the FQDN is recommended.
- The user logged into the Windows client must be a domain administrator or must have the Full Control permission on the share.
- For a View, the **Shares are Browsable** option must be toggled on for the shares to be discoverable by the Windows `net view` command. For more information, see Create a View.

Supported Actions Using MMC

- List Shares
- Get/set Share Permissions
- List Sessions
- List Active Open Files
- Close Active Open Files
- Create a Share
- Stop Sharing a View Alias

List Shares

1. Open the Shared Folders MMC snap-in:
   a. Click **Start > Run**.
   b. Type fsmgmt.msc and click **OK**.

      Shared Folders MMC snap-in opens.
2. Connect to the Cohesity cluster that has the SMB shares:
   a. Click **Action > Connect to another computer**.
   b. Type the FQDN of the Cohesity cluster. For example: MyCluster.cohesity.com
   c. Click **OK**.
3. In the left pane, click **Shares**. The right pane lists all SMB shares.

Get/set Share Permissions

1. Open the Shared Folders MMC snap-in:
   a. Click **Start > Run**.
   b. Type fsmgmt.msc and click **OK**.

      Shared Folders MMC snap-in opens.
2. Connect to the Cohesity cluster that has the SMB shares:
   a. Click **Action > Connect to another computer**.
   b. Type the FQDN of the Cohesity cluster. For example: MyCluster.cohesity.com
   c. Click **OK**.
3. In the left pane, click **Shares**. The right pane lists all SMB shares.
4. In the right pane, right-click the desired share and then click **Properties**.
5. To get the permissions a user or group has for the share:

   On the Share Permissions tab, click the desired user name or the group name. The permissions are displayed in the 'Permissions for' section.
6. To set the permissions a user or a group has for the share:

   On the Security tab, click the desired user name or the group name and then click **Edit**.

   > **Note:** Changes affect the SMB share only, not the ACL in the file system of the exported directory.

List Sessions

1. Open the Shared Folders MMC snap-in:
   a. Click **Start > Run**.
   b. Type fsmgmt.msc and click **OK**.
      Shared Folders MMC snap-in opens.
2. Connect to the Cohesity cluster that has the SMB shares:
   a. Click **Action > Connect to another computer**.
   b. Type the FQDN of the Cohesity cluster. For example: MyCluster.cohesity.com
   c. Click **OK**.
3. In the left pane, click **Sessions**. The right pane lists all active sessions to SMB shares.

List Active Open Files

1. Open the Shared Folders MMC snap-in:
   a. Click **Start > Run**.
   b. Type fsmgmt.msc and click **OK**.
      Shared Folders MMC snap-in opens.
2. Connect to the Cohesity cluster that has the SMB shares:
   a. Click **Action > Connect to another computer**.
   b. Type the FQDN of the Cohesity cluster. For example: MyCluster.cohesity.com
   c. Click **OK**.
3. In the left pane, click **Open Files**. The right pane lists all open files.

Close Active Open Files

1. Open the Shared Folders MMC snap-in:
   a. Click **Start > Run**.
   b. Type fsmgmt.msc and click **OK**.
      Shared Folders MMC snap-in opens.
2. Connect to the Cohesity cluster that has the SMB shares:
   a. Click **Action > Connect to another computer**.
   b. Type the FQDN of the Cohesity cluster. For example: MyCluster.cohesity.com
   c. Click **OK**.
3. In the left pane, click **Open Files**. The right pane lists all open files.
4. In the right pane, right-click the desired file and then click **Close Open File**.

Create a Share

You can create a View or a View alias.

1. Open the Shared Folders MMC snap-in:
   a. Click **Start > Run**.
   b. Type fsmgmt.msc and click **OK**.

      Shared Folders MMC snap-in opens.
2. Connect to the Cohesity cluster that has the SMB shares:
   a. Click **Action > Connect to another computer**.
   b. Type the FQDN of the Cohesity cluster. For example: MyCluster.cohesity.com
   c. Click **OK**.
3. Right-click **Shares** and select **New Share**.

   When the Shared Folder wizard opens, click **Next**.
4. **Folder path:** Enter a Storage Domain name and/or a View name or a folder path in a View. These must already exist on the Cohesity cluster.

   To create a View in the Storage Domain on the Cluster, for example, enter `c:\<storage-domain-name>`

   To create a View alias on the Cluster, for example, enter `c:\<storage-domain-name>\<view-name>\<view-path>`

   > **Note:** Using the Browse option to create a share is not supported.
5. **Share name:** Enter the `<share-name>`, which is either the View name or View alias name that will be created. This is based on the Folder path.
6. Select permissions and click **Finish**.

Stop Sharing a View Alias

> **Note:** Stopping sharing for a View in the Cluster is not supported.

1. Open the Shared Folders MMC snap-in:
   a. Click **Start > Run**.
   b. Type fsmgmt.msc and click **OK**.

      Shared Folders MMC snap-in opens.
2. Connect to the Cohesity cluster that has the SMB shares:
   a. Click **Action > Connect to another computer**.
   b. Type the FQDN of the Cohesity cluster. For example: MyCluster.cohesity.com
   c. Click **OK**.

3. Click **Shares**.
4. Right-click the View alias you want to stop sharing and click **Stop Sharing**.

Manage Share Aliases

A share is used to mount a View on a user's system. Cohesity automatically creates a default share name for each View in the cluster, using the View name as the share name. If a View is renamed, the default share name reflects the new name. The default share name provides access to the root View. To see a list of shares in the Cohesity cluster, navigate to **SmartFiles > Views** and click the **Shares** tab.

You can manually create a share alias to an existing folder in a View. Users can then use the alias to access the folder without providing (or even knowing) the full path to it. This allows you to expose only that folder and its files and sub-folders. The user does not see the parent folders.

Using an alias can shorten the path used in the mount command. Also, unlike the default share name, the alias does not change if the View is renamed.

Consider a View that contains nested folders as shown below:

```
viewName
       dirOne
               dirTwo
                       dirThree
```

For example, you can assign a share name of `Three` to the `\dirOne\dirTwo\dirThree` path and then use the share name as follows:

- Windows: `net use V: \\vip.eng.mycompany.com\Three`
- Linux: `mount -t nfs vip.eng.mycompany.com:/Three mnt`

You can create multiple shares for different folders in Views. Share names must be unique within the Cohesity cluster.

> **Note:** If a share name refers to a folder that gets deleted, the share will be discoverable but you cannot mount it until you recreate the folder.

Create a Share

1. Navigate to **SmartFiles > Views**.
2. In the **Views** page, click on a view name to navigate to the corresponding view details page.
3. Click the **Shares** tab.
4. Click **Create Share**.

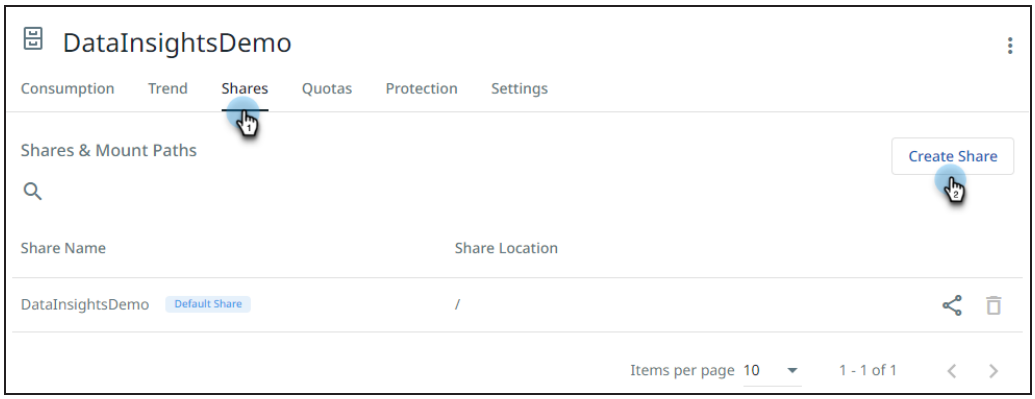

5. In the **Create Share** page, the root/folder and any sub-folders in the View are displayed. Click a folder and any sub-folder on which you want to create a share.
6. Provide the following details:
   - **Share Name**—Displays a modifiable name, based on the final folder in the path. Change the name as needed. The name must be unique within the Cohesity cluster.
   - **Auditing**—If auditing is enabled on the View, share-level auditing is automatically enabled.

     By default, **Inherit** is selected.

     For user-created shares, you can override the View-level setting. You can enable or disable auditing on a share. If you enable auditing on a share, the share is audited even if auditing is disabled on the View. If you disable auditing on a share, the share is not audited even if auditing is enabled on the View:

     > **Note:** You cannot modify the auditing setting of the default share from the share details page. To modify the auditing setting of the default share, edit the View and update the setting. For more information, see Enable or Disable Audit Logs for Views.

   - **Share Level Permissions**—For an SMB share, you can optionally define custom permissions for the share. The share-level permissions are stored on the cluster as view metadata and are independent of the NTFS permissions.

     Search for and select the principal (user or group) that you want to define permissions for. Select **Allow** or **Deny**. Click **Add**. Permissions adhere to the prescribed behavior of ACEs in a DACL.
   - **Super Users**—An SMB Super User is a privilege for AD users and groups, and local SMB users and groups. Select the AD domain and principal (user or group). Click + to add additional super users.

A user with this privilege has complete access to view or share content overriding any assigned ACLs. However, the privilege cannot override the global allowlist and file-level data lock permissions. The privilege is set at the share level for a View and not as NTFS root permissions. When the Super User privilege is assigned to a user for a specific View or share, the user has complete access to all the contents of the View or share. For more information, see SMB Super User.

> **Note:** The most restrictive permissions are applied. The permissions can be **Full Control**, **Modify**, **Read Only**, or **Super User**. For example, if NTFS permissions are Read and Write, but the share permission is Read, the Read permission is applied.

- **Override Global Allowlist**—Available only if the View supports the SMB protocol. Turn on the toggle if you want to create an allowlist for the share and override the Cohesity cluster's global allowlist.

  Click **Add Allowlist** and enter a subnet IP and subnet mask to specify a range of systems that have access to the share. Select the SMB permissions and enter an optional description. Click **Add**.

7. Click **Create Share**.

Edit a Share

You can change the global allowlist overrides for SMB views and share permissions for a share.

1. Navigate to **SmartFiles > Views** and click the **Shares** tab.
2. Click the share name you want to change.
3. Modify the share-level permissions and/or global allowlist overrides.

Delete a Share

You can delete a user-created share. You cannot delete a Cohesity-created share (the **Folder** column on the **Shares** tab contains only a forward slash / for these shares).

1. Navigate to **SmartFiles > Views** and click the **Shares** tab.
2. In the **Shares & Mount Paths** section, hover over the share that you want to delete and click the delete icon.
3. Confirm the deletion.

Mount Paths & API Endpoints

To view mount paths and API endpoints:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, click on a view name to navigate to the corresponding view details page.
3. Click the **Shares** tab.
4. In the **Shares & Mount Paths** section, click corresponding to the share name to view the **Mount Paths & API Endpoints** dialog:

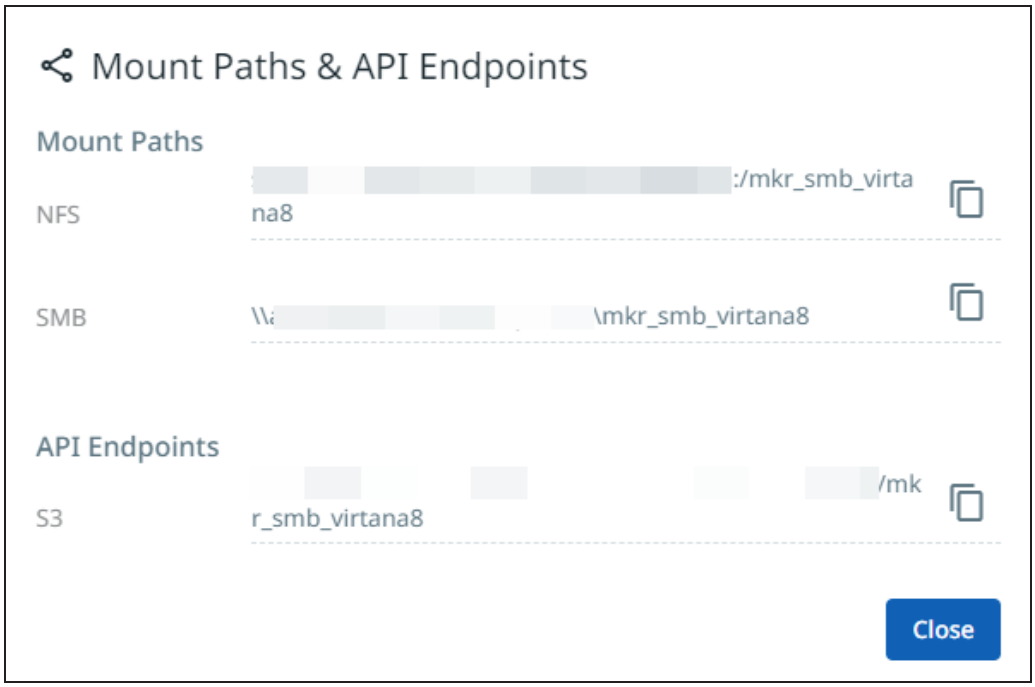

Troubleshooting File Share Issues

Use the following information to troubleshoot file share issues.

Error: Could not Mount Datastore on any Host

Symptom: Attempting to mount a view results in the error:

```
Could not mount datastore on any host [additional details: NFS mount
<FQDN>:<NFS Path> failed: The mount request was denied by the NFS server.
Check that the export exists and that the client is permitted to mount it.
```

Cause: A configuration error in the global or view allowlist.

Resolution: Perform the following actions:

- Ensure that the subnet of the host is included as a part of the subnets in allowlist for the view.

- Disable the **Access Based Enumeration** option for the view. If Access Based Enumeration is enabled, then when a user browses an SMB share, only the files and folders that a user has privilege to access are displayed.
- After configuring the global allowlist, ensure that the hostname of the Cohesity cluster is successfully resolved. DNS resolution errors will be reported as:

  ```
  Could not mount datastore on any host [additional details: Operation
  failed, diagnostics report: Unable to resolve hostname '<FQDN>']
  ```

  To resolve this issue in a VMware environment, add the DNS server that can resolve the FQDN using the following command:

  ```
  esxcli network ip dns server add  --server=<DNS Server IP>
  ```

Error: Skipped Backing Up {folder name}

Symptom: Protecting a physical server via SMB shares to a Cohesity view fails with the error:

```
Skipped backing up {folder name} due to: Failed to create a new connection
to 10.1.0.33/G$, error msg: smbc_wrapper_cli_full_connection failed: NT_
STATUS_IO_TIMEOUT
```

Cause: This error is expected if the option **Skip Files on Errors** is enabled in the Protection Group. The option is enabled by default and lets the Protection Group continue to run even if it encounters errors on files, such as permissions errors. The Protection Group attempts to protect the files on the next protection run.

Resolution: Disable the **Skip Files on Errors** option (note that any error on a file will cause the Protection Group to fail.) For more information, see Protect NAS.

Error: [kInvalidRequest]: Diff streamer could not visit any entity successfully

Symptom: Any of the following:

- After creating a NAS Protection Group to protect CIFS/SMB shares, the Protection Group fails with the error:

  ```
  [kInvalidRequest]: Diff streamer could not visit any entity
  successfully
  ```

- When registering a NAS source, clicking the **Register** button after entering the SMB user credentials may fail with the error:

  ```
  HTTP POST - Authorization failed
  ```

Cause: Typically an authentication failure. Specific reasons include:

- Credentials recently changed.
- Credentials were entered with an invalid syntax. For example, the Cohesity Dashboard expects domain\user syntax, but accepts user@domain. The format

user@domain results in several different issues when attempting to back up and must be avoided.

- Management traffic traverses a 1G interface, and the domain controller is not reachable from this interface.
- Domain account is being used to register the source for a Cohesity cluster which has not been registered with a domain.
- A NAS Protection Group is running on Windows 2003, which supports SMBv1. A mismatch of SMB versions may be occurring.
- The Active Directory server configuration may need to be changed. For example, trusted domains may need to be set to 0. Refer to Cohesity knowledge base article Slow or no Cohesity Dashboard login when authenticating with Active Directory domains.

Resolution: Perform the following actions:

Update the credentials:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Find the source in the list, click the Actions Menu ⋮ and select **Edit**.
3. Select the Edit icon next to the Protection Group.
4. Re-enter the username and password.
5. Click **Save**.

If a log message similar to the following is reported, investigate an SMB version mismatch.

```
E0514 00:34:59.163116  8201 file_backup_base_op.cc:3938] backup_1296951_
1/task_1296954Skipping backup of entity  due to error: [kInvalidRequest]:
Failed to get metadata for path , error msg: Failed to open: NT_STATUS_
INVALID_PARAMETER
```

Use the following command syntax to mount an SMB share manually to test the connection. Adjust the value in the `vers` field, as required.

```
/usr/bin/sudo USER=<username> /usr/bin/timeout --signal=SIGINT --kill-after-
r=1s 10s /usr/bin/mount -t cifs -o vers=1.0,soft,intr,user_xat-
tr,domain=<domain>,ro <share_path> <mount_path>
```

Error: nfs: server {server name} not responding, still trying

Symptom: The following error occurs when you connect to the Cohesity cluster with the NFS Client:

```
nfs: server {server name} not responding, still trying
```

Cause: The timeout value between the NFS Client and the Cohesity cluster is set too low.

Resolution: Increase the timeout value in the NFS client using the `timeo` mount option. For more information about NFS mount options, see the NFS Linux man page.

Some files are not displayed in the view

Symptom: The Windows NFS client does not display all the files in a view.

Cause: The view supports only the SMB protocol.

Resolution: Perform the following actions:

1. In the Cohesity Dashboard, select **SmartFiles > Views** .
2. Find the view in the list, click the Actions Menu ⋮ and select **Edit**.
3. Click **Show Advanced Settings** and select **All** from **View Protocol**.
4. Enable **Shares are Browsable**. If disabled, the shares will be hidden.
5. Disable **Access Based Enumeration for SMB**. If the Access Based Enumeration is enabled, then when a user browses an SMB share, only the files and folders that a user has privileges to access are displayed.
6. Click **Update View** to save the changes.

Error: Path not found

Symptom: The error `path not found` occurs when attempting to browse an SMB share for a view. Recovery tasks also fail.

Cause: When a Cohesity cluster is joined to an Active Directory, Cohesity SMB views use the computer name. This name is different from the Cohesity cluster name and what is configured in DNS. This difference will cause a "`path not found`" error and cause recovery tasks to fail.

For example, in the following figure, the machine account name for the Cohesity cluster is Cohesity (in the Cohesity Dashboard, select **Settings > Access Management** and click the **Active Directory** tab).

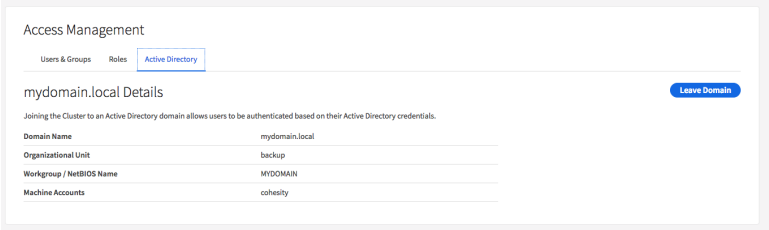

Looking at the Active Directory, the computer account will be created with the name Cohesity.

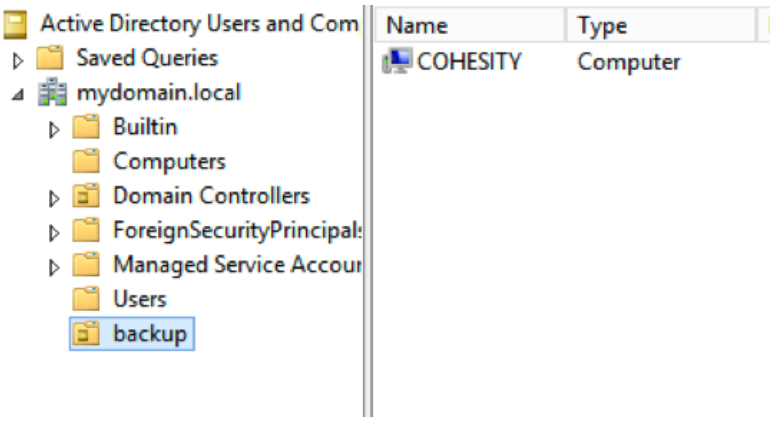

Since a DNS entry does not exist, the following network error will be returned. You can verify it by browsing `IP\\server name IP\share name`.

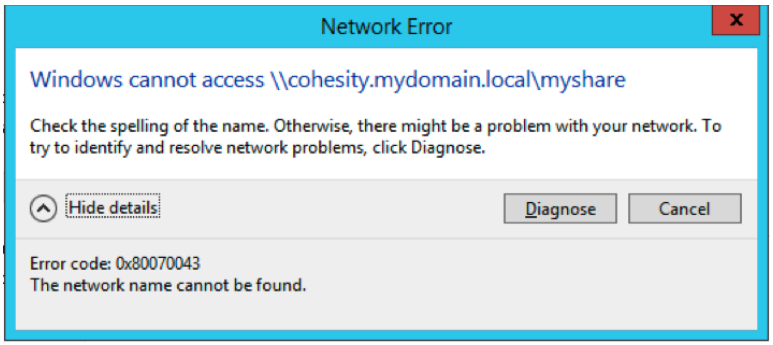

Attempting to ping Cohesity cluster or use `nslookup` on Cohesity cluster will fail.

```
C:\Users\Administrator&gt;nslookup cohesity.mydomain.local
Server:  UnKnown
Address:  ::1
```

Resolution: Perform one of the following actions:

- Create a DNS CNAME for resolving the machine account. From Windows DNS, right-click the new record, and then select the **CNAME** option.

  For example, in the following figure, the Cohesity cluster is **jj-backup** and has a DNS record. If no DNS record exists for the cluster's IP address, create a new A record and use the IP address of the cluster.

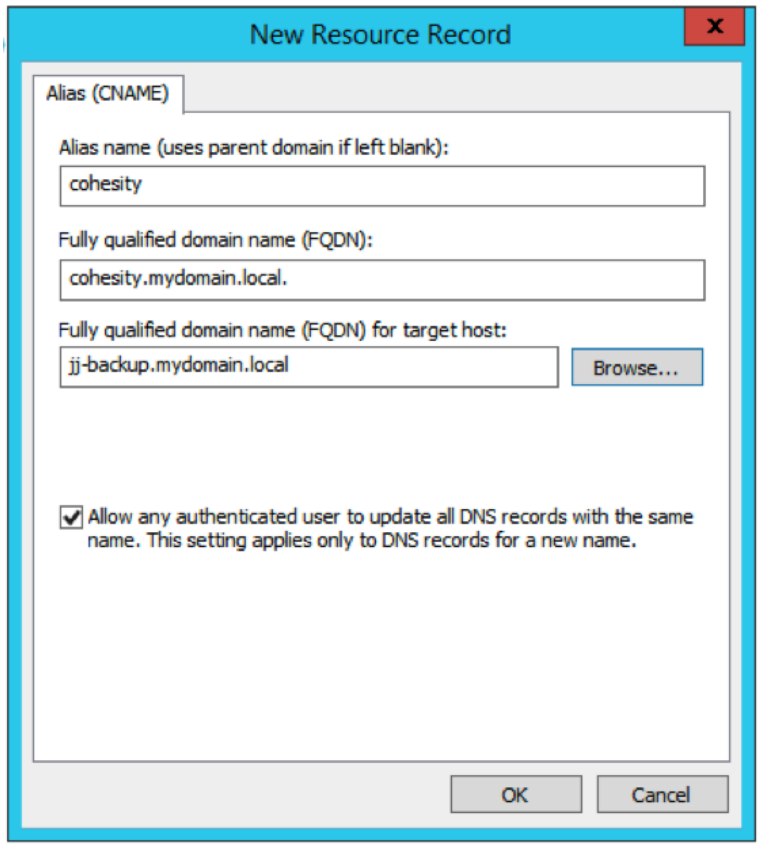

- Leave the domain AD (in the Cohesity Dashboard, select **Settings > Access Management > Active Directory** and click **Leave Domain**) and rejoin using a machine account that resolves in DNS.

  After the domain join completes, verify the Cohesity cluster with the `ping` and `nslookup` commands, and then test browsing a view via SMB.

Error: mount error(13): Permission denied

Symptom: Accessing an SMB share results in the error `mount error(13): Permission denied`, for example:

```
mount //app99/userdata to /home_cohesity_data/magneto/tmp/nas_mnt/source_
cohesity-3313822519568060-585114 failed: mount error(13): Permission denied
Refer to the mount.cifs(8) manual page (e.g. man mount.cifs)
```

Cause: The user credentials for the account accessing the SMB share may have changed.

Resolution: Update the Cohesity cluster with the new credentials.

Error: Windows cannot access {filename}

Issue: The "Enter network credentials" prompt is displayed on every mount or access attempt and the following error message is displayed:

```
Windows cannot access {filename}
You do not have permissions to access {filename}
```

Cause: The error is probably due to any of the following reasons:

- Client's IP address is not within the global allowlist.
- Access-Based Enumeration (ABE) is enabled
- AD group credentials are incorrect
- The LDAP server is unreachable

To confirm the IP address that is being denied, SSH to a node and run the following command:

```
grep -r <ip address> ~/logs/bridge_exec.*INFO*
```

Workaround: Depending on the root cause, use one of the following workarounds:

- Reconfigure the global allowlist to include the client's IP address, or override the global allowlist within the View configuration. For more information, see Add Subnets to Allowlists.
- Disable the **Access-Based Enumeration** option. For more information see, Create Views Using a Predefined Template.
- Refer to the Slow Cohesity Dashboard logins when authenticating with Active Directory domains KB article for guidance on how to troubleshoot an unreachable LDAP server.

Collect Information for SMB Troubleshooting

If you are unable to resolve an issue and need to contact Support, Cohesity recommends collecting the following information to assist Support with troubleshooting:

- Bridge logs: Search for LDAP errors. For information on capturing bridge logs, see Capture Logs.

  For example, `logs/bridge_exec.WARNING:E0525 12:33:20.113982 53125 ldap_request_executor.cc:165] Failed to acquire kerberos credentials for account: COHESITY$@TESTORG..ORG`

- Packet traces between the Cohesity cluster and the client: Use Wireshark to capture the traces on the client or use the `tcpdump` command on the cluster Node CLI.
  - Time of failure
  - Client IP file name and user name
  - Map using the IP address and not the fully qualified domain name to control the Node
- Packet traces between the Cohesity cluster and the domain controller for debugging authentication issues: Use the `tcpdump` command on the cluster Node CLI.

    - To capture only the SMB traffic on a Node:

      ```
      sudo tcpdump -i bond0 -w file.pcap -s 0 tcp port 445
      ```

      Cohesity recommends capturing packet traces without any port filters for a short duration in order to capture SMB, LDAP, Kerberos, and DNS to domain controller traffic.
    - To limit the capture to a single client, specify the IP address of the client in the command:

      ```
      sudo tcpdump -i bond0 -w file.pcap host <client-ip>
      ```

- Client information:
    - Client OS version
    - Other applications, such as anti-virus software, installed on the client.
- The UNC path.
- Owner of the file and any ACLs applied permissions
- Group membership of the user.

Network File System (NFS)

Cohesity cluster supports the Network File System (NFS) protocol. You can create an NFS view and mount it on NFS clients using the NFS protocol.

When you mount an NFS view on NFS clients, the clients need to authenticate and authorize before they access the view on the Cohesity cluster. The Cohesity cluster authenticates the credentials against configured AD Kerberos or MIT Kerberos and authorizes using LDAP.

For more information on multiprotocol access, see ID Mapping for Multi-Protocol Views.

> **Note:** Cohesity recommends using a Linux client with kernel version 4.x or higher.

Supported NFS Versions

The following NFS versions are supported on the Cohesity cluster:

- NFSv3
- NFSv4.1

Considerations for using NFSv4.1 Protocol

The following are the considerations when using NFSv4.1:

- You can select the NFSv4.1 protocol when creating a view in the Cohesity UI, to access the view over NFSv4.1. If the protocol was not enabled during view creation, you can modify the view properties to enable NFSv4.1.
- Both the NFS client and server must support NFSv4.1 to access an NFSv4.1 view
- You can create an NFSv3.0 and NFSv4.1 view with Kerberos and without Kerberos environments.
- In a Kerberos environment, NFS users of client machines should have corresponding accounts in both LDAP and KDC.
- LOCKT operations are not supported with NFSv4.1.
- Cohesity supports the pseudo file system, which allows you to mount the root directory. However, for security reasons, Cohesity does not support the `ls` operation after you mount the root directory.

Support for NFSv4.1 ACLs

An Access Control List (ACL) is a list of permissions associated with a file or directory. NFSv4.1 supports adding ACLs for a named user and special principals. Cohesity supports ACLs only on NFSv4.1 native views.

Example NFSv4.1 ACL

The following is an example of the NFSv4.1 ACL:

```
A:d:user@nfsdomain.org:rxtncy
```

The following section describes the example NFSv4.1 ACL:

- `A` is the ACE type and denotes Allow. This ACL allows the user or a group to perform actions requiring permissions.
- `d` is the inheritance flag and denotes that newly created subdirectories will inherit the ACE.
- `user@nfsdomain.org` is the ACE principal and denotes the named user or one of the following special principals:
  - OWNER@
  - GROUP@
  - EVERYONE@
- `rxtncy` are the permissions this ACE is allowing. You can specify one or more permissions.

Default NFSv4.1 ACL Behavior

By default, NFSv4.1 ACL denies any action if permission is not explicitly granted by an Allow ACE.

If a principal is attempting to perform an action over NFSv4.1 and if this action requires one or more permissions, an access check is performed. If an NFSv4.1 ACL is present, the system evaluates every ACE until a match is detected or until the end of the ACL. If every requisite permission is granted by Allow ACE(s) and is not forbidden by Deny ACE(s), the principal is allowed access. If any permission is denied by a Deny ACE, the system does not proceed with the evaluation, and the principal is denied access.

During the access check, if any permission is explicitly allowed or denied once, subsequent ACEs in the ACL cannot modify this permission.

NFSv4.1 ACL Limitations

- NFSv4.1 ACLs are not supported on unified views.
- POSIX ACLs are not supported on NFSv4 views.
- NFSv4.1 ACLs cannot be enforced on NFSv3 views.

NFS Security Modes for Kerberos Environments

The NFS security modes are available while creating a view with NFS protocol (NFSv3 or NFSv4.1). The security modes are applicable in a Kerberos environment.

The NFS security modes provide security to the data over the wire.

> **Note:** LDAP is a must for Kerberos (AD and MIT). However, if the security mode is Unix authentication (sec=sys) you do not need to have LDAP.

When you mount an NFS view on a client in the Kerberos environment, you must use one of the selected security modes.

The following NFS security modes are supported on the Cohesity cluster.

NFS Security Mode	Description
Unix Authentication	No support for authentication, NFS connections are authorized at a high level based on the Cohesity cluster configured subnets. For more information, see Mount a View Using NFS
Kerberos Authentication	Mount a view on a client using sec=krb5. Kerberos Authentication (krb5) supports only identity verification.
Kerberos Integrity	Mount a view on a client using sec=krb5i Kerberos Integrity (krb5i) supports data integrity services along with identity verification.

NFS Security Mode	Description
Kerberos Privacy	Mount a view on a client using sec=krb5p Kerberos Privacy (krb5p) supports data privacy along with integrity services and identity verification.

> **Note:** Unix authentication is the lowest security mode and krb5p is the highest security mode.

Considerations when Selecting NFS Security Modes

The following is a list of considerations you should consider for NFS security modes:

- By default, 'Unix Authentication' is selected when creating a view. However, you can choose the required security modes.
- You can only mount a view by using the selected security modes.

  For example, you selected Unix Authentication and Kerberos Privacy, so while mounting a view on a client, you can use any one of them.

  If you try to mount with other NFS security modes, the Cohesity cluster does not allow you to mount.

Prevent NFSv4.1 Client ID Collision

NFSv4.1 requires clients to identify themselves to servers with a unique string. File open and lock state shared between one client and one server is associated with this identity. To support robust NFSv4.1 state recovery and transparent state migration, this identity string must not change across client reboots.

If the hostname used by an NFSv4.1 client collides, the client is unable to mount over NFSv4.1. Hostname collision leads to client ID collision.

To prevent NFSv4.1 client ID collision, the NFS administrator must do the following:

1. Resolve hostname conflict.
2. Use the `nfs4_unique_id` parameter in the NFS client configuration.

   The `nfs4_unique_id` boot parameter specifies a unique string that can be used instead of the node name of a system when an NFS client identifies itself to a server. If the node name of the system is not unique, or it changes, its `nfs4_unique_id` remains the same, preventing collision with other clients or loss of state during NFS reboot recovery or transparent state migration.

NFSv 4.1 Unsupported Features

Cohesity does not support the following NFSv4.1 features:

- pNFS (Parallel NFS)
- Delegations - File and directory
- Volatile filehandle

> **Note:** Writing files larger than 1TB on Cohesity Views using the NFSv4.1 protocol may fail with the 'file too large' error. For more information, see the KB article.

NFS Exports

Mounting and accessing views as NFS exports requires the following high level steps.

1. Review the information about optimizing storage domains and views for your use case.

   See SmartFiles (File and Object Services) Overview.

2. Create storage domains and views for your use case.

   See Create or Edit Storage Domains and Create Views Using a Predefined Template.

3. Add the systems that will mount and access views to the allowlist.

   See Add Subnets to Allowlists.

4. Mount the views to systems so users and applications can store data in them.

   See Mount a View Using NFS.

Once the above steps are completed, you can use a view to store data such as files and directories. Other backup products, such as Veeam, can also store data in a view, as shown in the following figure.

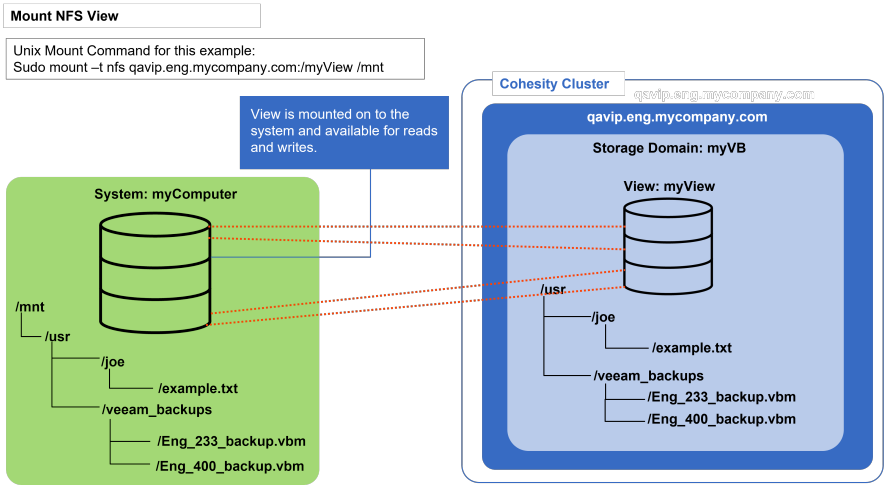

Add Subnets to Allowlists

Before you can mount and access a view, ensure the IP address of the system that is mounting the view is specified in a subnet that has been added to an allowlist.

Cohesity supports the following types of allowlists:

- **Global Allowlist**—The global allowlist applies to all views on the Cohesity cluster, but is overridden by a view allowlist or view share allowlist, if they exist.
- **View Allowlist**—Only the systems with IP addresses in the subnets listed in the view's allowlist have privileges to access and mount this view.

Update the Global Allowlist

To add a subnet to the Cohesity cluster's global allowlist:

1. Navigate to **SmartFiles > Views**.
2. Click the **Global Settings** tab.
3. In the **Global Subnet Allowlist** section, click **Add**.
4. In the **Add Allowlist** dialog, enter a Subnet IP in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).

   > **Note:** CIDR prefix values are allowed from 0 to 32.

5. Select any or all permissions as required.

   > **Note:** By default, the ability to set **S3 Permissions** for restricting access to specific subnets is enabled.

6. Select **NFS Root Squash** as applicable.
7. In the **Description** field, enter a description for the allowlist.

8. Click **Add**:

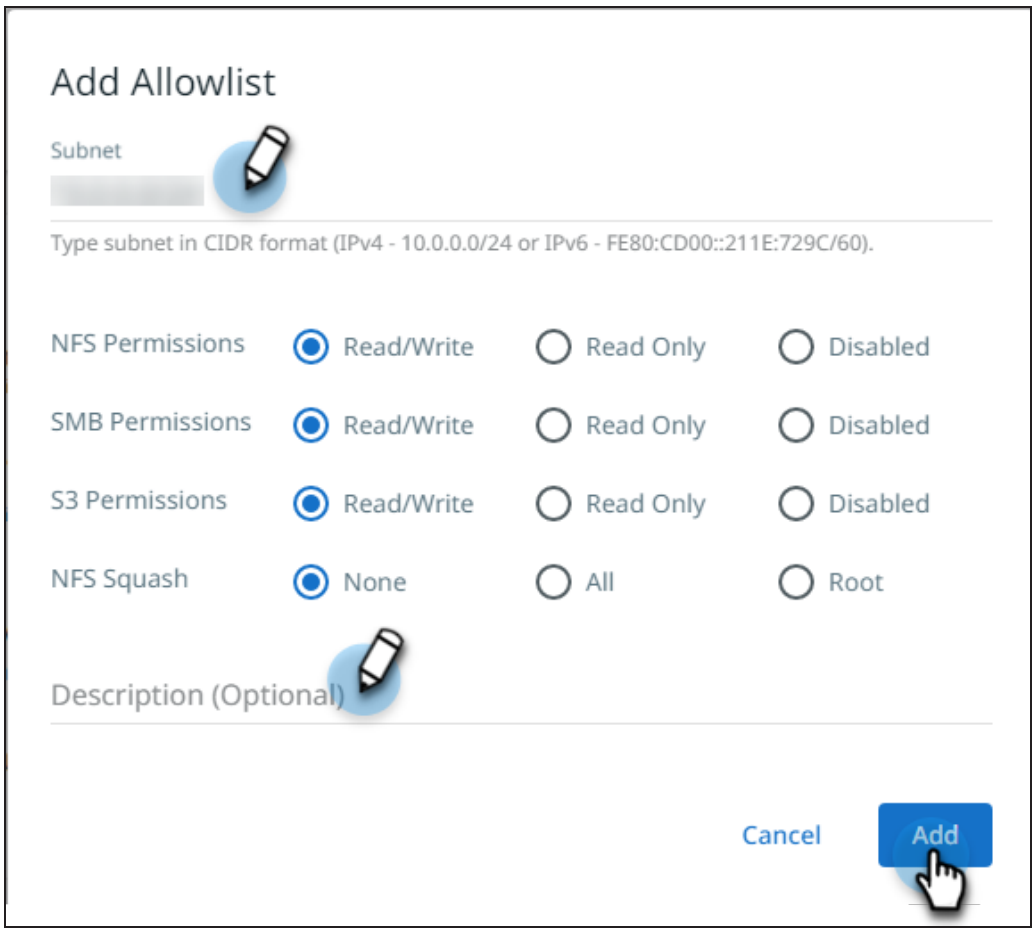

9. Repeat the preceding steps for specifying additional subnets in the global allowlist.

Update a View Allowlist

To add a subnet to an existing view's allowlist:

1. Navigate to **SmartFiles > Views**.
2. Locate the required View in the list, click the ⋮, and select **Edit** ( ).
3. In Advanced settings, click **Edit** for the Security category.
4. **Override Global IP Allowlist:** Turn on **Override Global IP Allowlist** if you want to specify a custom list for this view. Only the systems with IP addresses in the subnets listed in the view's allowlist have privileges to access and mount this view. If you enable the **Override Global IP Allowlist** toggle, the machines listed in the subnets of the global Cohesity cluster allowlist cannot access or mount this view unless the IP address of the machine is also listed in the View's Allowlist.

   Extend Global IP Allowlist: Allows you to extend the Global IP Allowlist.

   To add subnets to the allowlist of this view:

a. Under **IP Allowlist**, select: **Override Global IP Allowlist** or **Extend Global IP Allowlist**.
b. Click **Add**.
c. In the **Add Allowlist** screen, type subnet in CIDR format.
d. Select any or all permissions as required.
e. Select **NFS Squash** options.

The **NFS all squash** option is like NFS Root Squash. However, for given IP addresses, all users are mapped to the All Squash User/Group configured in the view. The default All Squash User for all views is "nobody".

> **Note:** If the **Override Global Allowlist** toggle is off, this view can only be accessed and mounted by systems specified by the global Cohesity cluster allowlist.

5. Click **Update View**.

For more information on global settings for each protocol, see the following topics:

- Global Settings for NFS
- Global Settings for SMB

Global Settings for NFS

In addition to the settings provided in Add Subnets to Allowlists, you can allowlist Netgroup that have permissions to all views.

After the Cohesity cluster is upgraded to 6.6.0c, note the following:

Existing NFS or SMB view—IP restriction is retained. If you access an existing NFS or SMB view using the S3 protocol, the request fails if the existing allowlist does not include the IP address of S3 clients. Before or after the upgrade, ensure to add the IP address or subnet at the view level if the NFS or SMB view is accessed using the S3 protocol.

Global Netgroup Allowlist

You can allowlist Netgroup that has permissions to all views.

Perform the following steps on the Cohesity cluster to allowlist Netgroup:

1. Navigate to **SmartFiles > Views**.
2. Click the **Global Settings** tab.
3. In the **Global Netgroup Allowlist** section, click **Add** and do the following:
    1. In the **Netgroup Name** field, enter a name.
    2. Select **NFS Permissions** (Read/Write, Read Only or Disabled).

3. Select **NFS Squash** permissions (None, Root, All).
4. Click **Add**.

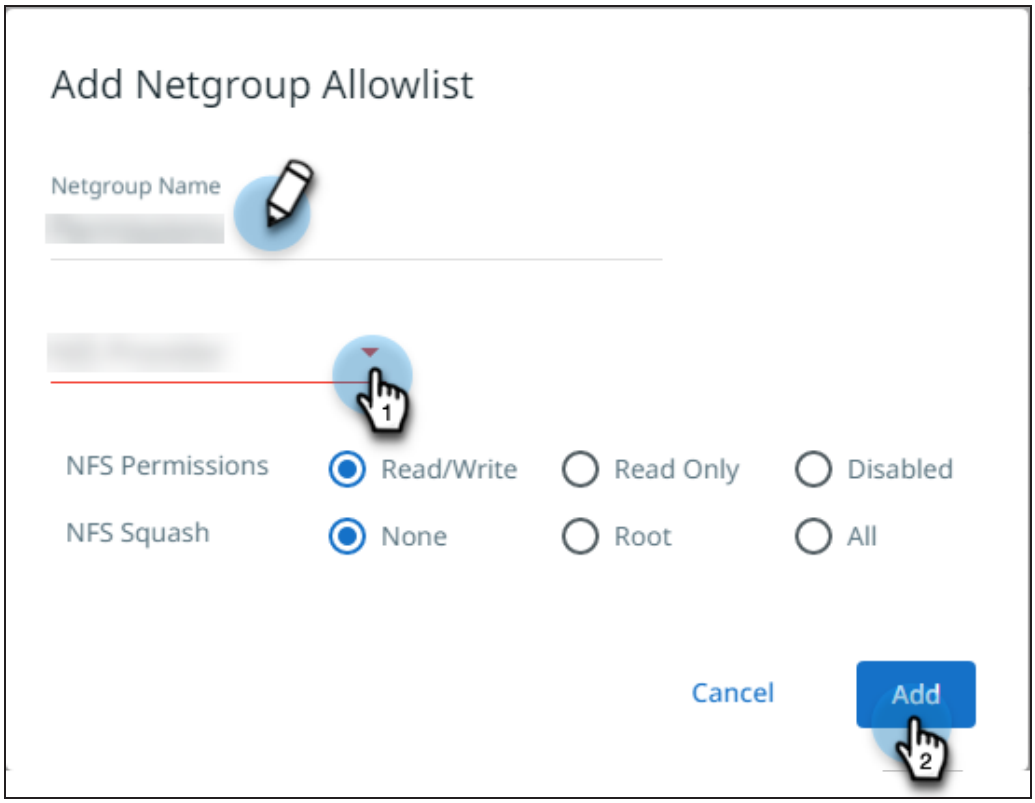

Mount a View Using NFS

Before storing data on a view, you have to mount it using NFS export on your Windows or Linux system.

Prerequisites

- The view must support the NFS protocol.
- The client system that will mount the view must be in a subnet in a Cohesity allowlist. Otherwise, the system will not have privileges to mount and access the view.
- Firewall port 2049 must be open for NFS operations on the Cohesity cluster. Port 2049 is the default port. You can change it as described Change the NFS Firewall Port.
- When mounting a view on a Windows system using NFS, the view must be browseable from NFS.

Considerations:

- NFS mount names and names of files contained in the mount support ASCII and UTF-8 character codes only.

- When data is deleted from a view, it may take up to a day for the disk space to become available again and visible from utilities such as `df`.
- To register an Oracle RAC or a RAC node as a physical server, the "host" command must be executed on each of the nodes of that RAC.
- Cohesity recommends using a Linux client with kernel version 4.x or higher.
- NFSv4.1 limitation:
    - LOCKT operations are not supported.

If you mount a view on a Windows system using NFS, you can read files from the mounted drive but not write to the view. To resolve this issue, change the UID and GID in the Windows Registry, as described in Mount a View with NFS protocol on Windows. (This issue does not apply if the view is mounted on Windows using SMB.)

NFS Client Mount Parameters

(missing or bad snippet)

The NFS mount options depend on the workload and the application. There is no standard recommendations for the mount options.

The following table lists the recommendations for NFS client mount parameters for file shares and backup target workloads in different operating systems.

NFSv3 Mount Parameters

Host	Recommended NFS Mount Parameters
Linux	**For Backup Target workload**: rw,relatime,vers=3,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none **For File Shares workload**: rw,relatime,vers=3,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none
UNIX (AIX / HP-UX / Solaris)	**For both File Shares & Backup Target workload**: bg,hard,intr,rsize=32768,wsize=32768,vers=3,proto=tcp,sec=sys,rw,acdirmax=30

NFSv4.1 Mount Parameters

Host	Recommended NFS Mount Parameters
Linux	**For Backup Target workload:** rw,relatime,vers=4.1,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys **For File Shares workload:** rw,relatime,vers=4.1,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys
UNIX (AIX / HP-UX / Solaris)	**For both File Shares & Backup Target workload:** bg,hard,intr,rsize=32768,wsize=32768,vers=4.1,proto=tcp,sec=sys,rw,acdirmax=30

Different versions of a host operating system can have different default NFS mount options.

The NFS mount options depend on the workload and the type and version of the host operating systems and applications. The mount parameters also vary based on various backup and restore operations and the deployment specifications. For example, Oracle deployment type, version and the workflow.

The following table lists the NFS client mount parameters while using Cohesity as the **Oracle Backup Target**.

Host	Deployment type	Version	Default mount options	Full/Incremental Backup	Log Backup	Overwrite Restore	Alternate Restore	Clone
Linux	RAC	12c	For NFSv3: rw,relative,vers=3,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none For NFSv4.1: rw,relative,vers=4.1,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys					For NFSv3: noac,local_lock=all For NFSv4.1: local_lock=all
		11g		noac				
	Standalone	12c						
UNIX (AIX / HP-UX / Solaris)	RAC	12c	For NFSv3: bg,hard,intr,rsize=32768,wsize=32768,vers=3,proto=tcp,sec=sys,rw,acdirmax=30 For NFSv4.1: bg,hard,intr,rsize=32768,wsize=32768,vers=4.1,proto=tcp,sec=sys,rw,acdirmax=30			noac		For NFSv3: noac,local_lock=all For NFSv4.1: local_lock=all
		11g		noac				
	Standalone	12c						
		11g						

The following table lists the NFS client mount parameters while using Cohesity as a **Backup Storage Target** for the backup applications such as **CommVault** or **Veeam**.

Host	Mount Parameters
CommVault/ Veeam	For NFSv3: noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,timeo=10000,hard,intr,nolock For NFSv4.1: noatime,vers=4.1,proto=tcp,rsize=1048576,wsize=1048576,timeo=10000,hard,intr

Mount a View with NFS Protocol

Perform the following steps to copy the view mount path:

1. In the Cohesity Dashboard, navigate to **File Services** > **Views** > **All Views**.
2. Browse for the view in the list
3. Click the Mount Paths icon of NFS.
4. In the Mount Paths screen, click the copy icon for **NFS mount path** and click **Close**.

For more information, see the following topics:

- Mount a View with NFS protocol using UNC Path
- Mount a View with NFS protocol on Linux/Unix
- Mount a View with NFS protocol on Windows
- Mount a View with Kerberos-Based NFS

Mount a View with NFS protocol using UNC Path

Some applications may require the NFS storage target in UNC path format.

For Example, while using Cohesity as a backup target for Veeam or CommVault, the View can be accessed with NFS protocol using the UNC Path in the below format:

```
\\<Cohesity_Cluster_Hostname_or_VIP>\<NFS_View_Name>
```

Mount a View with NFS protocol on Windows

> **Note:** If you are mounting a View on Windows using NFS, then you must disable NLM. To disable NLM, contact Cohesity Support.

To mount a view with NFS protocol on Windows:

1. On the Windows system, as the administrator, change the UID and GID in the Windows Registry. Otherwise, you will be able to read files from the mounted drive, but not write to the view. (This step is not applicable when mounting a Windows view using SMB.)

a. In the Windows Start menu, start the Registry Editor by searching for `regedit` and clicking it.
b. Navigate to the following path: `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\ClientForNFS\CurrentVersion\Default`
c. Right-click `Default` and select **New** > **DWord (32-bit) Value**.
d. Enter `AnonymousUid` for the name and leave the value as `0`.
e. Right-click `Default` and select **New** > **DWord (32-bit) Value** again.
f. Enter `AnonymousGid` for the name and leave the value as `0`.
g. Reboot the Windows system.

2. Install the NFS Client on your Windows system.
3. Mount the View with NFS protocol on the Windows client using the mount command.

   Sample Command:

   ```
   mount -o mtype=hard fileaccess=777 nolock <Cohesity_Cluster_Hostname_or_VIP>:/<NFS_View_Name> <Windows_Drive_Letter>
   ```

   Example:

   ```
   mount -o mtype=hard fileaccess=777 nolock host.example.com:/NFS1 Z:
   ```

4. Verify the mounted NFS file systems with the `mount` command.

   ```
   C:\Users\Administrator&gt;mount
   Local   Remote                                  Properties
   ------------------------------------------------------------------------
   ---
   z:      \\host.example.com\NFS1                 UID=0, GID=0
                                                   rsize=1048576,
   wsize=1048576
                                                   mount=hard, timeout=10.0
                                                   retry=1, locking=no
                                                   fileaccess=777, lang-ANSI
                                                   casesensitive=no, sec=sys
   ```

Mount a View with NFS protocol on Linux/Unix

To mount a view with NFS protocol on Linux or Unix:

1. Install the `nfs-common` and `rpcbind` modules on the Linux/Unix client, if they are not already installed:
2. Create a directory to mount the share.
3. Mount the NFS volume using the `mount` command.

- Specify the NFS version in the options parameter `vers` while mounting the view. For example, if you are mounting an NFSv4.1 View, specify `vers=4.1`. If you are mounting an NFSv3 View, specify `vers=3`.

Sample Command (for File Services workload)

```
mount -t nfs -o
rw,relatime,vers=3,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,t
imeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none
<Cohesity_Cluster_Hostname_or_VIP>:/<NFS_View_Name> <View Mount Path>
```

Example:

```
mount -t nfs -o
rw,relatime,vers=3,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,t
imeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none
host.example.com:/NFS1 /home/user/nfsdir
```

4. Make NFS mounts persist across reboots (Make an entry in the `/etc/fstab` file. Steps may vary depending upon the Host Operating system).

For more information on the NFS client mount parameters, see NFS Client Mount Parameters(missing or bad snippet)The NFS mount options depend on the workload and the application. There is no standard recommendations for the mount options.The following table lists the recommendations for NFS client mount parameters for file shares and backup target workloads in different operating systems.NFSv3 Mount ParametersHost Recommended NFS Mount Parameters Linux For Backup Target workload: rw,relatime,vers=3,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none For File Shares workload: rw,relatime,vers=3,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none UNIX (AIX / HP-UX / Solaris) For both File Shares & Backup Target workload: bg,hard,intr,rsize=32768,wsize=32768,vers=3,proto=tcp,sec=sys,rw,acdirmax=30 NFSv4.1 Mount ParametersHost Recommended NFS Mount Parameters Linux For Backup Target workload: rw,relatime,vers=4.1,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sysFor File Shares workload: rw,relatime,vers=4.1,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys UNIX (AIX / HP-UX / Solaris)For both File Shares & Backup Target workload:bg,hard,intr,rsize=32768,wsize=32768,vers=4.1,proto=tcp,sec=sys,rw,acdirmax=30Different versions of a host operating system can have different default NFS mount options.The NFS mount options depend on the workload and the type and version of the host operating systems and applications. The mount parameters also vary based on various backup and restore operations and the deployment specifications. For example, Oracle deployment type, version and the workflow.The following table lists the NFS client mount parameters while using Cohesity as the Oracle Backup Target.HostDeployment typeVersion Default mount optionsFull/Incremental BackupLog BackupOverwrite RestoreAlternate

RestoreCloneLinuxRAC12cFor NFSv3:rw,relatime,vers=3,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=noneFor NFSv4.1:rw,relatime,vers=4.1,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys For NFSv3: noac,local_lock=all For NFSv4.1: noac 11gnoac Standalone12c For NFSv3: local_lock=all UNIX (AIX / HP-UX / Solaris) RAC12cFor NFSv3:bg,hard,intr,rsize=32768,wsize=32768,vers=3,proto=tcp,sec=sys,rw,acdirmax=30For NFSv4.1:bg,hard,intr,rsize=32768,wsize=32768,vers=4.1,proto=tcp,sec=sys,rw,acdirmax=30 noac For NFSv3: noac,local_lock=all For NFSv4.1: noac 11gnoac Standalone12c For NFSv3: local_lock=all 11g The following table lists the NFS client mount parameters while using Cohesity as a Backup Storage Target for the backup applications such as CommVault or Veeam.HostMount ParametersCommVault/ Veeam For NFSv3: noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,timeo=10000,hard,intr,nolock For NFSv4.1: noatime,vers=4.1,proto=tcp,rsize=1048576,wsize=1048576,timeo=10000,hard,intr .

Mount a View with Kerberos-Based NFS

You can mount both NFS v3.0 and NFS v4.1 versions Views.

Prerequisites

- Before mounting a Kerberos-based view, configure the NFS client. For more information, see Configure Unix client for Mounting Kerberos Based View.
- The View must support the NFSv3.0 or NFSv4.1 protocol.
- The client system that will mount the View must be in a subnet in a Cohesity allowlist. Otherwise, the system will not have privileges to mount and access the view.
- Firewall port 2049 must be open for NFS operations on the Cohesity cluster. Port 2049 is the default port. You can change it as described Change the NFS Firewall Port.
- When mounting a view on a Windows system using NFS, the view must be browsable from NFS.

To mount a view with the Kerberos-based NFS, complete the following steps:

1. Install the `nfs-common` and `rpcbind` modules on the Kerberos client, if they are not already installed:
2. Create a directory to mount the share.
3. Mount the NFS volume using the mount command.
    - While mounting a view, specify the NFS version in the options parameter `vers`. For example, if you mount an NFSv4.1 View, specify `vers=4.1` and if you mount

an NFSv3 View, specify `vers=3`.

Sample Command (for File Services workload)

```
mount -t nfs -o
rw,relatime,vers=4.1,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp
,timeo=600,retrans=2,sec=krb5,mountvers=4.1,mountproto=tcp,local_
lock=none <Cohesity_Cluster_Hostname_or_VIP>:/<NFS_View_Name> <View
Mount Path>
```

Example:

```
mount -t nfs -o
rw,relatime,vers=4.1,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp
,timeo=600,retrans=2,sec=krb5,mountvers=4.1,mountproto=tcp,local_
lock=none host.example.com:/NFS1 /home/user/nfsdir
```

> **Note:** If you have selected specific NFS security modes for a View, you must use the same security modes when mounting a Kerberosed-based View. For example, if you have selected krb5i as the NFS security mode when creating a View, you must use the same.

4. Make NFS mounts persist across reboots (Usually done by making an entry in `/etc/fstab` file. Steps may vary depending upon the Host Operating system).

For more information on the NFS client mount parameters, see NFS Client Mount Parameters(missing or bad snippet)The NFS mount options depend on the workload and the application. There is no standard recommendations for the mount options.The following table lists the recommendations for NFS client mount parameters for file shares and backup target workloads in different operating systems.NFSv3 Mount ParametersHost Recommended NFS Mount Parameters Linux For Backup Target workload: rw,relatime,vers=3,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none For File Shares workload: rw,relatime,vers=3,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none UNIX (AIX / HP-UX / Solaris) For both File Shares & Backup Target workload: bg,hard,intr,rsize=32768,wsize=32768,vers=3,proto=tcp,sec=sys,rw,acdirmax=30 NFSv4.1 Mount ParametersHost Recommended NFS Mount Parameters Linux For Backup Target workload: rw,relatime,vers=4.1,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sysFor File Shares workload: rw,relatime,vers=4.1,rsize=32768,wsize=32768,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys UNIX (AIX / HP-UX / Solaris)For both File Shares & Backup Target workload:bg,hard,intr,rsize=32768,wsize=32768,vers=4.1,proto=tcp,sec=sys,rw,acdirmax=30Different versions of a host operating system can have different default NFS mount

options.The NFS mount options depend on the workload and the type and version of the host operating systems and applications. The mount parameters also vary based on various backup and restore operations and the deployment specifications. For example, Oracle deployment type, version and the workflow.The following table lists the NFS client mount parameters while using Cohesity as the Oracle Backup Target.HostDeployment typeVersion Default mount optionsFull/Incremental BackupLog BackupOverwrite RestoreAlternate RestoreCloneLinuxRAC12cFor NFSv3:rw,relatime,vers=3,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=noneFor NFSv4.1:rw,relatime,vers=4.1,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys For NFSv3: noac,local_lock=all For NFSv4.1: noac 11gnoac Standalone12c For NFSv3: local_lock=all UNIX (AIX / HP-UX / Solaris) RAC12cFor NFSv3:bg,hard,intr,rsize=32768,wsize=32768,vers=3,proto=tcp,sec=sys,rw,acdirmax=30For NFSv4.1:bg,hard,intr,rsize=32768,wsize=32768,vers=4.1,proto=tcp,sec=sys,rw,acdirmax=30 noac For NFSv3: noac,local_lock=all For NFSv4.1: noac 11gnoac Standalone12c For NFSv3: local_lock=all 11g The following table lists the NFS client mount parameters while using Cohesity as a Backup Storage Target for the backup applications such as CommVault or Veeam.HostMount ParametersCommVault/ Veeam For NFSv3: noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,timeo=10000,hard,intr,nolock For NFSv4.1: noatime,vers=4.1,proto=tcp,rsize=1048576,wsize=1048576,timeo=10000,hard,intr .

List Active NFS Connections

To list active NFS connections for views:

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. List the active NFS connections for views by using the Cohesity CLI `view list-nfs-connections` command:

   ```
   admin>view list-nfs-connections
   ```

> **Tip:** You can also use the CLI Help feature to view the `view list-nfs-connections` options by entering the following command: `view list-nfs-connections help`.

CLI Option	Example	Description
`names=`	`ViewName1,ViewName2`	Optional. Specify a comma separated list of view names. The active NFS connections for these views are displayed.

NLM Locking over NFS

File-level locking mechanisms are supported to restrict access to a file or portion of a file stored on the server to only one user or process at a time. Users can specify the byte range to lock or specify the entire file on the server to be locked. System administrator can use iris_cli commands for listing and clearing locks acquired on a file. Locks can be cleared from iris_cli if the client crashes while holding a lock for a file on the server.

NLM Locks are of two types: Shared or Exclusive.

- When a Shared or a Read lock is used, it prevents other users or processes from requesting a **write** lock to the file or a specific byte range of the file. The shared lock can be used by multiple users.
- Exclusive or Write Locks provides **exclusive** access for editing the file contents. When the exclusive lock is used, then no other user or process can lock that part of the file.
- Asynchronous NLM lock/unlock requests for files are not supported. For example, the client will not get a callback from the server in response to a lock/unlock request for a file.

View the NLM Locks

CLI Option	Description
view list-nlm-locks name=*<name>* file-path=*<file_path>*	To view the list of all the NLM locks based on the view and the filepath.
view list-nlm-locks name=*<name>* file-path=*<file_path>* verbose=true	To list all ranges locked for a file by all clients.
view clear-nlm-locks client-id=*<id>*	To clear all the existing NLM locks held by the client.

Change the NFS Firewall Port

NFS operations on the cluster use firewall port number 2049 by default. The following procedure explains how to change the port number.

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Change the Bridge service port to the desired port. To change the Bridge service port, contact Cohesity Support.
3. Add a new firewall port profile by using the following command. This procedure uses the profile name `new-nfs-1`, but you can change it.

   ```
   iris_cli firewall-profile add name=new-nfs-1 ports=<new_port>/tcp
   force=true
   ```

   Example:

   ```
   [cohesity@node-1 ~]$ iris_cli firewall-profile add name=new-nfs-1
   ports=60001/tcp force=true
   Success: Successfully added firewall profile.
   ```

4. Activate the attachment by using the following command:

   ```
   iris_cli firewall-profile activate profile-name=new-nfs-1 interface-
   groups=intf_group1 action=allow
   ```

   Example:

   ```
   [cohesity@node-1 ~]$ iris_cli firewall-profile activate profile-
   name=new-nfs-1 interface-groups=intf_group1 action=allow
   Success: Successfully activated the profile.
   ```

5. From the shell prompt, add the IPV4 rule by using the following command:

   ```
   sudo firewall-cmd --direct --add-rule ipv4 filter INPUT 0 -m tcp -p tcp
   --dport <new_port> -j ACCEPT
   ```

   Example:

   ```
   [cohesity@node-1 ~]$ sudo firewall-cmd --direct --add-rule ipv4 filter
   INPUT 0 -m tcp -p tcp --dport 60001 -j ACCEPT
   success
   ```

6. Add the IPV6 rule by using the following command:

```
sudo firewall-cmd --direct --add-rule ipv6 filter INPUT 0 -m tcp -p tcp
--dport <new_port> -j ACCEPT
```

Example:

```
[cohesity@node-1 ~]$ sudo firewall-cmd --direct --add-rule ipv6 filter
INPUT 0 -m tcp -p tcp --dport 60001 -j ACCEPT
success
```

7. Make the IPV4 rule permanent by using the following command:

```
sudo firewall-cmd --permanent --direct --add-rule ipv4 filter INPUT 0 -
m tcp -p tcp --dport <new_port> -j ACCEPT
```

Example:

```
[cohesity@node-1 ~]$ sudo firewall-cmd --permanent --direct --add-rule
ipv4 filter INPUT 0 -m tcp -p tcp --dport 60001 -j ACCEPT
success
```

8. Make the IPV6 rule permanent by using the following command:

```
sudo firewall-cmd --permanent --direct --add-rule ipv6 filter INPUT 0 -
m tcp -p tcp --dport <new_port> -j ACCEPT
```

Example:

```
[cohesity@node-1 ~]$ sudo firewall-cmd --permanent --direct --add-rule
ipv6 filter INPUT 0 -m tcp -p tcp --dport 60001 -j ACCEPT
success
```

9. Verify the port number changed in iptables by using the following command:

```
sudo iptables -nvL | grep <new_port>
```

Example:

```
[cohesity@node-1 ~]$ sudo iptables -nvL | grep 60001
    0     0 ACCEPT     tcp  --  *      *       0.0.0.0/0
0.0.0.0/0           tcp dpt:60001
```

10. Verify the port number changed in direct.xml by using the following command:

```
sudo grep <new_port> /etc/firewalld/direct.xml
```

Example:

```
[cohesity@node-1 ~]$ sudo grep 60001 /etc/firewalld/direct.xml
  <rule priority="0" table="filter" ipv="ipv4" chain="INPUT">-m tcp -p
tcp --dport 60001 -j ACCEPT</rule>
  <rule priority="0" table="filter" ipv="ipv6" chain="INPUT">-m tcp -p
tcp --dport 60001 -j ACCEPT</rule>
```

11. Restart the Bridge service by using the following command:

```
[cohesity@node-1 ~]$ iris_cli cluster restart-services service-
names=bridge
```

Objects (S3)

On a Cohesity cluster, you can create a SmartFiles S3 View. S3 Views are similar to Amazon S3 buckets. S3-compatible clients can use the S3 View to store and retrieve objects. As SmartFiles S3 manages data as objects, you can access files as S3 objects.

> **Tip:** You can use SmartFiles to automatically archive or tier data from a Cohesity cluster to Amazon S3 buckets. For these use cases, register an existing Amazon S3 bucket as an external target on the Cohesity cluster; Cohesity S3 Views are not required. For more information, see About Cloud Archival and About Cloud Tier.

> **Note:** Cohesity supports transferring payload in multiple Chunks (Chunked Upload) when using AWS Signature Version 4. This is required by default for the latest versions of S3 SDKs. For more information, see Chunk Upload Format.

Cohesity S3 supports the checksum algorithms, CRC32, CRC32C, CRC64NVME, SHA1, and SHA256. You can provide these algorithms in the S3 requests and store them on the S3 Objects Metadata. When the request does not specify any algorithm, Cohesity S3 uses CRC64NVME.

Prerequisite

- Firewall port 3000 (TCP, incoming to Cohesity cluster) must be open to support S3 protocol access to the views on the Cohesity cluster.

Considerations

- You must use one of the following accounts to create an S3 View:
  - A local Cohesity user.
  - An Active Directory user that was explicitly added to the Cohesity cluster and assigned a role. This user does not rely on an AD group for access to the Cohesity cluster.

> **Important:**
> You cannot create an S3 View using one of the following accounts:
>
> - An AD user that has Cohesity cluster access through an Active Directory group only
> - An SSO user
> - A Helios user

- To create a SmartFiles S3 View in a multi-tenant environment, log in to the Cohesity cluster as an Organization user. If you create the S3 View while impersonating an organization, the Service Provider administrator becomes the owner of the S3 View.
- Access Control Lists (ACLs) can be set on a bucket using the AWS CLI.
- You cannot use NFS to mount newly created S3 Views. However, if there are existing S3 Views that were configured to use NFS, you can mount such S3 Views using NFS.
- The maximum number of versions allowed per S3 object is 500,000.
- Cohesity recommends excluding any unsupported header(s) from your requests. By doing so, you can prevent any potential unintended consequences that may arise from using unsupported headers.
- **Download** – Download is not supported for S3 workflows including indexed browsing, non-indexed browsing, search, and recovery.
- **Recover**
  - Recovery of archival S3 objects is not supported.
  - Prefix-based recovery is not allowed for S3 objects.
  - S3 object recovery is allowed only across S3 Views.
  - Recovery from S3 2.0 View to S3 1.0 View is not supported.
  - Only object-based recovery is supported; prefix-based recovery is not supported.
- **S3 Indexing**

  The following are the limitations for S3 indexed browse. Since S3 object storage is inherently non-hierarchical (flat key space), but the browse displays a folder-style hierarchy, mapping between these models can introduce ambiguous scenarios.

  > **Note:** These limitations do not apply to non-indexed browse.

  - Ambiguity occurs between the S3 keys in the hierarchical browse View and native S3 tools in the following cases:

  - **Multiple slashes** – If a S3 key contains continuous multiple slashes, the folder View can appear different from the original key in other S3 tools.

    For example, */foo///xyz/pqr*
  - **File and folder reuse** – If the same S3 path is used both as a file and as a folder, the browse View can display only one version of the path.

    For example, */foo/pqr and /foo/pqr/abc*
  - **Leading slashes** – If one S3 key starts with a slash and another does not, they can still appear as the same path in the browse View.

    For example, *foo/pqr and /foo/pqr*
- For versioned buckets, only the latest version of an object corresponding to the snapshot is considered for indexed browsing, non-indexed browsing, search, and recovery workflows.
- Indexed browse does not support sorting, so results may appear in an order different than non-indexed browse.
- Search in indexed browse displays exact results only when the entire object name is specified.
- Search displays full object names and not prefixes.
- Search up to 10k objects is supported based on the maximum value selected. To display more results, you must refine the search query.

Get Started

The main steps for creating and using a view as a bucket are:

1. Sign in to the Cohesity cluster as a local user or an Active Directory user that was explicitly added to the Cohesity cluster and assigned a role.

   This step is required to establish ownership of your S3 view (bucket).
2. Retrieve your Cohesity S3 access keys from the Cohesity Dashboard and store them in your preferred location. For more information, see Manage Cohesity S3 Access Keys.

   The keys were automatically generated when your user account was created on the Cohesity cluster.
3. Create a view that supports the S3 protocol.

   You will be the owner of the view and will need your S3 access keys to access it through the S3 API and AWS CLI.
4. Optionally, install the AWS CLI so you can interact with the view, for example, to list buckets or objects. For more information, see Use the AWS CLI on Cohesity S3.

Create an S3 View

You can create an S3 View using a predefined template or a custom template.

You can create an S3 View using predefined templates for different use cases and scenarios. Predefined templates are provided by Cohesity. Cohesity provides the following predefined templates for **Object Services**:

- General
- Containers
- Hadoop
- Splunk SmartStore

To create a View using a predefined View template:

1. Navigate to **SmartFiles > Views** and do one of the following on the **Views** page:
    - Click **Create View**. On the **View Templates** side sheet, navigate to the **Predefined** tab and select a predefined template.
    - Go to the **Templates** tab. Click the **Type** filter and select **Predefined**. Click the **Category** filter and select **Object Services**. Hover over the predefined template that you want to use, click the Actions menu (⋮), and select **Create View**.
2. Enter a name for the View.
3. Based on the template you selected, the predefined template populates optimum preconfigured settings. However, if you want to update the settings, click **More Options** and update the settings. For more information about the available options, see S3 View Options.
4. Click **Create** to create the View.

Cohesity supports creating customized View templates and using a customized template to create S3 Views. For more information, see About Templates Tab.

To create an S3 View using a custom template:

1. Navigate to **SmartFiles > Views** and do one of the following on the **Views** page:
    - Click **Create View**. On the **View Templates** side sheet, navigate to the **Custom** tab and select a custom template.
    - Go to the **Templates** tab. Click the **Type** filter and select **Custom**. Click the **Category** filter and select **Object Services**. Hover over the custom template that you want to use, click the Actions menu (⋮), and select **Create View**.
2. Enter a name for the View.
3. Based on the template you selected, the custom template populates preconfigured settings. However, if you want to update the settings, click **More Options** and update

the settings. For more information about the available options, see S3 View Options.

4. Click **Create** to create the View.

Set S3 View as Logging Bucket to Other S3 Views

Cohesity allows you to designate one S3 View as the logging bucket for another S3 View using the S3 Server Logs predefined template in the Cohesity cluster.

Setting a S3 View as a logging bucket using the template includes the following steps:

1. Create the target logging View.
2. Create the source View and set the target View as the logging bucket.

To create the target logging View:

1. Navigate to **SmartFiles > Views** and perform one of the following on the **Views** page:
   - Click **Create View**. On the **View Templates** side sheet, navigate to the **Predefined** tab and under **Object Services**, select the **S3 Server Logs** template.
   - Go to the **Templates** tab. Click the **Type** filter and select **Predefined**. Click the **Category** filter and select **Object Services**. Hover over the **S3 Server Logs** template, click the Actions menu (⋮), and select **Create View**.

     > **Note:** All logging (target) buckets must be created using this template, as it automatically applies recommended defaults for settings such as bucket policy and QoS policy.

2. In the **Create View** window, enter the **View Name**, and select the **Storage Domain**, **Read/Write Protocol**, and **Bucket Owner**.

   The predefined template populates optimum preconfigured settings. However, if you want to update the settings, click **More Options** and update the settings. For more information about the available options, see S3 View Options.

3. Click **Create** to create the target View.

To create the source View and set the target View as the logging bucket:

1. Navigate to **SmartFiles > Views** and perform one of the following on the **Views** page:
   - Click **Create View** and select a predefined or custom template.
   - Go to the **Templates** tab. Click the **Type** filter and select a predefined or custom template. Hover over the template, click the Actions menu (⋮), and select **Create View**.
2. In the **Create View** window, enter the **View Name**.

Based on the template you selected, the optimum preconfigured settings are populated. However, if you want, you can update the settings. For more information about the available options, see S3 View Options.

3. Click **More Options** and navigate to **S3 Server Logs**.
4. Toggle On **Enable S3 Server Logging**.
5. Select the **target bucket** that you created.
6. Optionally, you can set the **Target Prefix**.
7. Select the **Log Key Format** for the logging bucket. The options include **Simple Prefix, Partitioned - Event Time**, and **Partitioned - Delivery Time**.
8. You can also add **Log Object ACLs**. For more information, see the **ACLs** section in S3 View Options.
9. Click **Create** to create the source View.

When you make any changes to the source View like modifying metadata, creating objects and so on, the logs are recorded in the target View.

S3 View Options

Storage Domain

A storage domain is a named storage location on the Cohesity cluster. A storage domain defines the policy for configurations such as deduplication, compression, erasure coding, and encryption.

Select a storage domain from the list or click **Add Storage Domain** and create a new storage domain. During a Protection Group Run, the Cohesity cluster stores the snapshots of the S3 View in this storage domain.

Bucket Policy

A bucket policy provides access to the objects stored in the bucket. It's important to note that bucket policies do not apply to objects owned by other accounts.

To add a bucket policy, click **Add Policy** and input the desired bucket policy in the provided text area.

The following is an example of the bucket policy:

```
{
"Version": "2012-10-17",
"Statement": [
{
"Sid": "AllowUser2ToWriteObject",
"Effect": "Allow",
"Principal": {
"CanonicalUser": [
"USER2@LOCAL@"
]
},
```

```
"Action": [
"s3:PutObject"
],
"Resource": "arn:aws:s3:::BUCKET-NAME/*"
},
{
"Sid": "AllowUser2ToGetObjectOnlyWhenTagsHavePublicSecurity",
"Effect": "Allow",
"Principal": {
"CanonicalUser": [
"USER2@LOCAL@"
]
},
"Action": [
"s3:GetObject"
],
"Resource": "arn:aws:s3:::BUCKET-NAME/*",
"Condition": {
"StringEquals": {
"s3:ExistingObjectTag/security": "public"
}
}
}
]
}
```

After adding a policy, you can click **Remove Policy** if you wish to delete a policy.

Bucket Owner

A bucket owner has certain permissions and control over the bucket and its contents. You can only add a local Cohesity user as a bucket owner. By default, the currently logged-in user is set as the bucket owner. You can designate a local user as the bucket owner by selecting the user from the User drop-down list.

> **Note:** You cannot change the bucket owner once the S3 View is created.

ACLs

S3 Access Control Lists (ACLs) allow you to regulate access to both buckets and objects. ACLs are applied to the buckets. ACLs determine the users or groups that are granted access. ACL permissions can only be granted to a local Cohesity user or one of the following S3 predefined groups:

- Authenticated Users group
- All Users group
- Log Delivery group

To add an ACL:

1. Click **Add ACL**.
2. Select a local Cohesity user or an S3 predefined group.
3. Select the bucket ACLs as necessary.
4. Click **Add**.

You can also edit or delete ACLs. To do so, click the vertical ellipsis corresponding to the user or group you wish to modify or remove and click **Edit** or **Delete** accordingly.

Object Keys

Cohesity S3 stores data based on Object Key patterns. When creating an S3 View, consider the object keys that you are most likely to store in the View, and choose the best pattern for optimal performance. You cannot edit object keys after the View is created.

Object Key Pattern—Select one of the following:

- **Object ID**: Object keys are mapped to internal object IDs.

  Important:
  - By default, **Object ID** is selected as the **Object Key Pattern**.
  - Cohesity recommends using Object ID as the Object Key Pattern to obtain better performance and scalability for the following operations:
    - `PutObject`
    - `ListObjects`
    - `ListObjectVersions`
  - Cohesity recommends using the Object ID as the Object Key Pattern for S3 views that are used as external targets.
  - If the **Object Key Pattern** of the S3 View is **Object ID**:
    - You cannot access the S3 View through SMB or NFS protocol.
    - Indexing is not supported and you cannot perform file-level recovery.
    - You can replicate the S3 View only if the remote cluster is running Cohesity cluster version 7.0. If the remote cluster is running an earlier version, replication fails.
    - You must use the **Incremental Forever** archival format while archiving the S3 View. You cannot archive the S3 View using the **Incremental with Periodic Full** format.
    - You can archive the S3 View to an external target that is registered on a Cohesity cluster running a version prior to 7.0. However, you can restore the S3 View only if the external target is registered on a Cohesity cluster that runs 7.0. If the external target is registered on a Cohesity cluster running a version prior to 7.0, restore fails.

- **Flat**: Optimal if most object keys are not organized hierarchically.

  If you create an S3 View using S3 API, by default, a Flat > Structured > Short View is created.

- **Flat Key Type**: Select one of the following:
  - **Random**: Optimal for object keys that are auto-generated, for example, MD5 hashes or UUIDs.
  - **Structured**: Optimal for object keys that use some structured pattern, like a date/time stamp. If most keys are less than 30 characters long, select Short. Otherwise, select Long.

  The Flat key pattern is supported on 6.4 and later Cohesity clusters. You cannot replicate the View to, or restore it from an archive on, a pre-6.4 Cohesity cluster.
- **Hierarchical**: Optimal if most object keys are organized in a logical hierarchy using '/' as a separator. Cohesity does not recommend using the **Hierarchical** View for latency-sensitive workloads.

> **Note:** Even though you choose a particular pattern, you can still store objects that have keys with a different pattern.

Read/Write Protocol

Select one of the following protocols:

- If you need read/write access to the View, select **S3**.
- Select **SWIFT** to create a Swift-only View that provides storage space for authorized Keystone users. The Swift-only View is tagged to a Project within a Domain. Only authorized Keystone users can access Swift-only Views. For more information, see Create a Swift-Only View.

> **Note:** If you select SWIFT, you cannot use **Object ID** as the **Object Key Pattern**. You can use **Flat** or **Hierarchical**.

Versioning

Click the **Versioning** section and turn on the toggle to keep multiple variants of an object in the same bucket. You can use versioning to preserve, retrieve, and restore every version of every object stored in your S3 bucket.

After you enable S3 Object Lock, the **Versioning** toggle is automatically turned on and you cannot disable Object Lock or suspend versioning on the bucket.

S3 Object Lock

Before proceeding, review the S3 Object Lock section.

1. Click the **S3 Object Lock** section and turn on the toggle:

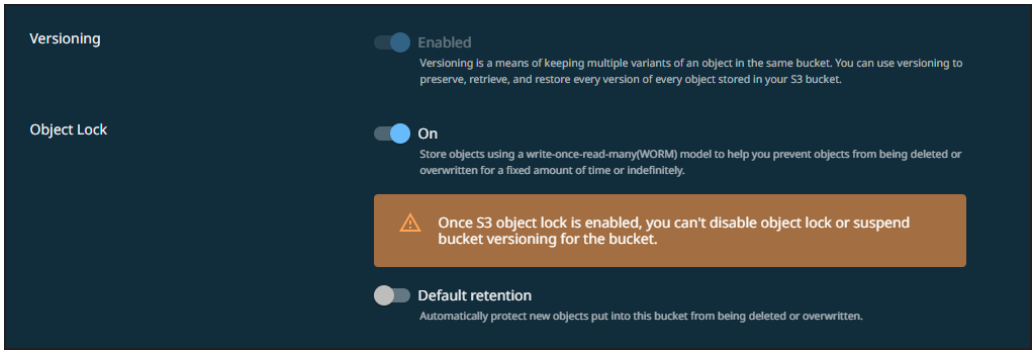

> **Note:** After you enable S3 Object Lock, the **Versioning** toggle is automatically turned on and you cannot disable Object Lock or suspend versioning on the bucket.

2. Turn on the **Default Retention** toggle and select one of the following retention modes to automatically protect new objects put into this bucket from being deleted or overwritten:
    1. **Governance**—Users with specific permissions can overwrite or delete protected object versions during the retention period.
    2. **Compliance**—No user can overwrite or delete protected object versions during the retention period.

    Select either **Days** or **Years** from the drop-down menu and enter the number of days or years to set the retention:

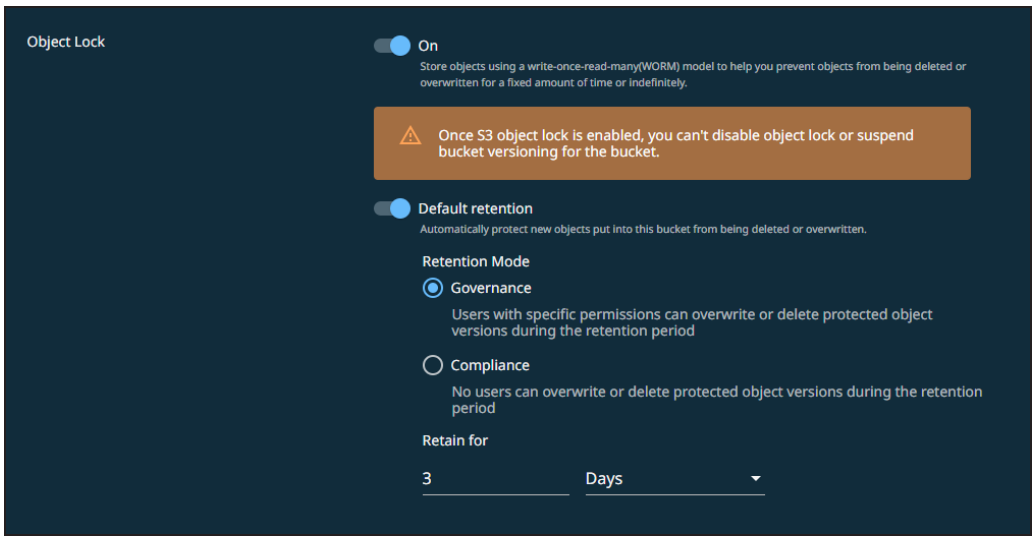

Lifecycle Management Rules

Before proceeding, review the S3 Lifecycle section.

1. Click the **Lifecycle Management Rules** section and click **Add Rule**.
2. In the **Add Lifecycle Management Rule** dialog, set up the rule.
3. Enter the name of the rule.
4. Select the **Rule Scope**:
    - Click **Apply to all objects in the bucket** to apply this rule to all objects in the bucket.
    - Click **Limit the scope of this rule using one or more filters** to limit the scope of this rule by adding one or more filters:
        - Enter a prefix to limit the scope of this rule to a single prefix.
        - Optionally, click **Add Object Tag** and add a key and value to apply the Lifecycle Management rule to a subset of objects that has all the tags specified in the rule. If an object has additional tags specified, the rule is still applied. For more information about object tags, see S3 Object Tags.
5. Select at least one rule action:
    - Select **Expire current versions of objects** and specify the number of days to expire objects identified in the rule.
    - Select **Permanently delete noncurrent versions of objects** and specify the number of days. Specify the number of days an object remains noncurrent before Cohesity executes the deletion process. Cohesity calculates this duration based on the time elapsed since the creation of its successor version, determining the object's noncurrent status.
    - Select **Remove expired object delete markers** to remove delete markers of expired objects. Cohesity will remove a delete marker only if there are no noncurrent versions associated with it.
    - Select **Delete incomplete multipart uploads** and specify the number of days to expire incomplete multipart uploads identified in the rule. Parts associated with these multipart uploads are deleted.
6. Click **Add**.

Protection

You can back up and recover S3 Views using local snapshots. Turn on the **Protection** toggle to specify the protection group and protection policy. If enabled, this S3 View uses this protection group and adheres to the selected protection policy. In the policy, you can also define how the objects should be replicated.

- **Protection Group**—Use this option to protect objects as a group, or individually for granular level control. You can either create a new protection group or select an existing protection group from the drop-down list.

  - **New Group**—Enter a new group name. This creates a new protection group when the View is created. Ensure that you specify a group name that is not in use.
  - **Existing Group**—Select a protection group from the drop-down list.
- **Storage Domain**—Select a storage domain to be used during protection. You can also create a new storage domain.
- **Protection Policy**—Select a policy from the drop-down list. The policy is read-only. If you want to update the policy, go to the **Data Protection > Policies** page to edit the policy.

Case Sensitive File or Folder Names

Turn on the **Case Sensitive File or Folder Names** toggle if you want the uppercase and lowercase versions of the same letter in a file or folder name to be treated differently on the View. For example, **A** is different from **a**.

> **Note:** You cannot edit this setting after you create the View.

Performance

- **QoS Policy**—Select the QoS policy. A QoS policy determines to which storage media data is written and the priority of I/O when contention occurs. For more information, see QoS Policies for Cohesity Views.
- **Pin View to SSD**—Turn on the toggle to pin the View to the SSD tier so that critical data and data that have to be retrieved quickly are not down-tiered to HDD or cloud. For more information, see Pin Views to SSD.

Security

Configure the security settings for the View. You can add and manage IP Allowlist for the S3 View:

1. Select **Override Global IP Allowlist** to ignore the global IP Allowlist and honor only specific entries provided in the View. Select **Extend Global IP Allowlist** to honor the global IP Allowlist and the specific entries provided in the View.
2. Click **Add**.
3. In the **Add Allowlist** dialog, enter the subnet in CIDR format.
4. Select **S3 Permissions**.
5. Optionally, you can enter a description for the IP Allowlist of the View.

Dedupe & Compression

Turn on the **Dedupe & Compression** toggle to always use post-processing, irrespective of the storage domain policy. This option is available only if deduplication and compression are enabled on the storage domain.

Logical Quota

Select one of the options to set the **Logical Quota** for the View:

- Inherit from storage domain
- Select **Override**, turn on the following options, and set the limit:
  - Logical Quota
  - Alert Threshold

> **Note:** Alert threshold must be less than logical quota.

Alert threshold notifies the System Administrator in advance and prevents system capacities from being exhausted without notification.

File Filtering

Turn on the **File Filtering** toggle if you want to block or allow files with specified file extensions from being written to this View. Select one of the following file filtering methods:

- **Denied File Types**—If you want to specify a list of file extensions for files that are blocked from being written to this View. For example, if you specify a list with the .exe and .dll extensions, files with these extensions cannot be written to this View.
- **Allowed File Types**—If you want to specify a list of file extensions for files that are allowed to be written to this View. For example, if you specify a list with the .cvs and .txt extensions, only files with these extensions can be written to this View.

  Enter file extensions separated with commas, spaces, tabs, or newlines, and click **Add**.

Audit Logs

Turn on the **Audit Logs** toggle to enable audit logs. For more information, see File Services Audit Logs.

S3 Server Logs

Toggle On the **Enable S3 Server Logging** option to configure the S3 Server Logs. This option allows you to set the target S3 View that acts as the logging bucket for source S3 Views.

1. Select the required **Target Bucket** on which the logs must be stored.
2. Optionally, you can set the **Target Prefix**.
3. Select the **Log Key Format** for the logging bucket. The options include **Simple Prefix**, **Partitioned - Event Time**, and **Partitioned - Delivery Time**. For more information, see TargetObjectKeyFormat and PartitionedPrefix.

4. You can also add **Log Object ACLs**. For more information, see the **ACLs** section in S3 View Options.

Description

Click **Description** and provide a description for the S3 View.

Manage Cohesity S3 Access Keys

Cohesity automatically generates Cohesity S3 access keys for all users created by using the Cohesity Dashboard. Users can see their own keys, generate a new key and copy it to the clipboard as described below.

Admin users who have the **Manage Users** privilege can generate keys for other users but cannot see or copy those keys to the clipboard. Note that when generating a new key, anyone using the old key will lose access until they retrieve and use the newly generated key.

Get the Access Key Using the Cohesity Dashboard

1. Navigate to **Settings > Access Management**.
2. Click your user name in the list of users and groups.
3. Optionally, click **Show Key** to display your key.
4. Click the copy icon next to the **Access Key ID** and paste the key from your clipboard to a secure location to use later.
5. Click the copy icon next to the **Secret Access Key** and paste it to a secure location to use later.

You can also view the **Status** and **Expiration Date** of the key.

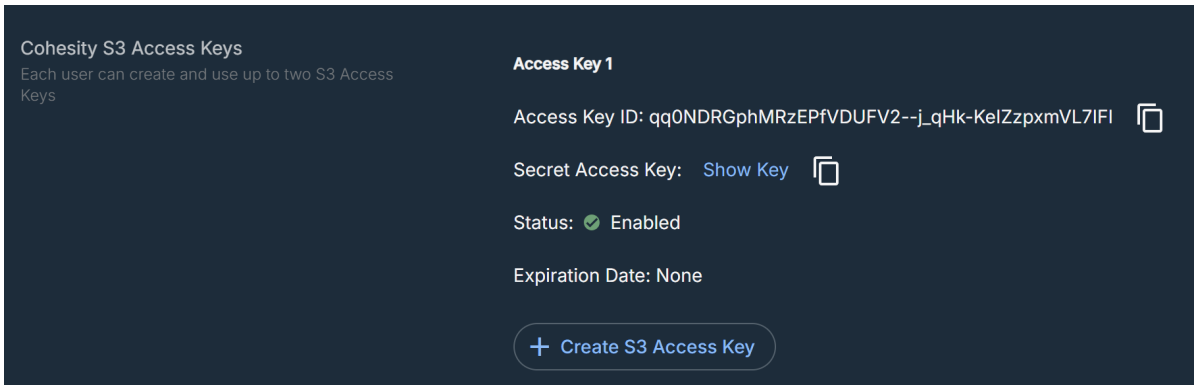

6. You can create two S3 Access Key pairs for each user. Both key pairs function identically and remain valid for authentication until they are disabled, expired, or deleted.

   To add a new access key:

1. Click the **+ Create S3 Access Key** button.
2. In the **Create S3 Access Key** dialog box, optionally toggle On the **Set Expiration** option and set the expiration date and time.
3. Click **Create S3 Access Key**.

To rotate the secret key:

1. Click the action menu (⋮) next to the access key name.
2. Click **Rotate Secret Key**.

To edit the S3 access key:

1. Click the action menu (⋮) next to the access key name.
2. Click **Edit S3 Access Key**.
3. In the **Edit S3 Access Key** dialog box, toggle On or Off the **Enabled** option to enable or disable the access key and toggle On or Off the **Set Expiration** option and set the expiration date and time.

   > **Note:** Once set, expiration time can be shortened but cannot be disabled or extended.

4. Click **Edit S3 Access Key**.

To delete the S3 access key:

1. Click the action menu (⋮) next to the access key name.
2. Click **Delete**.
3. In the **Delete S3 Access Key** dialog box, type **YES**. Once deleted, the key will no longer be valid for authentication.

   > **Note:** You cannot delete both the S3 access keys; at least one must be available. It can be disabled or expired but cannot be deleted.

4. Click **Delete S3 Access Key**.

Get the Access Key Using the Cohesity CLI

You can also get the access key information by using the following command in the Cohesity CLI:

```
admin>user ls
USER NAME                    : admin
ROLES                        : COHESITY_ADMIN
DOMAIN                       : LOCAL
CREATION TIME                : Sunday, 26-Mar-17 20:00:39 PDT
LAST UPDATED TIME            : Thursday, 30-Mar-17 09:37:13 PDT
```

```
EFFECTIVE START TIME          : Sunday, 26-Mar-17 20:00:39 PDT
EMAIL ADDRESS                 : admin@abc.com
S3 ACCOUNT ID                 : admin@LOCAL
S3 ACCESS KEY ID              : yQ5iFakef0exyvDRJGRuF0bfd/sIMag50YiTb4UIrfU=
S3 SECRET ACCESS KEY          : T+VIuvkdltrgFakes9PjHrgaaODMF5rJjnt3452113o=
```

Allow S3 Buckets on the Cohesity cluster

> **Note:** This task is required only if you will be using the S3 API to create S3 views on the Cohesity cluster. This task is not necessary if you use the Cohesity Dashboard to create the views.

Create or edit a storage domain on the Cohesity cluster that allows S3 buckets. Even though S3 clients are not aware of storage domains, views in the Cohesity cluster must be created within a storage domain.

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Create a storage domain that allows S3 buckets by using the Cohesity CLI `sd create` command:

   ```
   admin>sd create name=S3StorageDomain s3-buckets-allowed=true
   ```

> **Note:**
> It is mandatory to mention either the cluster-partition-id or the cluster-partition-name in the storage domain creation command.
>
> - Example command for storage domain creation using cluster-partition-id:
>   admin@127.0.0.1> sd create name=S3StorageDomain s3-buckets-allowed=true cluster-partition-id=*32*
> - Example command for storage domain creation using cluster-partition-name:
>   admin@127.0.0.1> sd create name=S3StorageDomain s3-buckets-allowed=true cluster-partition-name=*DefaultPartition*

Or edit an existing storage domain to allow S3 buckets using the `sd edit` command:

```
admin>sd edit name=S3StorageDomain s3-buckets-allowed=true
```

> **Note:** If multiple storage domains on the Cohesity cluster allow S3 buckets, Cohesity creates an S3 bucket in a random storage domain.

3. Create a view inside a storage domain that allows S3 buckets by using the Cohesity CLI `view create` command:

   ```
   admin>view create name=<viewName> storage-domain-
   name=<StorageDomainName>
   ```

   By default, all views are created with support for all protocols, but S3 access is read-only. If you want to disable NFS and SMB access to a view, use this command:

   ```
   admin>view create name=<viewName> storage-domain-
   name=<StorageDomainName> protocol=S3 view-category=ObjectServices
   ```

   Creating a view with `protocol=S3` means the view does not support NFS or SMB access.

Use the AWS CLI on Cohesity S3

You can use the AWS CLI to interact with buckets (S3 protocol views) on the Cohesity cluster.

1. Download and install AWS CLI.
2. Refer to the S3 commands set.

Cohesity cluster S3 support is detailed here.

3. Run the following command:

   ```
   aws configure
   ```

   and provide the access key and secret access key.

4. For the endpoint, use <Hostname or IP Address>:3000

   where Hostname or IP Address is available from a Cluster > Networking > VIPS page in the Cohesity Dashboard.

   For example: myCohesity.myCompany.com:3000

5. Cohesity supports both Signature Version 2 and Version 4, but recommends Version 4. Amazon has deprecated Version 2.

   The AWS CLI uses Version 4 by default for S3. To use Version 2, run the following command to configure it:

   ```
   aws configure set default.s3.signature_version s3
   ```

   To switch back to Version 4, run the following command:

   ```
   aws configure set default.s3.signature_version s3v4
   ```

Commands to List Buckets and Objects

> **Tip:** The commands below assume the Cohesity cluster has a valid SSL certificate installed. If the Cohesity cluster has the default self-signed certificate, add the `--no-verify-ssl` option to the commands.

Sample command to list buckets:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000  s3api list-buckets
```

Sample command to list objects in a bucket:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000  s3api list-objects-v2 --bucket=<bucket_name>
```

Manage S3 Server Certificate

The Cohesity S3 server shares the certificate with the Cohesity cluster. The default certificate is auto-generated and self-signed. To get the SSL certificate details of the cluster, execute the following command:

```
iris_cli cluster list-ssl-cert-details
```

For more information about Cohesity CLI, see Using the Cohesity CLI.

For more information about managing SSL certificates and the procedure to update the default certificate, see Update SSL Certificates Using Cohesity CA or Customer CA certificate for Agent to Cluster Communication.

Cohesity recommends using the AWS CLI to access the Cohesity S3 server without disabling certificate verification. Use the `--ca-bundle` option to specify the path of the certificate file in your host. The following is an example:

```
aws s3api put-object --bucket hello --endpoint-url https://example.com:3000
--key tmp --body myfile --profile 10.2.195.1 --ca-bundle cabundle.pem
```

For more information, see the *AWS CLI Configuration Variables* section in the AWS CLI Command Reference Guide.

If you have replaced the default certificate and if the new certificate has a certificate chain, you must combine all the certificates in this certificate hierarchy into one file and specify the path of this file in the `--ca-bundle` option.

> **Note:**
> The content in the combined certificate should be in the following order:
>
> 1. Server certificate
> 2. Intermediate certificates
> 3. Root certificate

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

Upgrade S3 Views to the Object ID Object Key Pattern

You can upgrade an S3 View to the **Object ID** object key pattern. Cohesity recommends using Object ID as the Object Key Pattern to obtain better performance and scalability for the following operations:

- PutObject
- ListObjects
- ListObjectVersions

Considerations

- Upgrading an S3 view to the **Object ID** object key pattern is irreversible.
- Replication of S3 views is not supported if the target cluster is running an older version.
- The upgrade process does not verify the archival policy beforehand. S3 views using **Object ID** as the object key pattern cannot be archived using the **Incremental with Periodic Full** format. As a result, archival jobs will fail for any existing S3 views using this format. Cohesity recommends creating new archival jobs with the **Incremental Forever** archival format before starting the upgrade.
- The S3 view upgrade process occurs in batches (a set of views) and is carried out in the background. The duration of the process depends on the availability of cluster resources.
- The migration may take longer than expected if:
  - The S3 view was not created according to the recommended guidelines. For example, the View may have been set up with a **Hierarchical** object key pattern, but the object keys are not organized hierarchically.
  - The S3 view includes FAT directories, each containing a significantly high number of objects at the same level.

Upgrade S3 Views

To upgrade S3 views:

1. Navigate to **SmartFiles > Views**.
2. Go to the **Upgrade** tab.
3. Check the upgrade status of the view(s) in the **Upgrade Status** column. If the upgrade status of one or more views is **Eligible for upgrade**, you can select multiple views at once and click **Enable Upgrade**:

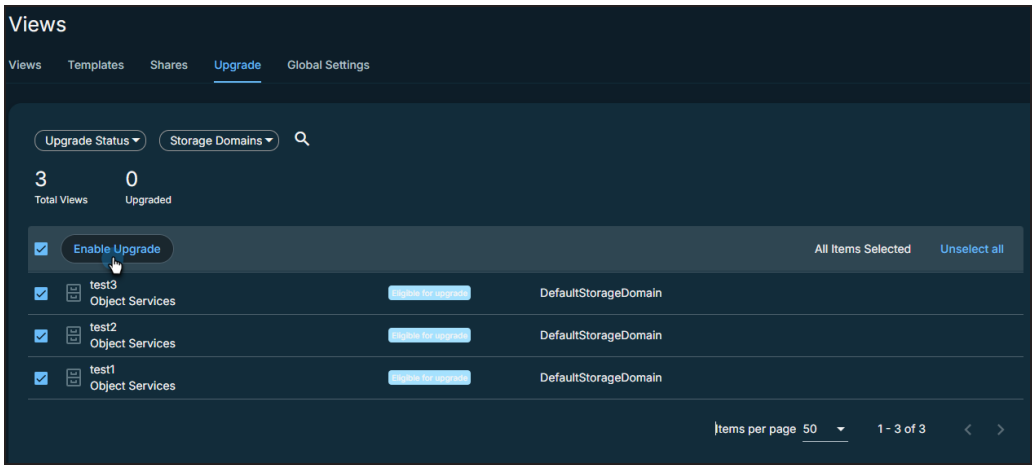

4. In the **Enable Upgrade** dialog, enter **YES** and click **Enable Upgrade**.

The upgrade process is initiated and the real-time progress is displayed in the progress bar.

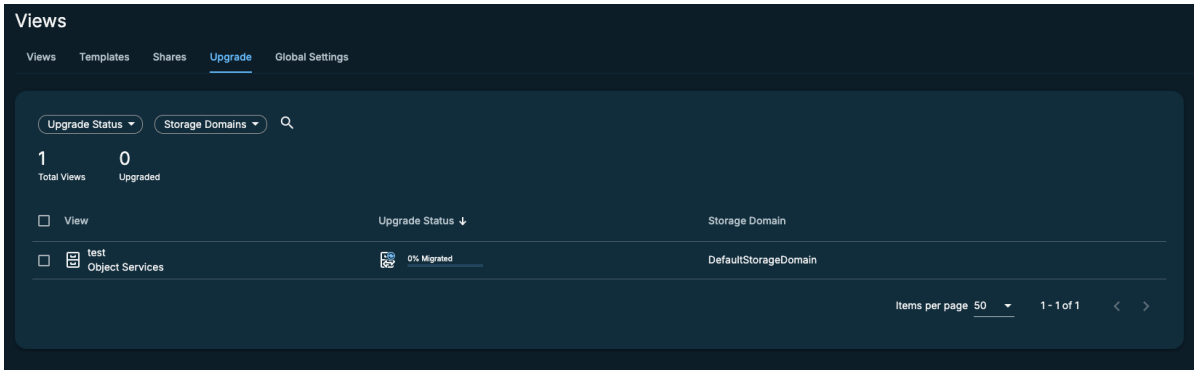

Once the background job starts the upgrade, the status changes to **Upgrade in progress**. You can pause the job temporarily and resume it later. However, you cannot cancel the upgrade after the job starts.

After the upgrade is complete, the status changes to **Upgrade Completed**:

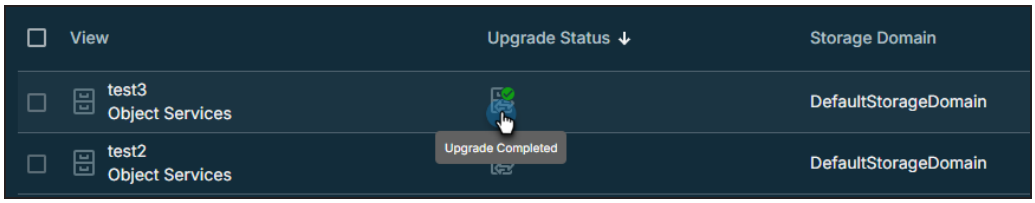

To check whether the upgrade is successful:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page, hover over the upgraded view, click the ⋮ icon, and select **Edit**.
3. In the **Object Keys** section, the **Object Key Pattern** for the view should be **Object ID**:

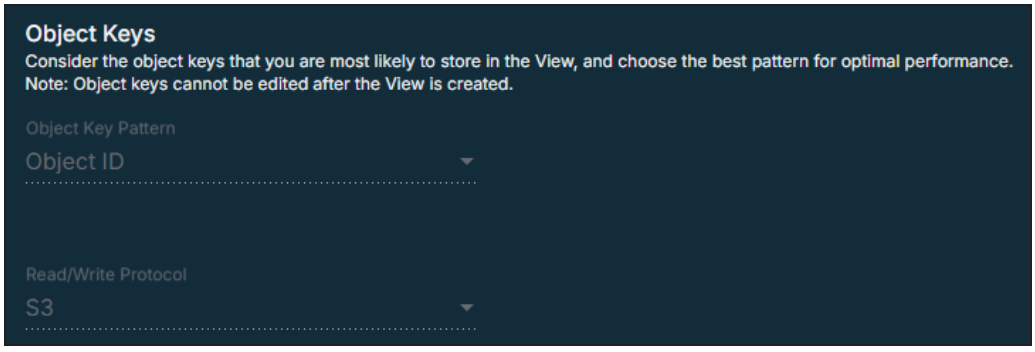

Supported S3 APIs

Cohesity supports the following S3 APIs:

API	Required Permissions
AbortMultipartUpload	s3:AbortMultipartUpload
CompleteMultipartUpload	s3:PutObject
CopyObject	• s3:PutObject • s3:GetObject • s3:GetObjectVersion
CreateBucket	The user must have the **Manage Cohesity Views** privilege assigned to their role to perform this operation. No additional S3 permissions are required.
CreateMultipartUpload	s3:PutObject
DeleteBucket	s3:DeleteBucket
DeleteBucketLifecycle	s3:DeleteBucketLifecycleConfiguration
DeleteBucketOwnershipControls	s3:PutBucketOwnershipControls
DeleteBucketPolicy	s3:DeleteBucketPolicy
DeleteObject	• s3:BypassGovernanceRetention • s3:DeleteObject • s3:DeleteObjectVersion
DeleteObjects	• s3:BypassGovernanceRetention • s3:DeleteObject • s3:DeleteObjectVersion
DeleteObjectTagging	• s3:DeleteObjectTagging • s3:DeleteObjectVersionTagging
GetBucketAcl	s3:GetBucketAcl
GetBucketLifecycleConfiguration	s3:GetLifecycleConfiguration

API	Required Permissions
GetBucketOwnershipControls	s3:GetBucketOwnershipControls
GetBucketPolicy	s3:GetBucketPolicy
GetBucketVersioning	s3:GetBucketVersioning
GetObject	• s3:GetObject • s3:GetObjectVersion • s3:GetObjectTagging • s3:GetObjectVersionTagging • s3:GetObjectLegalHold • s3:GetObjectRetention
GetObjectAcl	• s3:GetObjectAcl • s3:GetObjectVersionAcl
GetObjectLegalHold	s3:GetObjectLegalHold
GetObjectLockConfiguration	s3:GetBucketObjectLockConfiguration
GetObjectRetention	s3:GetObjectRetention
GetObjectTagging	• s3:GetObjectTagging • s3:GetObjectVersionTagging
HeadBucket	s3:ListBucket
HeadObject	• s3:GetObject • s3:GetObjectVersion
ListBuckets	No additional S3 permissions are required.
ListMultipartUploads	s3:ListBucketMultipartUploads
ListObjectsV2	s3:ListBucket
ListObjectVersions	s3:ListBucketVersions

API	Required Permissions
ListParts	s3:ListMultipartUploadParts
PutBucketAcl	s3:PutBucketAcl
PutBucketLifecycleConfiguration	s3:PutLifecycleConfiguration
PutBucketOwnershipControls	s3:PutBucketOwnershipControls
PutBucketPolicy	s3:PutBucketPolicy
PutBucketVersioning	s3:PutBucketVersioning
PutObject	• s3:PutObject • s3:PutObjectTagging
PutObjectAcl	• s3:PutObjectAcl • s3:PutObjectVersionAcl
PutObjectLegalHold	s3:PutObjectLegalHold
PutObjectLockConfiguration	s3:PutBucketObjectLockConfiguration
PutObjectRetention	s3:PutObjectRetention
PutObjectTagging	• s3:PutObjectTagging • s3:PutObjectVersionTagging
UploadPart	s3:PutObject
UploadPartCopy	• s3:PutObject • s3:GetObject • s3:GetObjectVersion

Supported S3 API Headers

This page provides a list of the S3 headers that are supported by Cohesity.

> **Note:** If unsupported headers are included in a request, they will be ignored.

Request Headers

Cohesity supports the following request headers:

- Authorization
- Content-Length
- Content-MD5
- Date
- Expect
- Host
- Range
- User-Agent—Ignored
- x-amz-bucket-object-lock-enabled
- x-amz-bypass-governance-retention
- x-amz-content-sha256
- x-amz-copy-source
- x-amz-copy-source-if-match
- x-amz-copy-source-if-modified-since
- x-amz-copy-source-if-none-match
- x-amz-copy-source-if-unmodified-since
- x-amz-copy-source-range
- x-amz-date
- x-amz-meta-*
- x-amz-object-lock-legal-hold
- x-amz-object-lock-mode
- x-amz-object-lock-retain-until-date
- x-amz-security-token—Ignored
- x-amz-tagging
- x-amz-tagging-directive
- x-amz-checksum-<algo> (<algo> can be crc32, crc32c, crc64vme, sha1, sha256)
- x-amz-checksum-type
- x-amz-checksum-algorithm
- x-amz-sdk-checksum-algorithm
- x-amz-trailer
- x-amz-decoded-content-length

Response Headers

Cohesity supports the following response headers:

- Connection
- Content-Length
- Content-MD5
- Date
- ETag
- Server
- x-amz-abort-date
- x-amz-abort-rule-id
- x-amz-copy-source-version-id
- x-amz-delete-marker
- x-amz-expiration
- x-amz-object-lock-legal-hold
- x-amz-object-lock-mode
- x-amz-object-lock-retain-until-date
- x-amz-tagging-count
- x-amz-version-id
- x-amz-checksum-<algo> (<algo> can be crc32, crc32c, crc64vme, sha1, sha256)
- x-amz-checksum-type

Chunk Upload Format

The following formats of chunked uploads are supported at Cohesity S3. You can specify the format to be used by your request by adding the x-amz-content-sha256 header with one of the following values:

Header Value	Description
Actual payload checksum value	The actual checksum of your object used only when uploading data in a single chunk.
UNSIGNED-PAYLOAD	Used when uploading the object as a single unsigned chunk.
STREAMING-UNSIGNED-PAYLOAD-TRAILER	Used when sending an unsigned payload over multiple chunks. A trailing header is included after the chunk is uploaded.

Header Value	Description
STREAMING-AWS4-HMAC-SHA256-PAYLOAD	Used when sending a payload over multiple chunks that are signed using AWS4-HMAC-SHA256. This produces a SigV4 signature.
STREAMING-AWS4-HMAC-SHA256-PAYLOAD-TRAILER	Used when sending a payload over multiple chunks that are signed using AWS4-HMAC-SHA256. This produces a SigV4 signature. In addition, the digest for the chunks is included as a trailing header.

By default, the latest versions of AWS S3 SDK send S3 PutObject and S3 UploadPart request with STREAMING-UNSIGNED-PAYLOAD-TRAILER header value.

> **Note:** The header values, STREAMING-AWS4-ECDSA-P256-SHA256-PAYLOAD and STREAMING-AWS4-ECDSA-P256-SHA256-PAYLOAD-TRAILER are not supported.

Considerations

- When an algorithm required by the server is not specified in a signature V4 request, by default the CRC64NVME algorithm is used.
- For a multipart upload request, if the multi part is initiated (CreateMulitpartUpload) with an algorithm, the following UploadPart requests must also use the same algorithm, else the UploadPart request fails.
- When an algorithm is not specified during the multi part initiation in a signature V4 request, by default the CRC64VME algorithm is used. The following UploadPart requests must also use the same algorithm.
- For the PutObject and UploadPart S3 requests using signature V4, both the client and server compute checksum for the object. Client SDK also chunks the data before sending it to server. This can cause a minor performance impact.
- By default, the latest SDK client sends request in chunked upload. To avoid chunked upload, set the configuration, "request_checksum_calculation = WHEN_REQUIRED" or equivalent field in your S3 client.

  > **Note:** This can make the data integrity checks of your request less strict.

- By default, the latest AWS S3 SDK clients include a header, x-amz-checksum-mode:ENABLED, in the GetObject request. In such cases, the server may send checksum of the object, and the latest S3 SDK clients compute the checksum received. This is a client-side latency that can slow down the GetObject responses.

  As a workaround, you can perform either of the following:

- Avoid setting the header to x-amz-checksum-mode.
- In your S3 client configuration, set "response_checksum_validation = WHEN_REQUIRED" or an equivalent field. This ensures that the SDK client does not include x-amz-checksum-mode:ENABLED in the header by default.

> **Note:** This can make the data integrity checks of your request less strict.

S3 Object Lock

On a Cohesity cluster, you can create a Cohesity view and use it as an Amazon S3 bucket. S3-compatible clients can use the view for object storage. S3 Object Lock enables you to store objects using a write-once-read-many (WORM) model. Cohesity supports S3 Object Lock, so you can prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely.

You can enable Object Lock on new or existing SmartFiles S3 views. Cohesity recommends that you enable Object Lock when you create a new view. If you enable Object Lock on an existing SmartFiles S3 view, Object Lock does not implicitly protect objects that were created before the lock was enabled on the view. To protect objects that were created before the lock was enabled, use the `PutObjectRetention` or `PutObjectLegalHold` APIs and explicitly enable Object Lock on these objects.

> **Note:** If you enable Object Lock on a view, versioning is automatically enabled for the view and you cannot suspend versioning.

Object Lock uses versioning so multiple variants of an object are available in the same bucket. Even if the bucket owner deletes an object in a versioning-enabled bucket, the object version is protected and you can restore any version of the object stored in the bucket. From a security perspective, Object Lock provides an extra layer of protection against object version changes, inadvertent deletion, and malicious modification of data.

In AWS, a policy defines permissions. You can associate a policy with an identity or resource. To restrict access to SmartFiles S3 views, you can attach resource-based policies so that you can specify who has access and what actions they can perform.

For more information about Object Lock, versioning, and policies, see Amazon S3 User Guide.

Retention Period

A retention period is a fixed period of time during which an object remains locked. During this period, the object is WORM-protected and cannot be overwritten or deleted. To add a retention period:

- Object Lock must be enabled on the bucket that contains the object
- The user must have the `s3:PutObjectRetention` permission

Retention Modes

Object Lock provides two retention modes: governance and compliance. Consider the following points before you choose the retention mode:

- Governance mode—You cannot overwrite or delete an object version or alter its lock settings unless you have special permissions. The bucket owner or a user granted permissions through S3 Bucket Policy can modify the retention period or delete the object using the `bypass-governance-retention` parameter in the request.

  > **Note:**
  > - When a user with adequate permissions executes a `DeleteObject` operation in a versioning-enabled bucket, a unique version delete marker is generated, safeguarding the previous version. To permanently delete a particular version, you can invoke the `DeleteObject` operation and provide the specific `VersionId`.
  > - To view all versions of an object, execute the `ListObjectVersions` operation. The response includes all accessible versions within the S3 system.
  >
  >   For more information about S3 versioning and deleting object versions, see Amazon S3 User Guide.

- Compliance mode—No user can overwrite or delete a protected object version. If an object is locked in compliance mode, you cannot change the retention mode or shorten the retention period.

Legal Hold

A legal hold prevents an object version from being overwritten or deleted. The legal hold provides the same protection as a retention period, but it has no expiration date. The legal hold remains in place until you explicitly remove it. To place or remove legal holds:

- Object Lock must be enabled on the bucket that contains the object
- The user must have the `s3:PutObjectLegalHold` permission

Important Points to Note

- If you enable Object Lock on a view, you can delete the view only if the view does not contain any objects.

- If you enable Object Lock, you cannot disable Object Lock or suspend versioning on the bucket.
- You can customize the Object Lock when cloning S3 views, using the clone API. Disabling Object Lock allows you to delete objects in the cloned view even if Object Lock is enabled in the source view.
  For information on the API, contact your Cohesity account team.

Support for New Amazon S3 APIs

Cohesity introduces support for the following Amazon S3 APIs:

- GetObjectLockConfiguration
- PutObjectLockConfiguration
- GetObjectLegalHold
- PutObjectLegalHold
- PutObjectRetention
- GetObjectRetention

Cohesity enhances support for the following Amazon S3 APIs:

- CreateBucket—Support for the `x-amz-bucket-object-lock-enabled` request header.
- PutBucketVersioning
- PutObject—Support for the following request headers:
  - `x-amz-object-lock-mode`
  - `x-amz-object-lock-retain-until-date`
  - `x-amz-object-lock-legal-hold`
- GetObject—Support for the `VersionId` parameter and the following response headers:
  - `x-amz-object-lock-mode`
  - `x-amz-object-lock-retain-until-date`
  - `x-amz-object-lock-legal-hold`
- HeadObject—Support for the `VersionId` parameter and the following response headers:
  - `x-amz-object-lock-mode`
  - `x-amz-object-lock-retain-until-date`
  - `x-amz-object-lock-legal-hold`
- DeleteObject—Support for the `x-amz-bypass-governance-retention` request header.

- CreateMultipartUpload—Support for the following request headers:
  - `x-amz-object-lock-mode`
  - `x-amz-object-lock-retain-until-date`
  - `x-amz-object-lock-legal-hold`

  For more information about these APIs, see Amazon S3 API Reference.

Server Access Log for Cohesity S3 Views

Server access logging provides records in the AWS S3 access log format for the requests to Cohesity S3 Views. This functionality can be accessed using the AWS **get-bucket-logging** and **put-bucket-logging** APIs.

To configure access logging on the source bucket:

1. Setup a target bucket where the logs will be published.
2. Create a bucket policy on the target bucket allowing logging service to publish log objects.

   A bucket policy must be set to enable writes on the target bucket. This can be set using the put-bucket-policy command.

   Sample Bucket Policy

   ```
   {
   "Id":"abc",
   "Statement":
   [{
   "Action":["s3:PutObject"],
   "Effect":"Allow",
   "Principal":{"Service":"logging.s3.cohesity.com"},
   "Resource":["urn:csf:s3:::<view>/*"],
   "Sid":"S3ServerAccessLogsPolicy"
   }],
   "Version":"2012-10-17"
   }
   ```

3. Create a bucket logging configuration on the source bucket using the put-bucket-logging command.

   Sample Bucket Logging Configuration

   ```
   {
   ```

```
"LoggingEnabled":
{
"TargetBucket": "<view>",
"TargetPrefix": "logs/",
"TargetObjectKeyFormat":
{
"PartitionedPrefix":
{
"PartitionDateSource": "EventTime"
}
}
}
}
```

Enable Object Lock From the Cohesity Cluster GUI

From the Cohesity Cluster GUI, you can enable Object Lock and Versioning on new or existing SmartFiles S3 views. If you enable Object Lock, versioning is automatically enabled for the view and you cannot suspend versioning.

To enable S3 Object Lock:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page:
    - Click **Create View** and proceed to the next step.
    - If it is an existing S3 view, click the Actions menu (⋮) and select **Edit** to edit the view. Proceed to Step 5.
3. Select a predefined template or custom template. Ensure that you select one of the following from the **Object Services** category to create an S3 view:
    - General
    - Containers
    - Hadoop
    - Splunk SmartStore
4. Click **More Options**.
5. Click the **S3 Object Lock** section and turn on the toggle.

> **Note:** After you enable S3 Object Lock, the **Versioning** toggle is automatically turned on and you cannot disable Object Lock or suspend versioning on the bucket.

6. Turn on the **Default Retention** toggle and select one of the following retention modes to automatically protect new objects put into this bucket from being deleted or overwritten:
   1. **Governance**—Users with specific permissions can overwrite or delete protected object versions during the retention period.
   2. **Compliance**—No user can overwrite or delete protected object versions during the retention period.

   Select either **Days** or **Years** from the drop-down menu and enter the number of days or years to set the retention.

For detailed instructions about creating an S3 View, see Create an S3 View.

Enable Object Lock Using AWS CLI

Ensure that the following prerequisites are met:

- Install curl on a Linux host which has connectivity with the Cohesity cluster.
- Install AWS CLI on the Linux host.
- AWS access key ID and secret access key of the user who is creating the bucket.

To enable Object Lock using AWS CLI, execute the following command:

```
aws s3api put-object-lock-configuration --endpoint-url https://cluster-fqdn-or-ip-address:3000 --bucket "bucket-name" --object-lock-configuration '{"ObjectLockEnabled": "Enabled"}' --no-verify-ssl --profile profile-name
```

View Object Lock Status

To determine whether Object Lock is enabled on a view, navigate to **SmartFiles > Views**. On the **Views** page, a lock symbol is displayed next to the view name if Object Lock is enabled.

The following figure shows a view on which Object Lock is enabled:

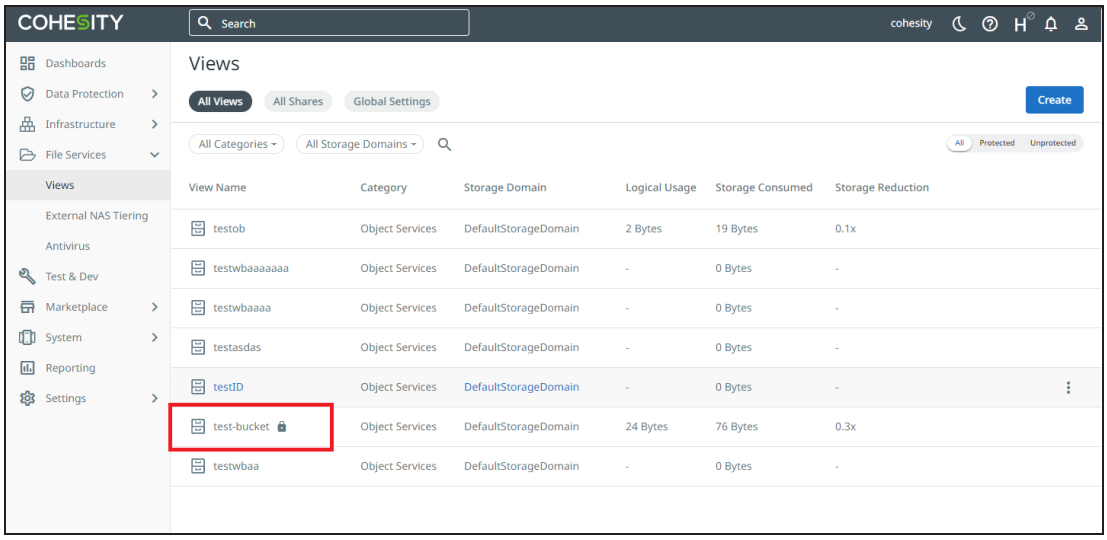

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

S3 Lifecycle

Cohesity supports configuring the lifecycle of your S3 objects so that they are stored cost-effectively throughout their lifecycle. S3 Lifecycle configuration includes a set of rules that define actions applicable to a specified group of objects. You can enable lifecycle management on new or existing S3 Views.

AWS S3 supports two types of actions:

- Transition actions—Define when objects transition to another S3 storage class.

  > **Note:** Cohesity does not support transition actions.

- Expiration actions—Define when objects should expire. Cohesity supports expiration actions and deletes expired objects on your behalf.

You can organize S3 objects using prefixes or tags. A prefix value enables you to group similar objects together in a bucket. You can create multiple S3 Lifecycle configuration rules for a bucket and apply a rule to a subset of objects.

Lifecycle Configuration

The S3 Lifecycle configuration is specified as XML and consists of a set of predefined rules. Each rule can include the following elements:

- **ID**—Uniquely identifies a rule. If you do not specify an ID in the Lifecycle configuration, Cohesity S3 automatically generates a unique identifier.
- **Status**—Indicates whether the rule is enabled or not.
- **Filter**—Identifies a subset of objects on which the rule should be applied. You can filter objects by prefixes or tags. This is an optional element and can be specified if required.
- **Actions**—Specifies a set of actions that must be applied to the filtered objects. Cohesity supports two expiration actions.

  For more information, see Expiration Actions.

The following is a sample lifecycle configuration using AWS CLI:

```
aws s3api put-bucket-lifecycle-configuration --bucket <view_name> --endpoint-url https://<host_name>:3000 --
no-verify-ssl --lifecycle-configuration \
'{
 "Rules":[
   {
     "ID":"id0",
     "Status":"Enabled",
     "Filter":{
       "Prefix":"obj"
     },
     "Expiration":{
       "Days":1
     }
   },
   {
     "ID":"id1",
     "Status":"Enabled",
     "Filter":{
       "Prefix":"mpu"
     },
     "AbortIncompleteMultipartUpload":{
       "DaysAfterInitiation":1
     }
   }
 ]
}'
```

Expiration Actions

Cohesity enables you to define when objects should expire. Cohesity supports the following three expiration actions:

- **Expiration**—Expires the objects identified in the rule. Objects are deleted at midnight UTC.

  Expiration action is performed based on the object versioning state:

  - For **versioned** buckets, Cohesity S3 adds a delete marker and retains the current version of the object as a noncurrent version.
  - For **non-versioned** buckets, Cohesity S3 permanently removes the object. This is not applicable to the **ExpiredObjectDeleteMarker** option.

- For **versioning-suspended** buckets, Cohesity S3 creates a delete marker with null as the version ID. This delete marker replaces any object version with a null version ID in the version hierarchy, which effectively deletes the object.

You can configure one of the following parameters for this action:

- **Days**—Specify the number of days after which the object should expire. Lifetime is measured from the object's creation time. Lifetime is calculated by adding the number of days specified to the object's creation time and rounding off the resulting time to midnight UTC of the next day. The following is an example:
  - Number of days specified—2
  - Object creation time—1/20/2022 11:30 AM UTC

    The object is deleted at midnight UTC on 1/23/2022
- **Date**—Specify the date on which the object should expire.
- **ExpiredObjectDeleteMarker**—Indicates if Cohesity removes a delete marker without any noncurrent versions. When enabled, the delete marker expires; when disabled, no action is taken. You cannot specify this option with Days or Date in a Lifecycle Expiration Policy.

- **NoncurrentVersionExpiration**—Specifies when noncurrent object versions expire. Once it has expired, Cohesity permanently deletes the noncurrent object versions. You can set this lifecycle configuration action on a versioning-enabled View or a versioning-suspended View to delete noncurrent object versions after a set period in the object's lifecycle.
  - **NoncurrentDays**—Specifies the number of days an object is noncurrent. Once this specified period elapses, Cohesity automatically deletes the noncurrent object versions from the View.
- **AbortIncompleteMultipartUpload**—Manages the lifecycle of incomplete multipart uploads.
  - **DaysAfterInitiation**—Specify the number of days after initiation of multipart uploads Cohesity can abort these incomplete multipart uploads.

Support for Upgrade, Archive, and Replication

Expiration actions are also supported on existing Cohesity SmartFiles S3 views. If you have successfully upgraded to a version that supports S3 Lifecycle, you can configure these actions on views created in a previous version, for example, 6.5.1.

If S3 Lifecycle is enabled on a view and if the view is replicated to or archived and restored to another cluster that does not support S3 Lifecycle, the lifecycle configuration does not take effect. After the cluster is upgraded to a version supporting the S3 Lifecycle feature, the lifecycle configuration takes effect.

Add Lifecycle Management Rules From the Cohesity Cluster GUI

From the Cohesity cluster GUI, you can add Lifecycle Management rules on new or existing SmartFiles S3 views.

To add Lifecycle Management rules:

1. Navigate to **SmartFiles > Views**.
2. On the **Views** page:
    - Click **Create View** and proceed to the next step.
    - If it is an existing S3 view, click the Actions menu (⋮) and select **Edit** to edit the view. Proceed to **Step 5**.
3. Select a predefined template or custom template. Ensure that you select one of the following from the **Object Services** category to create an S3 view:
    - General
    - Containers
    - Hadoop
    - Splunk SmartStore
4. Click **More Options**.
5. Click the **Lifecycle Management Rules** section and click **Add Rule**.
6. In the **Add Lifecycle Management Rule** dialog, set up the rule.
7. Enter the name of the rule.
8. Select the **Rule Scope**:
    - Click **Apply to all objects in the bucket** to apply this rule to all objects in the bucket.
    - Click **Limit the scope of this rule using one or more filters** to limit the scope of this rule by adding one or more filters:
        - Enter a prefix to limit the scope of this rule to a single prefix.
        - Optionally, click **Add Object Tag** and add a key and value to apply the Lifecycle Management rule to a subset of objects that has all the tags specified in the rule. If an object has additional tags specified, the rule is still applied. For more information about object tags, see S3 Object Tags.
9. Select at least one rule action:
    - Select **Expire current versions of objects** and specify the number of days to expire objects identified in the rule.
    - Select **Permanently delete noncurrent versions of objects** and specify the number of days. Specify the number of days an object remains noncurrent before Cohesity executes the deletion process. Cohesity calculates this duration

based on the time elapsed since the creation of its successor version, determining the object's noncurrent status.

- Select **Remove expired object delete markers** to remove delete markers of expired objects. Cohesity will remove a delete marker only if there are no noncurrent versions associated with it.
- Select **Delete incomplete multipart uploads** and specify the number of days to expire incomplete multipart uploads identified in the rule. Parts associated with these multipart uploads are deleted.

10. Click **Add**.

Supported S3 APIs and Response Headers

Cohesity introduces support for the following S3 Lifecycle APIs:

- PutBucketLifecycleConfiguration

  Cohesity also supports the PutBucketLifecycle API.
- GetBucketLifecycleConfiguration
- DeleteBucketLifecycle

The following APIs now support the `x-amz-expiration` response header:

- PutObject
- CopyObject
- CompleteMultipartUpload
- GetObject
- HeadObject

The following APIs now support the `x-amz-abort-date` and `x-amz-abort-rule-id` response headers:

- CreateMultipartUpload
- ListParts

For more information about these APIs, see Amazon S3 API Reference documentation.

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

S3 Bucket Ownership Controls

To store data in Amazon S3, you must create a bucket and upload data to that bucket as objects. An object consists of file data and any metadata that describes the file. An AWS S3 bucket is a container for objects. By default, the object creator owns the object and can

manage access through Access Control Lists (ACLs). However, as a bucket owner, you can use buckets to manage access to all objects in the bucket. You can disable ACLs on the S3 bucket and take ownership of all the objects in the bucket. You can use policies to control access to your data.

On a SmartFiles S3 View, you can use one of the modes to manage access on objects:

- **ObjectWriter**—This is the default mode of operation and remains in effect if object ownership controls are not configured on the bucket. The object creator automatically becomes the owner and remains the owner of the object.
- **BucketOwnerPreferred**—This mode is similar to **ObjectWriter** with the exception that the bucket owner becomes the object owner when the object is created with the `bucket-owner-full-control` canned ACL. Enabling or disabling this mode does not change the ownership of existing objects.
- **BucketOwnerEnforced**—This mode disables the ACLs on the bucket and objects. All existing objects/MPUs start being owned by the bucket owner. The owner of the object/MPU metadata does not change. If the mode is disabled, the previous owners start owning the objects/MPUs again. When enabled, all the new objects/MPUs are owned by the bucket owner. Ownership of these newly created objects/MPUs does not change when the mode is disabled. No canned ACLs `except bucket-owner-full-control` are allowed while creating the object/MPU. The grant specification is allowed only if it translates to full control for the bucket owner. This mode can be enabled only when the bucket is privately owned by the bucket owner and does not have any other grants. When this mode is enabled, bucket and object ACLs cannot be updated.

Support for Amazon S3 APIs

Cohesity introduces support for the following Amazon S3 APIs:

- PutBucketOwnershipControls
- GetBucketOwnershipControls
- DeleteBucketOwnershipControls

For more information about these APIs, see Amazon S3 API Reference.

Example Usage

To set bucket ownership controls, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api put-bucket-ownership-controls --bucket <test-bucket> --ownership-controls='{"Rules": [{"ObjectOwnership": "BucketOwnerPreferred"}]}'  --profile aws
```

To verify whether bucket ownership controls are set, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api get-bucket-ownership-controls --bucket <test-bucket> --profile aws
```

The following is a sample response:

```
{
    "OwnershipControls": {
        "Rules": [
            {
                "ObjectOwnership": "BucketOwnerPreferred"
            }
        ]
    }
}
```

To delete bucket ownership controls, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api delete-bucket-ownership-controls --bucket <test-bucket> --profile aws
```

To verify whether bucket ownership controls are deleted, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api get-bucket-ownership-controls --bucket <test-bucket> --profile aws
```

If bucket ownership controls are deleted, the following error message is displayed:

```
An error occurred (OwnershipControlsNotFoundError) when calling the GetBucketOwnershipControls operation: The bucket ownership controls were not found.
```

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

S3 Bucket Policy

To store data in Amazon S3, you must create a bucket and upload data to that bucket as objects. An object consists of file data and any metadata that describes the file. An S3 bucket is a container for objects. By default, only the creator can access the bucket or object. You can use a resource-based policy or an ACL to grant other users access to the bucket or object:

- If the policy is attached to a resource such as an S3 bucket, it is called a resource-based policy.
- An ACL is a list of grants identifying the grantees and the permissions granted. You can use ACLs to grant basic read/write permissions to other users.

Benefits of Using Bucket Policy

A bucket policy provides the following benefits over an ACL:

- ACL is restricted to a specific resource. For example, a bucket owner can grant access to a specific object by storing the ACL in the metadata of that object. If the owner

wants to grant permissions to a set of objects having the prefix `log_`, the owner has to grant permissions using ACL for each object. This can be cumbersome. You can use a bucket policy and grant permissions to all objects having the prefix `log_`.

- You can use a bucket policy to grant permissions for certain bucket subresources such as `lifecycle` and `versioning`. You cannot use ACL to grant permissions for these subresources.
- In a bucket policy, you can specify conditions, which permit the users to perform actions only if those conditions are met. For example, you can allow a user to list objects in a bucket only if the value of the `max-keys` parameter is less than 10.

Elements of a Bucket Policy

Bucket policy uses access policy language to define access permissions and specify conditions. The access policy is JSON-based and contains the following elements:

- Principals
- Resources
- Actions
- Effect
- Conditions

Principals

Cohesity S3 supports the following principals:

- Canonical user
- Anonymous user

Resources

You can use one of the following formats:

- `arn:aws:s3:::bucket_name/key_name`—AWS S3 uses Amazon Resource Name (ARN) to identify a resource.
- `urn:csf:s3:::bucket_name/key_name`—Cohesity S3 uses this format to identify a resource.

Actions

A set of predefined operations for each resource. Each action corresponds to a specific S3 operation such as `s3:PutObject` or `s3:PutObjectAcl`.

Effect

You can use one of the following values for the Effect element:

- `Allow`
- `Deny`

The default value is `Deny`.

Conditions

You can specify the conditions for a policy using this element. A `Condition` element can contain multiple conditions and each condition can contain multiple key-value pairs.

Cohesity S3 supports the following condition operators:

- String—`StringEquals`, `StringEqualsIgnoreCase`, `StringNotEquals`, and `StringNotEqualsIgnoreCase`.
- Numeric—`NumericEquals`, `NumericNotEquals`, `NumericLessThan`, `NumericLessThanEquals`, `NumericGreaterThan`, and `NumericGreaterThanEquals`.
- If a condition has multiple keys or values, you can use the `ForAllValues` or `ForAnyValue operator`.
- You can add the `IfExists` operator to the end of any condition operator name except the `Null` condition.

Cohesity S3 supports the following condition keys:

Condition Key	Description	Type
`s3:ExistincgObjectTag/<tag-key>`	Filters access by existing object tag key and value	String
`s3:RequestObjectTagKeys`	Filters access by the tag keys to be added to objects	ArrayOfString
`s3:RequestObjectTag/<tag-key>`	Filters access by the tag keys and values to be added to objects	String
`s3:max-keys`	Filters access by the maximum number of keys returned in a `ListBucket` request	Numeric
`s3:prefix`	Filters access by key name prefix	String
`s3:delimiter`	Filters access by the delimiter parameter	String
`s3:x-amz-acl`	Filters access by canned ACL in the request's `x-amz-acl` header	String

The following table lists the supported actions and the corresponding condition keys:

Action	Supported Condition Keys
ListBucket	• s3:prefix • s3:max-keys • s3:delimiter
ListBucketVersions	• s3:prefix • s3:max-keys • s3:delimiter
PutObject	• s3:x-amz-acl • s3:RequestObjectTagKeys • s3:RequestObjectTag/<tag-key>
PutObjectAcl	• s3:x-amz-acl • s3:ExistingObjectTag/<tag-key>
PutObjectVersionAcl	• s3:x-amz-acl • s3:ExistingObjectTag/<tag-key>
DeleteObjectTagging	• s3:ExistingObjectTag/<tag-key>
DeleteObjectVersionTagging	• s3:ExistingObjectTag/<tag-key>
GetObject	• s3:ExistingObjectTag/<tag-key>
GetObjectVersion	• s3:ExistingObjectTag/<tag-key>
GetObjectAcl	• s3:ExistingObjectTag/<tag-key>

Action	Supported Condition Keys
GetObjectTagging	• s3:ExistingObjectTag/<tag-key>
GetObjectVersionAcl	• s3:ExistingObjectTag/<tag-key>
GetObjectVersionTagging	• s3:ExistingObjectTag/<tag-key>
PutObjectTagging	• s3:ExistingObjectTag/<tag-key> • s3:RequestObjectTagKeys • s3:RequestObjectTag/<tag-key>
PutObjectVersionTagging	• s3:ExistingObjectTag/<tag-key> • s3:RequestObjectTagKeys • s3:RequestObjectTag/<tag-key>
BypassGovernanceRetention	• s3:RequestObjectTagKeys • s3:RequestObjectTag/<tag-key>

Example of an S3 Bucket Policy

The following is a sample S3 bucket policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowUser2ToWriteObject",
      "Effect": "Allow",
      "Principal": {
        "CanonicalUser": [
          "USER2@LOCAL@"
        ]
      },
      "Action": [
        "s3:PutObject"
      ],
      "Resource": "arn:aws:s3:::BUCKET-NAME/*"
    },
    {
      "Sid": "AllowUser2ToGetObjectOnlyWhenTagsHavePublicSecurity",
```

```
      "Effect": "Allow",
      "Principal": {
        "CanonicalUser": [
          "USER2@LOCAL@"
        ]
      },
      "Action": [
        "s3:GetObject"
      ],
      "Resource": "arn:aws:s3:::BUCKET-NAME/*",
      "Condition": {
        "StringEquals": {
          "s3:ExistingObjectTag/security": "public"
        }
      }
    }
  ]
}
```

Support for Amazon S3 APIs

Cohesity introduces support for the following Amazon S3 APIs:

- PutBucketPolicy
- GetBucketPolicy
- DeleteBucketPolicy

For more information about these APIs, see Amazon S3 API Reference.

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

S3 Object Tags

Cohesity supports S3 object tagging. Object tags are unique key-value pairs that are applied to individual S3 objects. You can create, update, or delete object tags at any time during the lifetime of an S3 object. You can use object tags to categorize storage. For example, you can add details such as data owners and users, business unit names, media types, data format, or security classification.

Considerations

- You can add up to 10 tags for each object.
- Tag keys must be unique within each object.
- A tag key can be up to 128 Unicode characters in length.
- A tag value can be up to 256 Unicode characters in length.
- Keys and values are case-sensitive.

Support for Amazon S3 APIs

Cohesity introduces support for the following Amazon S3 APIs:

- PutObjectTagging
- GetObjectTagging
- DeleteObjectTagging

Cohesity enhances support for the following Amazon S3 APIs:

- PutObject
- GetObject
- CreateMultipartUpload
- CopyObject
- PutBucketLifecycleConfiguration

Support for New Request and Response Headers

Cohesity introduces support for the following new request headers:

- `x-amz-tagging`
- `x-amz-tagging-directive`

Cohesity introduces support for the following new response header:

- `x-amz-tagging-count`

For more information about these APIs, see Amazon S3 API Reference.

Example Usage

To update tags on an existing object, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api put-object-tagging --bucket <test-bucket>  --key obj2 --tagging '{"TagSet": [{ "Key": "Color", "Value": "Blue" }, { "Key": "Project", "Value": "Alpha" }]}' --profile vmrobo2
```

To get the entire tag set from an existing object, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api get-object-tagging --bucket <test-bucket> --key obj2 --profile vmrobo2
```

The following is a sample response:

```
{
    "TagSet": [
        {
            "Key": "Color",
            "Value": "Blue"
        },
```

```
        {
            "Key": "Project",
            "Value": "Alpha"
        }
    ]
}
```

To delete all tags from an existing object, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api delete-object-tagging --bucket <test-bucket>   --key obj2 --profile vmrobo2
```

To create an object with tags, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api put-object --bucket <test-bucket> --key obj1 --body ~/test/obj1.txt --profile vmrobo2  --tagging 'Color=Red&Project=Beta'
```

The following is a sample response:

```
{
"ETag": "\"0accc61f62bb422443c3f0f2d351ca16\""
}
```

To get the object tag count, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api get-object --bucket <test-bucket> --key obj1 --profile vmrobo2 outfile
```

The following is a sample response:

```
{
    "LastModified": "2022-10-21T13:51:13+00:00",
    "ContentLength": 1024,
    "ETag": "\"0accc61f62bb422443c3f0f2d351ca16\"",
    "Metadata": {},
    "TagCount": 2
}
```

To add tags during the initiation of multipart upload, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api create-multipart-upload --bucket <test-bucket> --profile vmrobo2 --key obj4 --tagging 'Color=Red&Project=Beta'
```

To verify whether the tags are added, execute the GetObjectTagging API. The following is a sample response:

```
{
    "Bucket": "test-bucket",
    "Key": "obj4",
    "UploadId": "4706"
}
```

> **Note:** Once you complete this multipart upload, you can read/update tags on the resultant object.

To copy an existing object and associated tags, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api copy-object --copy-source test-bucket/obj1 --key obj2 --bucket <test-bucket> --profile vmrobo2 --tagging-directive COPY
```

> **Note:** Even if you do not specify the `--tagging-directive COPY` header, the `CopyObject API` copies all object tags.

The following is a sample response:

```
{
    "CopyObjectResult": {
        "ETag": "\"0accc61f62bb422443c3f0f2d351ca16\"",
        "LastModified": "2022-10-21T13:58:44+00:00"
    }
}
```

To copy an existing object and add new tags on the copy, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api copy-object --copy-source test-bucket/obj1 --key obj2 --bucket <test-bucket> --profile vmrobo2 --tagging-directive REPLACE  --tagging '{"TagSet":[{"Key":"Owner1", "Value":"Blue"}]}'
```

The following is a sample response:

```
{
    "CopyObjectResult": {
        "ETag": "\"0accc61f62bb422443c3f0f2d351ca16\"",
        "LastModified": "2022-10-21T13:58:44+00:00"
    }
}
```

To specify object tags while creating a new lifecycle configuration for the bucket, execute the following command:

```
aws --endpoint-url https://<cluster_ip or hostname>:3000 s3api put-bucket-lifecycle-configuration --bucket <test-bucket> --lifecycle-configuration
```

The following is a sample response:

```
'{
  "Rules": [
    {
```

```
      "ID": "id0",
      "Status": "Enabled",
      "Filter": {
        "Tag": {
          "Key": "k0",
          "Value": "v0"
        }
      },
      "Expiration": {
        "Days": 1
      }
    }
  ]
}'
```

Use Object Tags with S3 Lifecycle Management

You can use object tags to enable fine-grained object lifecycle management by specifying a tag-based filter in a lifecycle management rule. The `PutBucketLifecycleConfiguration` API has been enhanced to accept `Tag` and `And` filters.

Use Object Tags with Bucket Policy

You can use object tags to enable fine-grained access control of permissions by specifying object tags in the condition of an access policy. You can use bucket access policies to restrict access to data with specific tags, group objects with unique compliance requirements, or create lifecycle policies for objects with specific tags.

Cohesity supports the following three condition keys:

- `s3:ExistingObjectTag/<tag-key>`—Use this condition key in the policy to grant users permission to perform specific operations on objects. The condition limits the permission to only objects that have the tag key and tag value specified in `s3:ExistingObjectTag/<tag-key>`.
- `s3:RequestObjectTagKeys`—Use this condition key to restrict tag keys that the user wants to allow on objects. This is useful when adding tags to objects using the `PutObjectTagging` and `PutObject` requests.
- `s3:RequestObjectTag/<tag-key>`—Use this condition key to restrict tag keys and tag values the user wants to allow on objects. This is useful when adding tags to objects using the `PutObjectTagging` and `PutObject` requests.

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

Support for UploadPartCopy API

You can use the `UploadPartCopy` API to upload a part by copying data from an existing object as the data source. In your request, specify the following:

- Data source by adding the `x-amz-copy-source` request header
- Byte range by adding the `x-amz-copy-source-range` request header

For more information, see Amazon S3 REST API Documentation.

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

S3 Access Control Lists

Cohesity supports Access Control Lists (ACLs) for both buckets and objects. You can specify ACL grantees to define fine-grained access control. Cohesity supports all the features of AWS S3 ACL.

> **Note:** Cohesity recommends using bucket policies instead of ACLs for managing access control. Bucket policies provide more granular control and flexibility in defining access permissions for buckets and objects. For more information about bucket policies, see S3 Bucket Policy.

The table below provides an overview of the ACL permissions and their corresponding API permissions when ACL permissions are granted on a bucket or an object:

ACL Permission	API Permissions Granted on a Bucket	API Permissions Granted on an Object
`READ`	• `s3:ListBucket` • `s3:ListBucketVersions` • `s3:ListBucketMultipartUploads`	• `s3:GetObject` • `s3:GetObjectVersion`
`WRITE`	• `s3:PutObject` • `s3:DeleteObject` The bucket owner can create, overwrite, or delete any object in the bucket.	Not applicable
`READ_ACP`	• `s3:GetBucketAcl`	• `s3:GetObjectAcl` • `s3:GetObjectVersionAcl`

ACL Permission	API Permissions Granted on a Bucket	API Permissions Granted on an Object
`WRITE_ACP`	• `s3:PutBucketAcl`	• `s3:PutObjectAcl` • `s3:PutObjectVersionAcl`
`FULL_CONTROL`	Equivalent to granting `READ`, `WRITE`, `READ_ACP`, and `WRITE_ACP` ACL permissions. Additionally, the following API permissions are granted: • `s3:GetObjectRetention` • `s3:PutObjectRetention` • `s3:BypassGovernanceRetention` • `s3:GetObjectLegalHold` • `s3:PutObjectLegalHold` • `s3:PutObjectTagging` • `s3:PutObjectVersionTagging` • `s3:GetObjectTagging` • `s3:GetObjectVersionTagging` • `s3:DeleteObjectTagging` • `s3:DeleteObjectVersionTagging`	Equivalent to granting `READ`, `READ_ACP`, and `WRITE_ACP` ACL permissions.

The following additional permissions can only be assigned to the bucket owner:

- `s3:DeleteObjectVersion`
- `s3:PutBucketObjectLockConfiguration`
- `s3:GetBucketObjectLockConfiguration`

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

S3 Virtual-Hosted–Style Requests

Cohesity supports virtual-hosted–style requests. In a virtual-hosted–style request, the bucket name is part of the domain name in the URL. For example, `https://bucket-name.cohesity.com/key-name`.

> **Note:** This feature allows Cohesity SmartFiles to function as storage for Snowflake external tables. For more information about Cohesity SmartFiles integration with Snowflake Data Cloud, see Extend Snowflake Data Access Securely and Cost Effectively.

Prerequisites

- Add a DNS A record in the DNS management software by specifying the IP address of the Cohesity cluster and its fully qualified domain name (FQDN). The cluster IP address must be DNS-resolvable. The following is an example of a DNS record:

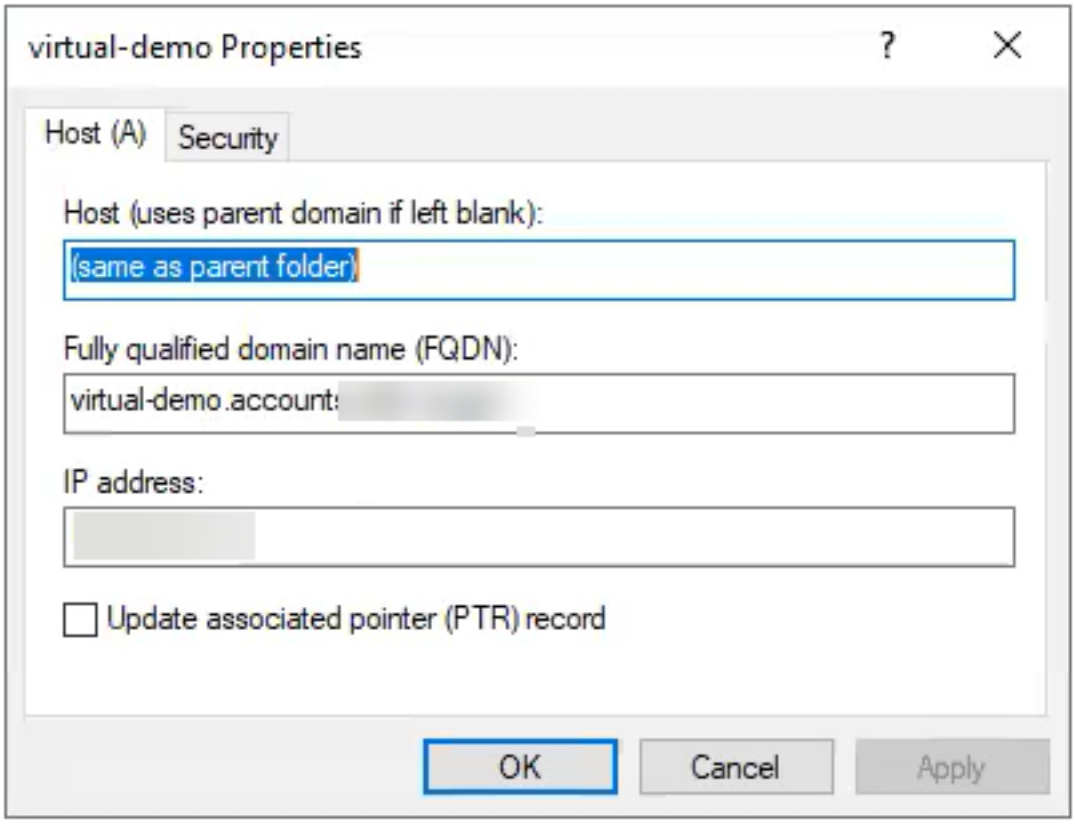

- To enable the resolution of any bucket name added to the FQDN to the Cohesity cluster, include a DNS A record with a wildcard (*) in the FQDN. The following is an

example of (*) in the FQDN:

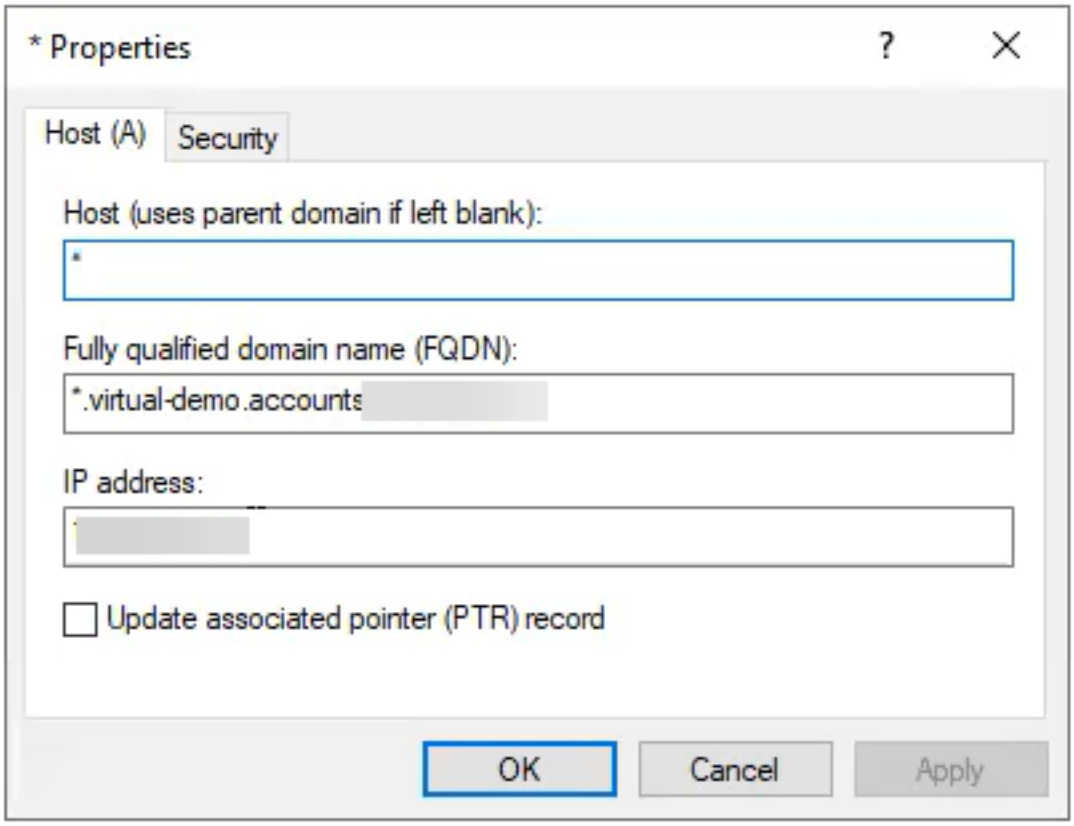

Enable Virtual Hosted Style Requests

You can only enable virtual-hosted-style requests through the CLI.

To enable virtual-hosted-style requests, do the following:

1. Log in to Cohesity CLI. For more information, see Using the Cohesity CLI.
2. Enter the following command:

   ```
   iris_cli cluster update-s3-virtual-hosted-domain-names virtual-hosted-domain-names=<domain names>
   ```

   > **Note:** You can add multiple domain names separated by a comma.

3. Enter the username and password.
4. Verify whether the virtual-hosted-style requests are enabled by executing the following command:

   ```
   iris_cli cluster info
   ```

   If the feature is enabled, the following field is displayed:

   ```
   S3 VIRTUAL HOSTED DOMAIN NAMES
   ```

Important Points to Note

- If the Host header value ends with one of the specified hostnames, the request is interpreted as following the virtual-hosted-style request, and the initial section of the header is treated as the bucket name. Consequently, the URL is considered to have an object key only.
- If the Host header value is exactly the same as one of the hostnames, the request is considered to be a path-style request.
- You can access an S3 bucket using a path-style request or a virtual-hosted-style request.

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

Presigned URLs

Presigned URLs allow temporary access to a specific S3 resource and operation without requiring additional authentication or authorization. All authentication information, including the request signature, is included directly in the query string. This type of URL is referred to as a presigned URL.

Presigned URLs have an expiration time. You can create and manage presigned URLs through AWS CLI or AWS SDKs. You can set expiration times up to seven days from the time of URL signing.

> **Note:** Cohesity supports only the AWS4-HMAC-SHA256 algorithm.

Cohesity supports the following query parameters:

- X-Amz-Algorithm
- X-Amz-Credential
- X-Amz-Date
- X-Amz-Expires
- X-Amz-SignedHeaders
- X-Amz-Signature

For more information about managing presigned URLs through AWS CLI or AWS SDKs, see Sharing objects with presigned URLs.

For multi-part uploads with presigned URLs, complete the following steps:

1. **Initiate Multi-part Upload**: Authorize the initiation using HTTP Headers, and obtain a UploadId in the response.
2. **Generate Presigned URLs**: Create presigned URLs for each part using the obtained UploadId.
3. **Upload Parts**: Use the corresponding presigned URL for each part. These uploads are authorized using query string parameters. The response for each upload contains an ETag.
4. **Complete the Upload**: Authorize the completion using HTTP Headers. This step requires the ETags received in the PUT requests from the previous step.

Related Topics

- Supported S3 APIs
- Supported S3 API Headers

Add Subnets to Allowlists

Before you can mount and access a view, ensure the IP address of the system that is mounting the view is specified in a subnet that has been added to an allowlist.

Cohesity supports the following types of allowlists:

- **Global Allowlist**—The global allowlist applies to all views on the Cohesity cluster, but is overridden by a view allowlist or view share allowlist, if they exist.
- **View Allowlist**—Only the systems with IP addresses in the subnets listed in the view's allowlist have privileges to access and mount this view.

Update the Global Allowlist

To add a subnet to the Cohesity cluster's global allowlist:

1. Navigate to **SmartFiles > Views**.
2. Click the **Global Settings** tab.
3. In the **Global Subnet Allowlist** section, click **Add**.
4. In the **Add Allowlist** dialog, enter a Subnet IP in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).

   > **Note:** CIDR prefix values are allowed from 0 to 32.

5. Select any or all permissions as required.

> **Note:** By default, the ability to set **S3 Permissions** for restricting access to specific subnets is enabled.

6. Select **NFS Root Squash** as applicable.
7. In the **Description** field, enter a description for the allowlist.
8. Click **Add**:

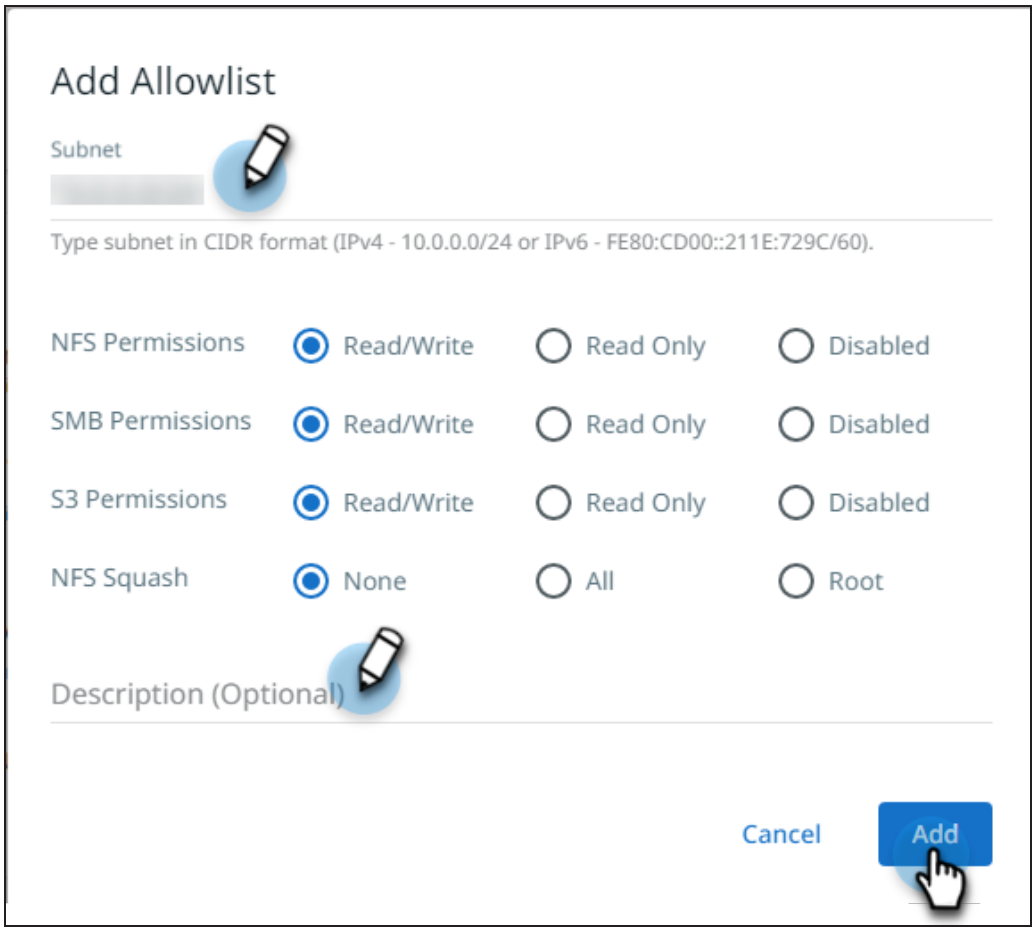

9. Repeat the preceding steps for specifying additional subnets in the global allowlist.

Update a View Allowlist

To add a subnet to an existing view's allowlist:

1. Navigate to **SmartFiles > Views**.
2. Locate the required View in the list, click the ⋮, and select **Edit** ( ).
3. In Advanced settings, click **Edit** for the Security category.

4. **Override Global IP Allowlist:** Turn on **Override Global IP Allowlist** if you want to specify a custom list for this view. Only the systems with IP addresses in the subnets listed in the view's allowlist have privileges to access and mount this view. If you enable the **Override Global IP Allowlist** toggle, the machines listed in the subnets of the global Cohesity cluster allowlist cannot access or mount this view unless the IP address of the machine is also listed in the View's Allowlist.

   Extend Global IP Allowlist: Allows you to extend the Global IP Allowlist.

   To add subnets to the allowlist of this view:

   a. Under **IP Allowlist**, select: **Override Global IP Allowlist** or **Extend Global IP Allowlist**.
   b. Click **Add**.
   c. In the **Add Allowlist** screen, type subnet in CIDR format.
   d. Select any or all permissions as required.
   e. Select **NFS Squash** options.

   The **NFS all squash** option is like NFS Root Squash. However, for given IP addresses, all users are mapped to the All Squash User/Group configured in the view. The default All Squash User for all views is "nobody".

   > **Note:** If the **Override Global Allowlist** toggle is off, this view can only be accessed and mounted by systems specified by the global Cohesity cluster allowlist.

5. Click **Update View**.

For more information on global settings for each protocol, see the following topics:

- Global Settings for NFS
- Global Settings for SMB

SmartFiles Dashboard

The SmartFiles dashboard provides several key metrics for better performance visibility. You can proactively monitor the performance and manage your SmartFiles deployments more effectively. The SmartFiles dashboard provides:

- Consumption metrics to track the total usable capacity and storage consumed
- Performance metrics such as throughput, IOPS, and latency
- File distribution by type
- Clients connected across nodes

To view the **SmartFiles** dashboard:

1. Log in to the Cohesity cluster.
2. Click **Dashboards** and select **SmartFiles** from the drop-down list

   The following illustration shows an example of the **SmartFiles** dashboard:

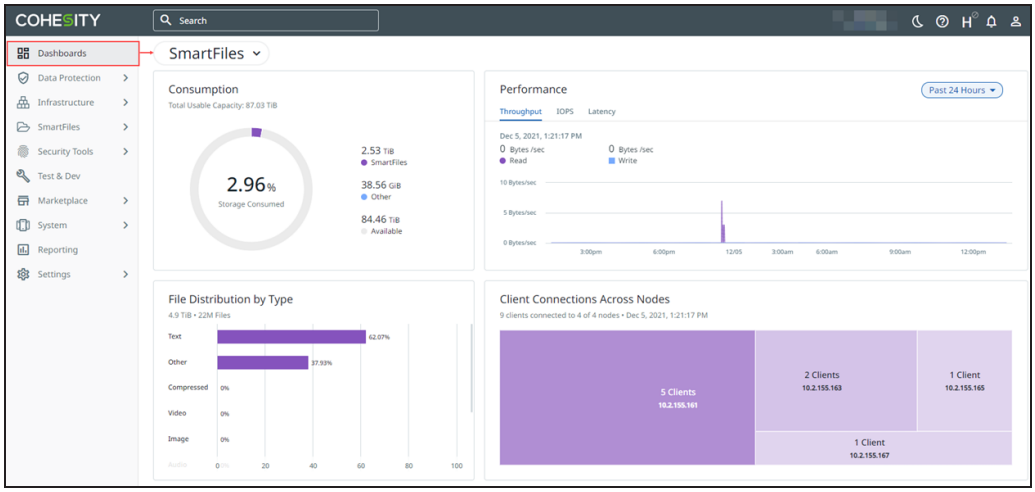

Consumption

The **Consumption** widget in the SmartFiles dashboard provides a summary of the total usable capacity and the following metrics:

- Storage consumed by SmartFiles
- Storage consumed by other services
- Storage available for consumption

The following illustration shows an example of the **Consumption** widget. Click the widget to drill down to the **Consumption** page:

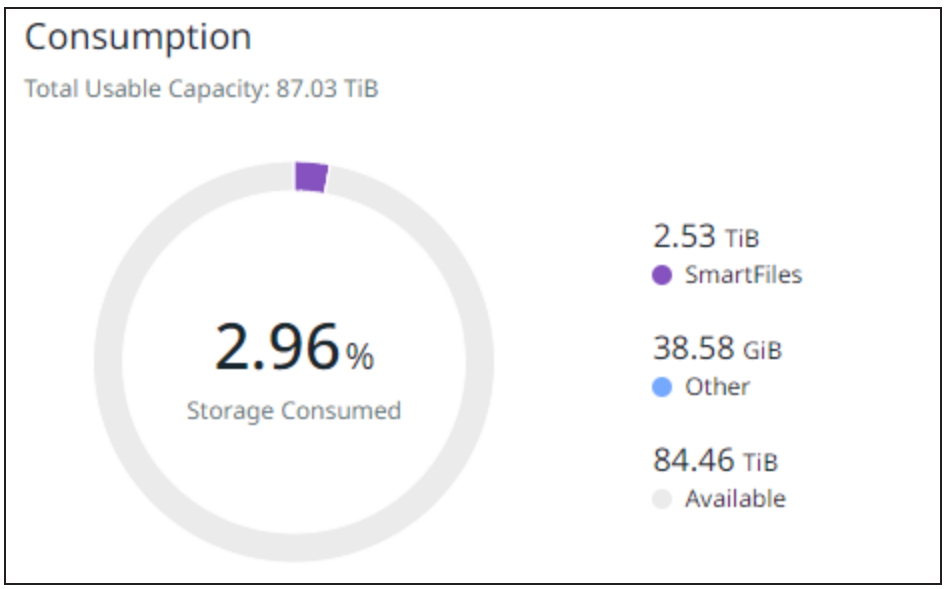

For more information, see Consumption Page.

Performance

The **Performance** widget in the SmartFiles dashboard displays performance metrics such as throughput, IOPS, and latency. You can view the performance data for the past 1 hour, past 12 hours, past 24 hours, past 7 days, or past 30 days. Click the widget to drill down to the **Performance** page.

Throughput

The **Throughput** chart displays the amount of data transferred through external network interfaces in the cluster. Hover over the chart to view the throughput for read and write operations at a particular date and time. The following illustration shows an example of the **Throughput** chart:

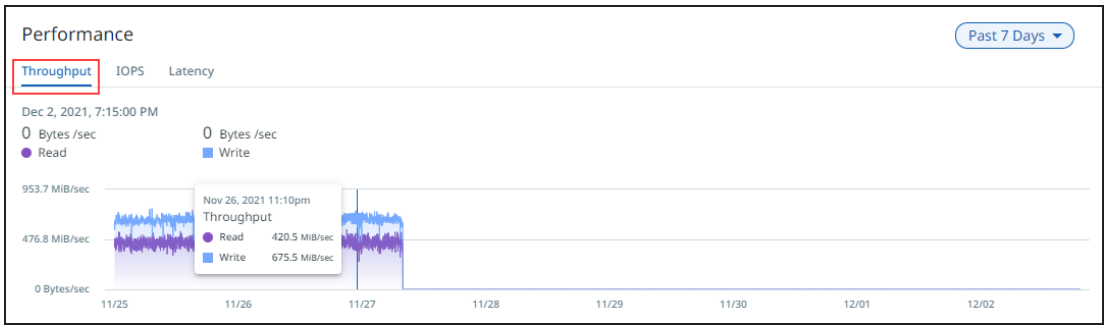

IOPS

The **IOPS** chart displays the average number of read, write, or metadata operations from external clients serviced per second. Hover over the chart to view the IOPS for read and write operations at a particular date and time. The following illustration shows an example of the **IOPS** chart:

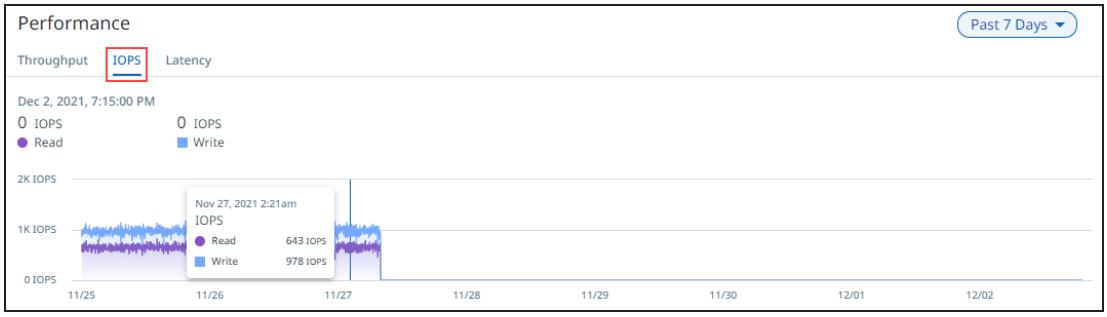

Latency

The **Latency** chart displays the average amount of time required for the cluster to process incoming and outgoing external client requests. Clients might observe additional latencies caused by the network. Hover over the chart to view the latency for read and write

operations at a particular date and time. The following illustration shows an example of the **Latency** chart:

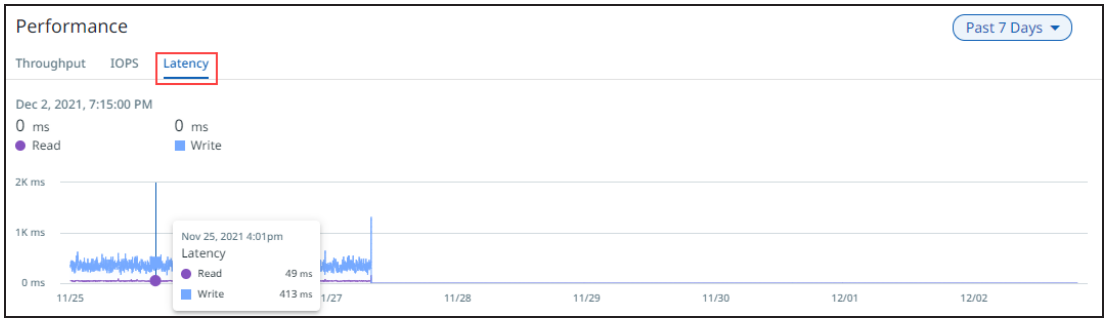

For more information, see Performance Page.

File Distribution by Type

The **File Distribution by Type** widget in the SmartFiles dashboard displays the distribution of files by type on SmartFiles views. The widget displays the total storage consumed and the total number of files. On the chart, hover over the file type to view the storage consumed and the number of files for the corresponding file type.

To view the **File Distribution by Type** widget:

1. Log in to the Cohesity cluster.
2. Click **Dashboards** and select **SmartFiles** from the drop-down list.

   The **File Distribution by Type** widget appears in the SmartFiles dashboard:

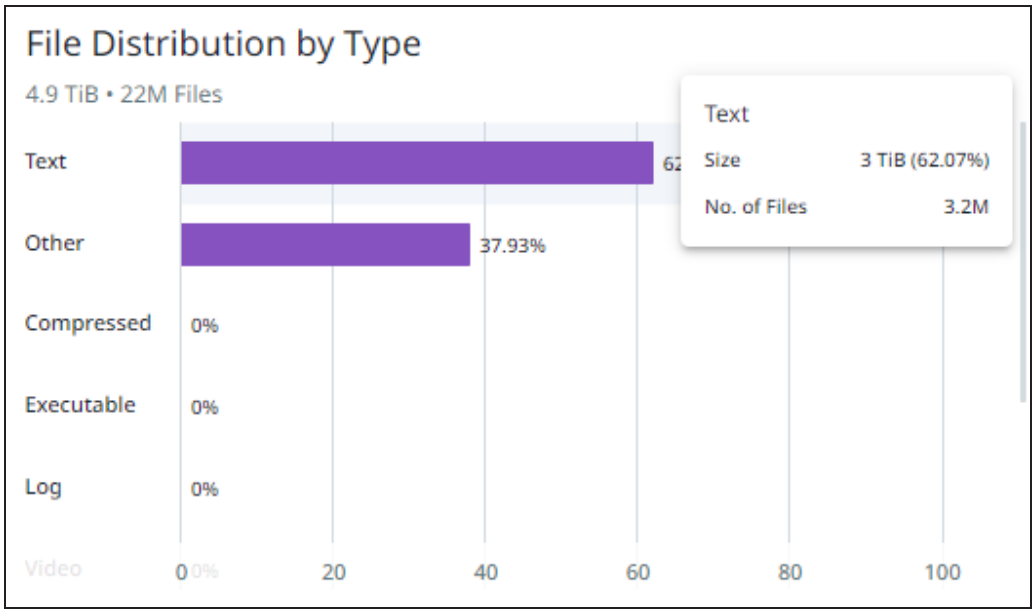

   In the **File Distribution by Type** widget, file types are categorized into the following groups:

- **Text**—.doc, .docx, .pdf, .txt, .ods, .xls, .xlsx, .ppt, .pptx, and .key
- **Image**—.jpeg, .jpg, .png, .bmp, and .gif
- **Audio**—.mp3, .wav, .mp4, .aif, and .mid
- **Video**—.mpg, .mov, .wmv, .rm, and .3gp
- **Log**—.log, .info, .error, and .warning
- **Compressed**—.7z, .rar, .rpm, .gz, .z, .zip, .pkg, and .deb
- **Database**—.csv, .dat, .db, .dbf, .mdb, .sql, and .xml
- **Executable**—.apk, .bat, .bin, .exe, .jar, .py, .cgi, and .com

Client Connections Across Nodes

The **Client Connections Across Nodes** widget in the SmartFiles dashboard displays a summary of clients connected across nodes. Hover over a section in the widget to view the following details:

- Node IP address
- Total number of clients connected
- Total number of NFS clients
- Total number of SMB clients

The following illustration shows an example of clients connected across nodes:

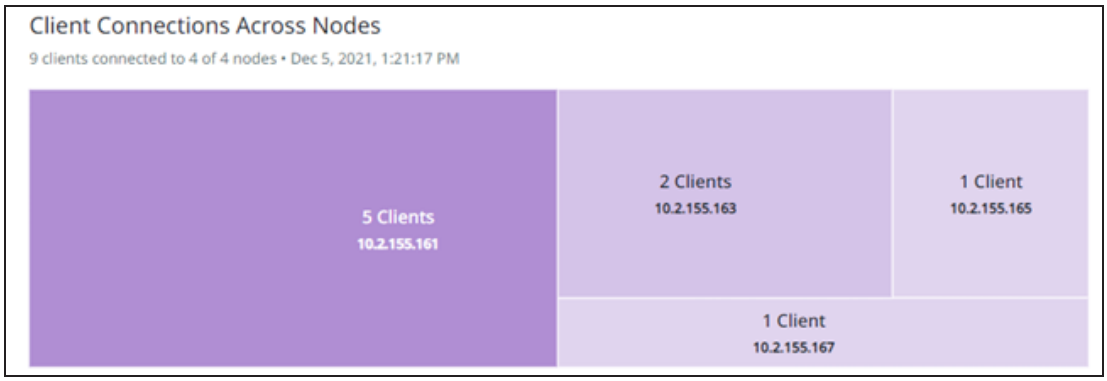

Click any node IP address in the widget to drill down to the **Client Connections** page. The client(s) connected to the specific node are displayed on the table.

For more information, see Client Connections Page.

Consumption Page

The **Consumption** page includes the following sections:

- Glance bar
- Trend chart
- Total consumption per View

To view the **Consumption** page:

1. Log in to the Cohesity cluster.
2. Go to **SmartFiles > Consumption**:

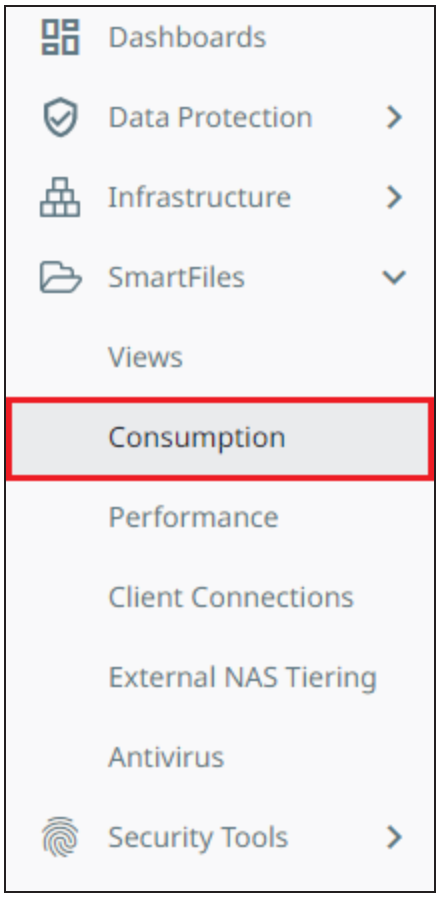

Glance Bar

The glance bar provides an overall summary of the consumption metrics. The following metrics are displayed:

- Logical—Logical size of your active file system. Logical usage is updated every four hours.
- Data Reduction—Size reduction ratio. This is the logical data divided by physical.
- Physical—Storage capacity consumed after data reduction operations such as deduplication and compression.

- Resiliency Impact—Storage consumed by resiliency operations such as erasure coding and redundancy factor.
- Storage Consumed—Final storage consumed after including resiliency impact. Storage consumed is updated every two days.
- Files—Number of files stored.
- Directories—Number of directories stored.

The following is an illustration of the glance bar:

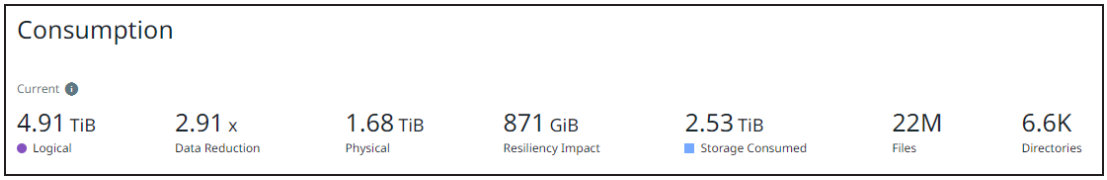

Trend Chart

The trend chart displays the trend of total consumption across all views in the past 30 days. Hover over the chart to view the logical and storage consumed data at a particular date and time:

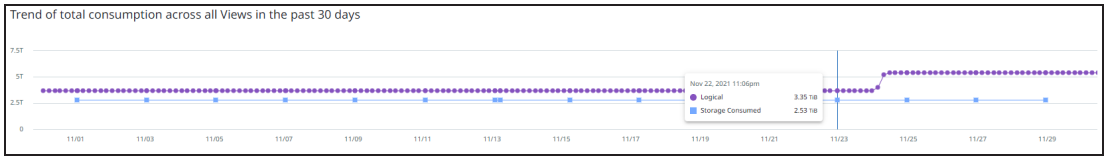

Total Consumption Per View

The Total consumption per View table lists all the views and displays the total consumption and the percentage of logical quota consumed for each view.

If there are a large number of views, you can enter the view name in the search bar to instantly filter the views. Click the view name to drill down to the view details page.

The following is an illustration of the table:

Consumption Growth Tab

The **SmartFiles > Consumption** page now includes a **Growth** tab which displays various metrics and trends to measure the growth in storage consumption for the:

- Past 24 hours
- Past 7 days
- Past 30 days

Glance Bar

The glance bar displays the following details:

- **Logical**—Growth in the logical size of your active file system.
- **Physical**—Growth in the physical storage.
- **Storage Consumed**—Growth in the final storage consumed after including resiliency impact.
- **Files**—Growth in the number of files stored.
- **Directories**—Growth in the number of directories stored.

Trend Chart

The trend chart displays the trend of consumption growth across all views for the time range you selected.

Click **Logical & Storage Consumed** to view the trend for logical and storage consumed. Hover over the chart to view the logical and storage consumed data at a particular date and time.

Click **Files & Directories** to view the growth trend for files and directories. Hover over the chart to view the files and directories stored at a particular date and time.

Growth Per View

The **Growth per View** table lists all the views and displays the consumption growth of each view. If there are a large number of views, you can enter the view name in the search bar to instantly filter the views. Click the view name to drill down to the view details page.

The growth is displayed in bytes and percentages. Click **Bytes** or **%** to update the table data. You can sort the columns in ascending or descending order. For example, if you sort the **Files** column, you can identify the view storing the most number of files and the growth.

Performance Page

The **Performance** page provides the overall trend for all views, top views, the overall trend for all clients, and top clients.

To view the **Performance** page:

1. Log in to the Cohesity cluster.
2. Go to **SmartFiles > Performance**:

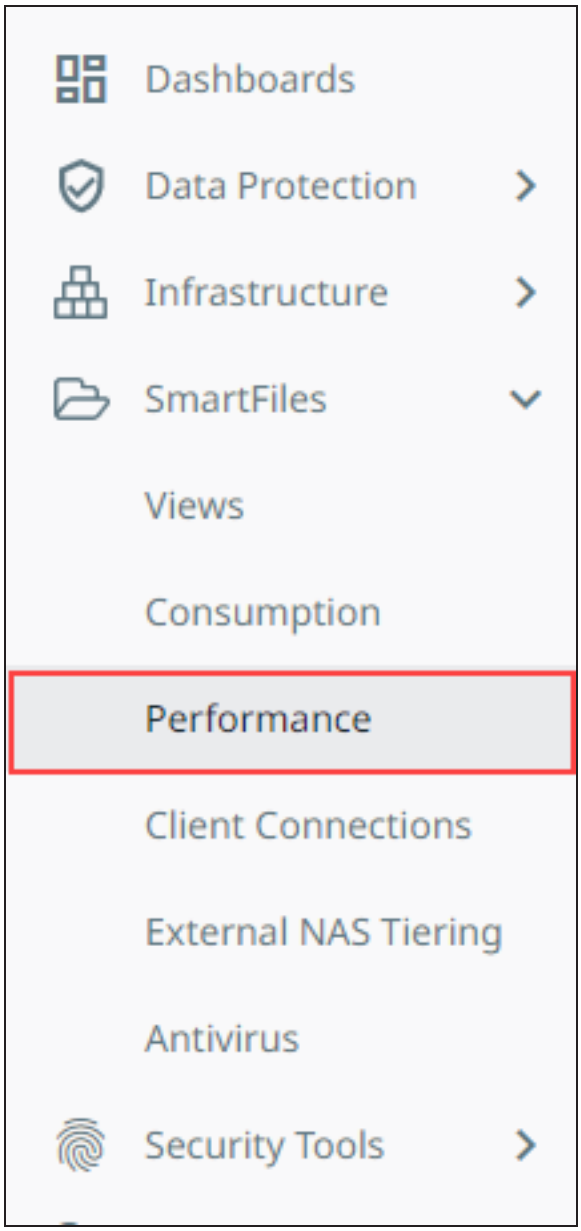

View Performance by Views

The **Performance** page includes multiple filters that allow you to pare down data. To view performance by views, click the **By** filter and select **View**:

- **Protocol**—Select the protocol: **NFS**, **SMB**, or **S3**.
- **Metrics**—After you select the protocol, you can select the metric(s). The options listed in the **Metrics** drop-down list are populated based on the protocol you select. The data for each metric is graphically represented so that you can better understand trends and patterns. For more information about metrics, see:

  - NFS Metrics
  - SMB Metrics
  - S3 Metrics
- **View**—Select a view to see the trends specific to the view.
- **Range**—Select a time range: Past 1 Hour, Past 12 Hours, Past 24 Hours, Past 7 Days, or Past 30 Days.

The following illustration shows an example of the overall trend of all views:

To view the top views, click **Top Views**:

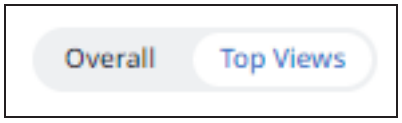

The time range for **Top Views** can be Past 1 Hour, Past 12 Hours, or Past 24 Hours.

You can sort NFS and SMB views by Read (Peak) or Write (Peak) and you can sort the S3 views by Put (Peak) or Get (Peak).

View Performance by Clients

The **Performance** page includes multiple filters that allow you to pare down data. To view performance by clients, click the **By** filter and select **Client**:

- **Protocol**—Select the protocol: **NFS**, **SMB**, or **S3**.
- **Metrics**—After you select the protocol, you can select the metric(s). The options listed in the **Metrics** drop-down list are populated based on the protocol you select. The data for each metric is graphically represented so that you can better understand trends and patterns.

> **Note:** Client statistics for Throughput, Latency, IOPS, and Error metrics are aggregated for all protocols, irrespective of the selected protocol.

For more information about metrics, see:

- NFS Metrics
- SMB Metrics
- S3 Metrics

- **Client**—Select a client to view the trends specific to the client.
- **Range**—Select a time range: Past 1 Hour, Past 12 Hours, Past 24 Hours, Past 7 Days, or Past 30 Days.

The following illustration shows an example of the overall trend of all clients:

To view the top clients, click **Top Clients**:

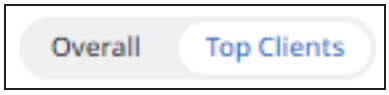

The time range for **Top Clients** can be Past 1 Hour, Past 12 Hours, or Past 24 Hours.

View Performance by Nodes

The **Performance** page includes multiple filters that allow you to analyze the performance of specific entities. To view performance by nodes, click the By filter and select **Node**:

- **Node** – Select a node to see the trends specific to the node.
- **Disk** – Select a disk to see the trends specific to the disk in the selected node.
- **Protocol** - Select the protocol: *NFS, SMB, or S3*.

> **Note:** You cannot filter by Disk and Protocol together.

- **Metrics** - Select the required metric(s). The options listed in the Metrics drop-down list are populated based on the protocol you select. If you do not select a protocol, only the metrics that apply to all protocols are displayed. The data for each metric is graphically represented so that you can better understand trends and patterns. For more information about metrics, see:
  - NFS Metrics
  - SMB Metrics
  - S3 Metrics
- **Range** - Select a time range: Past 1 Hour, Past 12 Hours, Past 24 Hours, Past 7 Days, or Past 30 Days.

The following illustration shows an example of the overall trend of all clients:

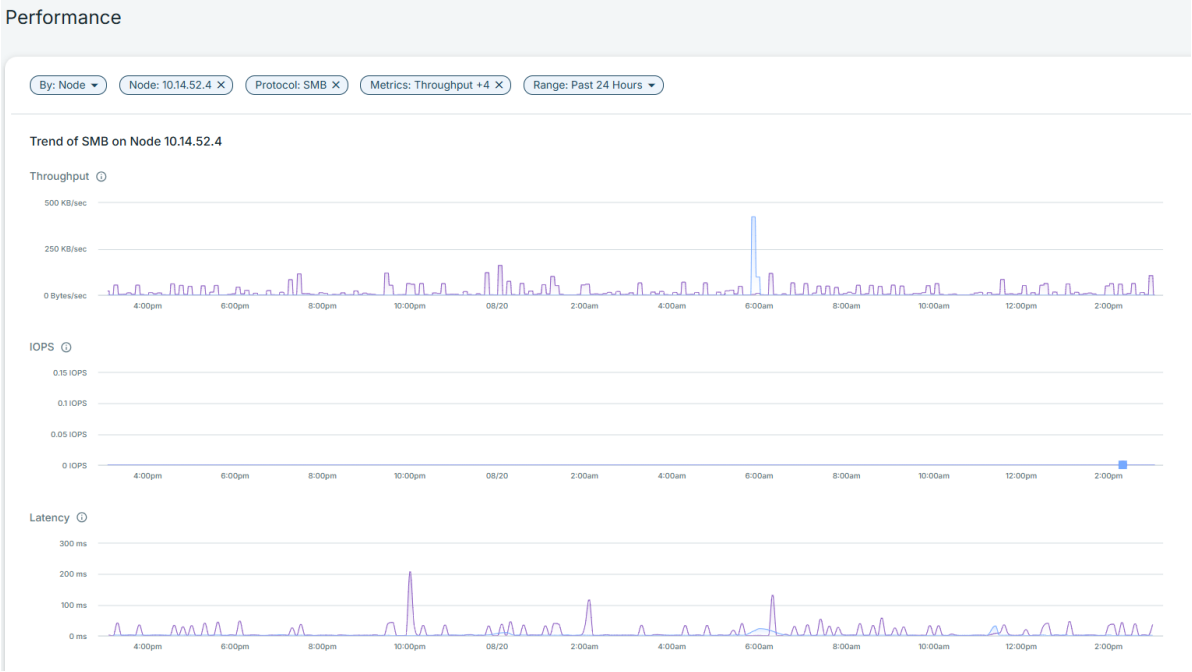

NFS Metrics

Metric	Description
IOPS	The average number of read, write, or metadata operations from external clients serviced per second.
Throughput	The amount of data transferred through external network interfaces in the Cohesity cluster.

Metric	Description
Latency	The average amount of time required for the Cohesity cluster to process incoming and outgoing external client requests. Clients might observe additional latencies caused by the network.
Error	The number of data read and write errors generated in serving external client requests.
Mount	IOPS and latency for NFS client requests to mount file systems over the network.
Unmount	IOPS and latency for NFS client requests to unmount file systems over the network.
Create File	IOPS and latency for NFS client requests to create a file.
Make Directory	IOPS and latency for NFS client requests to create a directory.
Remove File	IOPS and latency for NFS client requests to remove a file.
Remove Directory	IOPS and latency for NFS client requests to remove a directory.
Read Directory	IOPS and latency for NFS client requests to read the filenames in a directory.
Read Directory Plus	IOPS and latency for NFSv3 client requests to read the filenames and attributes in a directory.
Rename	IOPS and latency for NFS client requests to rename a directory.
Lookup	IOPS and latency for NFS client requests to search a directory for a file.
Symlink	IOPS and latency for NFS client requests to create a symbolic link in a remote file system.

SMB Metrics

Metric	Description
IOPS	The average number of read, write, or metadata operations from external clients serviced per second.
Throughput	The amount of data transferred through external network interfaces in the Cohesity cluster.
Latency	The average amount of time required for the Cohesity cluster to process incoming and outgoing external client requests. Clients might observe additional latencies caused by the network.

Metric	Description
Error	The number of data read and write errors generated in serving external client requests.
Session Setup	IOPS and latency for SMB client requests for a new authenticated session.
Logoff	IOPS and latency for SMB client requests to log off from the session.
Tree Connect	IOPS and latency for SMB client requests to request access to a share.
Tree Disconnect	IOPS and latency for SMB client requests to disconnect from a share.
Open	IOPS and latency for SMB client requests to open a file.
Close	IOPS and latency for SMB client requests to close a file.
Create File	IOPS and latency for SMB client requests to create a file.
Make Directory	IOPS and latency for SMB client requests to create a directory.
Delete File	IOPS and latency for SMB client requests to remove a file.
Remove Directory	IOPS and latency for SMB client requests to remove a directory.
Rename	IOPS and latency for SMB client requests to rename a directory.
Read Directory	IOPS and latency for SMB client requests to read a directory entry.
Get Information	IOPS and latency for SMB client requests to get information on a file.
Set Information	IOPS and latency for SMB client requests to set information on a file.

S3 Metrics

Metric	Description
IOPS	The average number of put, get, list, or delete operations from external clients serviced per second.
Throughput	The amount of data transferred through external network interfaces in the Cohesity cluster.

Metric	Description
Latency	The average amount of time required for the Cohesity cluster to process incoming and outgoing external client requests. Clients might observe additional latencies caused by the network.
Error	The number of put, get, list, or delete errors generated in serving external client requests.
Size	Sum of get and put data transferred through external client requests.

Client Connections Page

The **Client Connections** page provides details of all clients connected across nodes.

To view the **Client Connections** page:

1. Log in to the Cohesity cluster.
2. Go to **SmartFiles > Client Connections**:

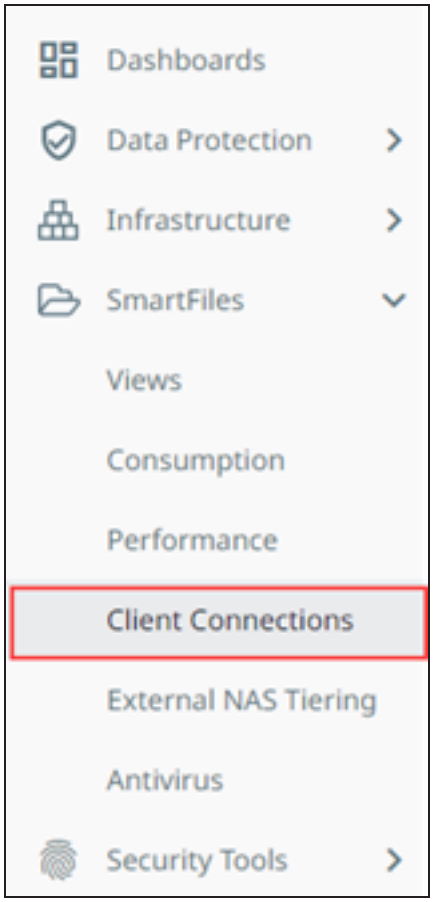

View Client Connection Client Details

The **Client Connections** page is displayed. The table supports the following filters:

- **Protocol**—Select **NFSv3** or **SMB**.

  > **Note:** In the 6.7 release, data about NFSv4.1 clients are not displayed on this page.

- **View**—Select a view to see the client(s) connected to the view.
- **Node**—Select a node to see the client(s) connected to the node.
- **Search Client IP**—Hover over the search icon to view the **Search Client IP** search bar. Enter the IP address of the client to filter the clients based on the IP address:

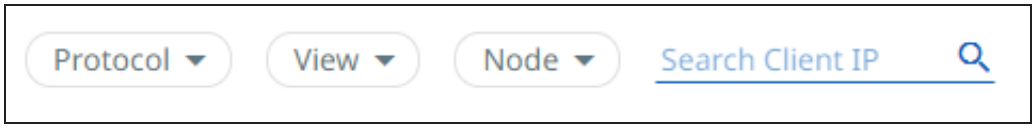

The following illustration shows an example of the **Client Connections** table:

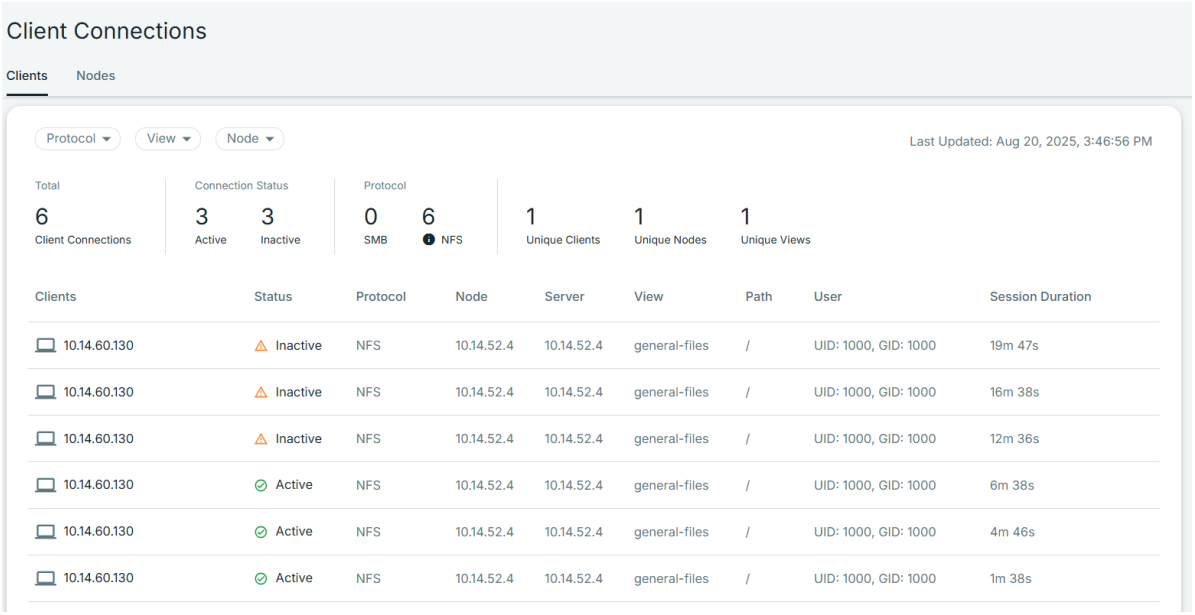

The table displays the following details:

- **Clients**—IP address of the client.
- **Protocol**—Protocol of the client: NFS v3 or SMB 3.0.
- **Node**—IP address of the node.
- **VIP**—Virtual IP address of the node.
- **View**—Name of the view.
- **Path**—Mount path details.
- **User**—User or group details.

- **Session Duration**—Duration for which the client is connected. You can sort this column according to the client session duration. You can sort the column and identify:
  - Clients that have been connected for the longest period of time
  - Clients that have been connected for the shortest period of time

View Client Connection Node Details

The table supports the following filter:

- **Protocol**—Select **NFS** or **SMB**.

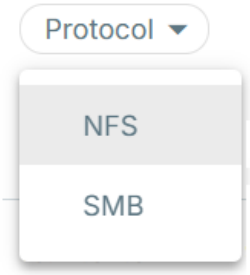

The following illustration shows an example of the Nodes table:

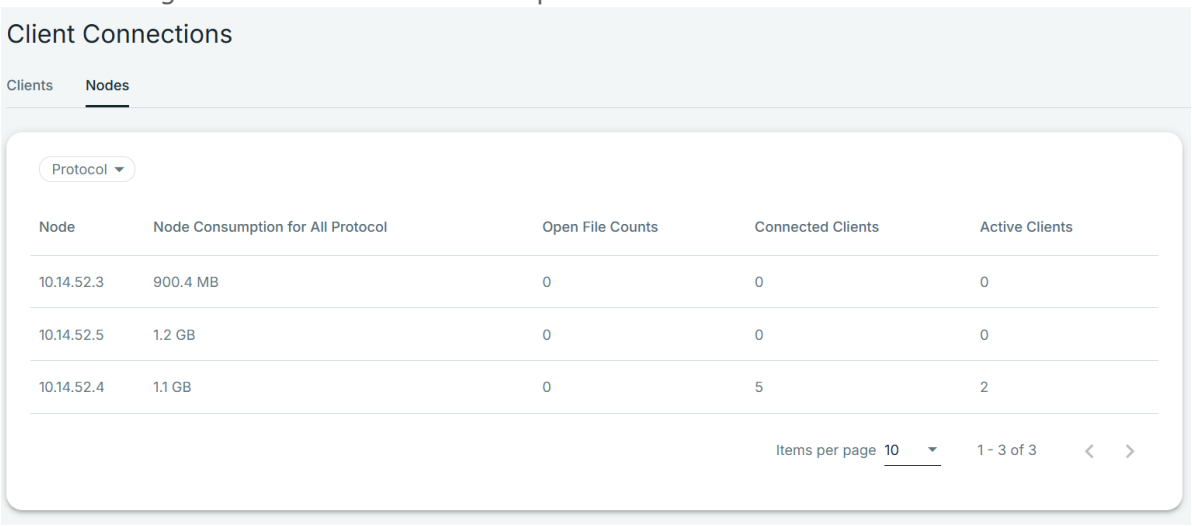

The table displays the following details:

- **Node** – IP address of the node.
- **Node Consumption for All Protocol** - Storage capacity consumed by the node for all the protocols.
- **Open File Counts** - Number of open files.
- **Connected Clients** - Number of connected clients.
- **Active Clients** - Number of active clients.

OpenStack Swift Storage Protocol

Cohesity supports the OpenStack Swift Storage protocol. Cohesity Swift services function as a third-party Swift implementation for the OpenStack environments.

You can create a Cohesity view, assign it to the Swift Only protocol, and then authenticated Swift clients can use Swift APIs to access the view as they would use an OpenStack container.

The OpenStack Keystone identity service is used to authenticate and authorize the Swift clients.

> **Note:** Cohesity cluster communicates with the Keystone Server using the identity API v3. So, the Keystone server must support the identity API v3. For more information, see OpenStack API Documentation.

You must add a Keystone server and configure default Swift options to use the Swift protocol when the Organizations feature is disabled.

> **Note:** You can manually add the Cohesity Swift service endpoint in the Keystone server for clients using the OpenStack command line to discover the Cohesity Swift service.

OpenStack Project and Domain Concepts

Project

Projects are an equivalent concept of tenancy in the Openstack environment. From Swift perspective, Containers tagged with the project ID can be given ownership to Users / User groups using the concept of roles. When you create a Swift Only view, you create a Swift container and tag that view to a Project in your Openstack environment. The view that you created is associated with the Project that you defined when creating the Swift Only view.

Domain

A Domain is a namespace within the Openstack environment. All the entities such as Projects, Users, and Roles belong to a namespace by default. In the OpenStack environment, a domain isolates one business entity from other business entities.

Swift Authentication and Authorization

Swift protocol authentication and authorization go through a Keystone server. All users are created and managed independently in the external Keystone server. Each user has at least one proper role and project assigned under some domain to operate on the Cohesity

Swift service. A user who is a Swift operator can create Swift Only View using cURL. All authorized Swift users can provide access to other users using Access Control Lists (ACL) based on the container owned by the Swift user. ACLs allow access to Swift containers/objects.

For example, a Swift user who owns a container A, can provide access to container A using ACL.

> **Note:**
> - Cohesity supports only Container level ACLs. For more info, see OpenStack Access Control Lists (ACLs).
> - Cohesity Swift service grants access to the users based on the ACL setting on the container when the user does not have the required roles assigned. A Swift container owner can provide access to the container to other authorized Swift users using ACL.

Keystone uses token-based authentication. Before doing any operation on Swift, a user must request a token from the Keystone server by providing its user credentials. If verification of user credentials succeeds, the user is assigned a token. Then the user could use the token to send a request to Swift service.

The Swift service authenticates itself with a Keystone server and obtains a token. To authenticate the user's token, Swift service sends the user token to the Keystone server. After the Keystone server verifies the user's token, Swift service will be able to grant authorization based on the information returned by the Keystone server.

Multitenancy Support (Organizations)

Cohesity Swift implementation supports multitenancy by assigning a Keystone config to each tenant.

For a multi-tenant Cohesity cluster (with the Organizations feature enabled), consider the following:

- You must add a Keystone server and complete Swift config to use the Swift protocol in both organizations and single tenancy environments.
- Assign the Keystone server to the Cohesity cluster and Swift user.
- If the Organization feature is enabled on the Cohesity cluster, you can have multiple Keystone servers and can be assigned to multiple tenants.

> **Note:** You cannot disable the Organizations feature once it is enabled.

- If the Organizations feature is not enabled on the Cohesity cluster, you can have only one Keystone server.
- You can assign only one Keystone configuration per organization.
- You must choose an existing Keystone config and complete Swift config for the required tenant if you want to use Swift service in the tenant.

Supported Swift APIs

The following tables indicate which Swift APIs are supported and unsupported by the Cohesity Swift service. For more information about all Swift APIs, see Object Storage API.

You can use Swift APIs to create Swift Only view from the Cohesity Dashboard and also using cURL commands. The Swift Only view can be created for both cluster with multitenancy and without multitenancy.

To use Swift APIs, you must first fetch a token from a Keystone server and use it in the Object Storage API.

Operation on the Discoverability

Operation	Supported
GET	NO

Operations on the Accounts

Operation	Supported
GET ACCOUNT	YES
HEAD ACCOUNT	YES
POST ACCOUNT	YES

Operations on the Containers

Operation	Supported
GET CONTAINER	YES
PUT CONTAINER	YES

Operation	Supported
DELETE CONTAINER	YES
POST CONTAINER	YES
HEAD CONTAINER	YES

Operations on the Objects

Operation	Supported
GET OBJECT	YES
PUT OBJECT	YES
DELETE OBJECT	YES
HEAD OBJECT	YES
POST OBJECT	YES
COPY OBJECT	NO

Operations on the Endpoints

Operation	Supported
GET ENDPOINT	NO

Following are the common functionalities supported by the above APIs:

- Pseudo-hierarchical folders and directories
- Page through large lists of containers or objects
- Serialized response formats: JSON

> **Note:** Serialized response formats in .XML are not supported.

Add a Keystone Server

The Keystone Server in the OpenStack environment provides secure controlled access to all of the OpenStack services. To create a Swift Only view, you must add a Keystone config for authentication and access management.

Prerequisites:

- Keystone Authentication URL
- Keystone admin credentials

To add a Keystone Server:

1. In the Cohesity Dashboard, select **Settings > Access Management**.
2. Click the **Keystone** tab.
3. On the Keystone page, click **Add Keystone Server**.
4. In the **Edit Keystone Server** page, complete the following:
    1. Type Keystone configuration name.
    2. Type or copy-paste the Keystone Authentication URL.
    3. Provide Admin Credentials.
        1. Admin Domain
        2. User name
        3. Password
    4. Provide Admin Scope, type Project name and Domain name.

    > **Note:** Admin scope is a configuration that defines Cohesity to get the admin token

5. Click **Add Keystone Server**.

Manage Swift Roles for the Single Tenancy

You can manage Swift roles for a single-tenant(default tenant) by adding or updating the Swift roles.

Swift roles are already defined roles on the Keystone server. These roles are used by Swift clients to send requests to a Cohesity cluster. When you add a Swift role, the Cohesity Swift service verifies that the role exists on the Keystone server.

To add or modify Swift roles:

1. In the Cohesity Dashboard, select **Settings > Access Management**.
2. Click the **Keystone** tab.

3. In the **Default Swift Options** section of the Keystone page, click the edit icon.
4. In the **Edit Default Swift Options** screen, add or delete Swift roles.
5. Click **Save**.

Register Cohesity Swift Service on OpenStack

Prerequisites

- You must add a KeyStone server and assign the Keystone server to the organization.
- Swift configuration must be complete for the organization.

You can register Cohesity Swift Service on the Keystone server for better integration of the OpenStack system.

To register the Swift Service on the Keystone server:

1. Log in to the Cohesity cluster.
2. Run the following command:

```
iris_cli
```

3. Run the following command:

```
swift register tenant-id=<tenant_id>
```

> **Note:**
> If the Organizations feature is not enabled on the Cohesity cluster, to register the swift service, use the following command:
>
> ```
> swift register tenant-id=''
> ```

As part of the registration process the following steps are performed on the Keystone server:

- The Cohesity cluster sends a request to the Keystone server to create a Region for the tenant. The Region name is a combination of the cluster-id and tenant-id.
- Ensure a Service Id is available for the Swift service on the Keystone server. If a Service Id is not available, then the Cohesity cluster sends a request to the Keystone server to create a Service Id.
- The Cohesity cluster sends a request to the Keystone server to create Endpoints for the Service Id in the created Region.

Unregister Cohesity Swift Service on OpenStack

You can unregister Cohesity Swift Service on the Keystone server.

To unregister the Swift Service on the Keystone server:

1. Log in to the Cohesity cluster.
2. Run the following command:

```
iris_cli
```

3. Run the following command:

```
swift unregister tenant-id=<tenant_id>
```

> **Note:**
> If the Organizations feature is not enabled on the Cohesity cluster, to unregister the swift service, use the following command:
>
> ```
> swift unregister tenant-id='
> ```

As part of the unregistration process, the Cohesity cluster deletes the Endpoints, Service Id, and Region on the Keystone server.

Create a Swift-Only View

A Swift Only view provides storage space for authorized Keystone users. The Swift Only view is tagged to a Project within a Domain. Only authorized Keystone users can access Swift Only views.

> **Note:** You cannot convert a view created with other protocols to Swift Only view later. Nor can you convert a Swift Only view to another protocol.

After creating a Swift Only view (container), you can use a Swift API compatible client or cURL to send requests to a Cohesity cluster. You must use the 3001 port to perform any operations supported by the Swift APIs. For example, the Swift client can be OpenStack CLI, Swift CLI, or cURL. You must log in to the Keystone server to use the Swift CLI and OpenStack CLI.

> **Note:** You must have OpenStack Object Storage Framework installed to send requests to the Cohesity cluster when using OpenStack and Swift command lines.

Prerequisites:

- Keystone Config and Swift config are configured before creating a Swift Only view.
- Project and Domain already exist on the Keystone server.

- Details of a Keystone Project and Domain for which you want to provide the view ownership.

To create a Swift only view:

1. In the Cohesity Dashboard, select **SmartFiles > Views**.
2. To create a new view, click **Create > View** at the top right of the page.
3. In the **Create View** page, type a name for a new **View**.
4. Select a Storage Domain or use the default Storage Domain.
5. Select the **Swift only** option for **View Protocol**.

   > **Note:** You cannot convert a View created with the All option to Swift Only View later.

6. Provide **Project** name and **Domain** for the View ownership.

   > **Note:** When you provide the Swift Ownership, you are tagging the View to an existing Keystone Project.

7. (Optional) Provide a Keystone username for the Swift Ownership.

   > **Note:** This field will be used to set the initial container level ACL.

8. Select **Object Key Pattern:**This option cannot be changed after the view is created. For more information, see Object Keys.
9. **Case Sensitive:** Toggle on if you want the uppercase and lowercase versions of the same letter in a file or folder name to be treated differently on this view, for example: `A` is different from `a`.

   You cannot edit this setting after you create the view.

   Optional—Cohesity recommends entering a description for the view, for example, the reason why the view is being created.
10. Provide a description for the View.
11. Select QoS Policy for performance. **QoS Policy:** A QoS Policy determines to which storage media data is written and the priority of I/O when contention occurs. Select an appropriate QoS policy for the view. By specifying a priority and the type of workload you expect, the Cohesity cluster can optimize how it processes data in regard to latency, throughput, and priority.For details and recommendations, see QoS Policies for Cohesity Views.

12. **Security:**
    - **Override Global IP Allowlist:** Turn on **Override Global IP Allowlist** if you want to specify a custom list for this view. Only the systems with IP addresses in the subnets listed in the view's allowlist have privileges to access and mount this view. If you enable the **Override Global IP Allowlist** toggle, the machines listed in the subnets of the global Cohesity cluster allowlist cannot access or mount this view unless the IP address of the machine is also listed in the View's Allowlist.

      Extend Global IP Allowlist: Allows you to extend the Global IP Allowlist.

      To add subnets to the allowlist of this view:

      a. Under **IP Allowlist**, select: **Override Global IP Allowlist** or **Extend Global IP Allowlist**.
      b. Click **Add**.
      c. In the **Add Allowlist** screen, type subnet in CIDR format.
      d. Select any or all permissions as required.
      e. Select **NFS Squash** options.

      The **NFS all squash** option is like NFS Root Squash. However, for given IP addresses, all users are mapped to the All Squash User/Group configured in the view. The default All Squash User for all views is "nobody".

      > **Note:** If the **Override Global Allowlist** toggle is off, this view can only be accessed and mounted by systems specified by the global Cohesity cluster allowlist.

13. **Dedupe & Compression:** Toggle on to always use post-processing, irrespective of storage domain policy. This applies only when the storage domain has enabled deduplication or compression.
14. **Quota:** If applicable, the view inherits quota and alert settings from its storage domain by default. You can optionally **Override** these.

    Optionally toggle on **Logical Quota** and set the view's logical data capacity limit. (Logical data is fully hydrated and expanded.) Capacities must use whole numbers only.

    You can also choose to receive an alert when the view's capacity reaches a specified threshold.

    (After the view is created, you can define user quotas.)
15. **File DataLock:** Toggle on to make the files in this view immutable based on the settings described below. File level DataLock is typically used for compliance and

regulatory purposes.

> **Note:** The File DataLock option is only applicable for NFS and SMB View types.

- **File DataLock Mode:** Select one of the following modes:

  Compliance: Once File DataLock is enabled for the view, it cannot be disabled by anyone. However, a Cohesity user with Data Security privileges can edit the other DataLock settings. Compliance mode means the view complies with SEC Rule 17a-4.

  Enterprise: A Cohesity user with Admin privileges or Data Security privileges can rename the view, delete the view, switch to Compliance mode, and edit the File DataLock settings. This mode is useful for giving storage administrators additional control over the view.

  > **Note:** Enterprise mode is supported on 6.4.1 and later Cohesity clusters. If the view is replicated to a pre-6.4.1 Cohesity cluster, Enterprise mode is automatically changed to Compliance mode. If the view is then replicated back to a 6.4.1 or later Cohesity cluster, it retains the Compliance mode setting.

- **Default Retention:** Indicate the default time period that locked files remain locked if a user or application does not specify a retention default. When the time period ends, files can be deleted but not changed.
- **Autolock Files:** Indicate whether files are automatically locked after the specified period of I/O inactivity (based on the last modified time stamp).
- **Manually Lock Using:** By default, manual file locking occurs when a user or application sets the file's read-only attribute. Click the edit icon to set the following:

  Select **Read-Only** to lock files when the file permission is set to read only, for example, by setting the Windows file properties attribute or using the Linux `chmod` command.

  Select **FutureATime** to lock files when the last access time (atime) is set, for example, by using Windows PowerShell commands or the Linux `touch` command.

  Optionally set a **Minimum Retention** and/or a **Maximum Retention** period to ensure that a file's last access time cannot be set outside this period. A request to do so results in a permission denied error if manual locking mode is

FutureATime. Otherwise, the request will be silently rejected and the file will be locked for the minimum retention duration.

- **Override - Lock all files until:** Indicate whether the lock on files expires on a specific date. This setting is typically used for legal hold purposes. Files in a lock expired state cannot be deleted until this date.

  In **Compliance** mode, a user with Data Security privileges can extend the override date into the future, but no one can shorten the date.

  In **Enterprise** mode, a user with Admin or Data Security privileges can extend or shorten the override date.

16. **File System:** Toggle on **File Filtering** if you want to block or allow files with the specified file extensions from being written to this view. Select one of the following options:
    - Select **Denylist** if you want to specify a list of file extensions for files that are blocked from being written to this view. For example, if you specify a list with the `.exe` and `.dll` extensions, files with these extensions are blocked from writing to this view.
    - Select **Allowlist** if you want to specify a list of file extensions for files that are allowed to be written to this view. For example, if you specify a list with the `.cvs` and `.txt` extensions, only files with these extensions can be written to this view.
      - Enter file extensions separated with commas, spaces, tabs or newlines and click **Add**
17. Click **Create View**.

Assign a Keystone Config to an Organization

After you create an Organization on the Cohesity cluster, you must map a Keystone config to use the Cohesity Swift service in the Organization.

> **Note:** You can also assign a Keystone config when adding a new Organization.

To map a Keystone Server to an Organization:

1. In the Cohesity Dashboard, select **Settings > Access Management**.
2. In the **Organizations** page, select a required organization from the list.
3. In the selected organization page, click the edit pencil icon on the top-right corner of the screen.
4. In the **Edit Organization** page, click the **Assign AD/LDAP/Keystone** link.
5. In the **Assigning AD/LDAP** page, the **Keystone** tab displays the registered Keystone servers. Select one Keystone server and click **Assign**.

6. In the **Swift Roles** section of the **Edit Organization** page, click the pencil icon. In the **Edit Default Swift** options, add or modify the required Swift roles and click **Save**.

   > **Note:** You can only add the roles defined on the OpenStack server.

7. To further isolate network traffic for the organization, assign VLAN IDs. VLANs can be added and removed later. Using VLANs is optional but recommended.

   > **Note:** All the swift requests whose tenancy cannot be determined will fail.

   > **Note:** If the **Enable for all Organization** option is selected for a VLAN and the VLAN is assigned to an Organization, the Swift clients cannot access Cohesity Swift Service through this shared VLAN.

8. Click **Save**.

Accessing Swift Services

After you configure a Keystone config on the Cohesity cluster and complete the configuration steps, you can access the Cohesity Swift service by using Swift APIs.

You can use a Swift API compatible client to send requests to a Cohesity Swift service by using the 3001 port to perform any operations supported by Swift APIs. For example, the Swift client can use OpenStack, Swift CLI, or cURL. For more information on Swift APIs, see Object Storage API.

You can use the following endpoint format with the Swift APIs:

Endpoint format:

https://<host>:3001/v1/AUTH_<project_id>

The host can be any of the following:

- Cohesity cluster hostname
- Cohesity cluster node IP
- VLAN hostname
- VIP

> **Note:** To use Swift APIs, you must first fetch a token from a Keystone server and use it in the Object Storage API. For more information, see API Examples using Curl.

For information about creating a Container on a Cohesity cluster API, see Object Storage API.

Examples:

Request a token from a Keystone server (10.1.22.33)

```
keystone_server='10.1.22.33'

curl -w "\n" -i -H "Content-Type: application/json" -d '
{ "auth": { "identity": {
    "methods": [ "password" ],
    "password": { "user": {
      "name": "admin",
      "domain": { "id": "default" },
      "password": "bd4dd51646714fd3"
    }}},
    "scope": { "project": { "domain": { "id": "default" }, "name": "admin" } }
  }
}' "http://${keystone_server}:5000/v3/auth/tokens"
```

```
Response (body is omitted)
```

```
HTTP/1.1 201 CREATED
Date: Thu, 02 Apr 2020 21:10:49 GMT
Server: Apache
X-Subject-Token: gAAAAABehlTavUkjXGlOndMO6S8C5XnuIF7ihxzRP0kQNdKAV0MuX6mUS-
YiASZRg2KLC-XTxhtTtslfnDBd-
ya0DpRbAmfxd9LPcroYKifqZFTnDCR4M1rKwErm11nhWDwcPcvmodArBITB6bRH2PvYNNZ1W0u-
NH_aKPJsuJWJvrAL6hBy2gs
Vary: X-Auth-Token
x-openstack-request-id: req-2eaa860c-b592-44f0-aa5e-9cb7e066f67f
Content-Length: 7627
Content-Type: application/json
```

Create a container named container0

Request

```
cohesity_server='10.2.33.44'
token='gAAAAABehlTavUkjXGlOndMO6S8C5XnuIF7ihxzRP0kQNdKAV0MuX6mUS-YiASZRg2KLC-
XTxhtTtslfnDBd-
ya0DpRbAmfxd9LPcroYKifqZFTnDCR4M1rKwErm11nhWDwcPcvmodArBITB6bRH2PvYNNZ1W0u-
NH_aKPJsuJWJvrAL6hBy2gs'
project_id='b7aa1b1c21784ac1bfc0ddc9815c27b5'
container_name='container0'

curl -k -i -X PUT -H "X-Auth-Token: ${token}"
"https://${cohesity_server}:3001/v1/AUTH_${project_id}/${container_name}"
```

Response

```
HTTP/1.1 201 Created
Server: CohesitySwift
Transfer-Encoding: chunked
Date: Thu, 02 Apr 2020 21:01:56 GMT
Connection: keep-alive
```

Put an object named object0 in the container container0

Request

```
object_name='object0'
curl -k -i -X PUT -d "abcdef" -H "Content-Type: text/html; charset=UTF-8" -H "Content-
Length: 6" -H "X-Auth-Token: ${token}"
"https://${cohesity_server}:3001/v1/AUTH_${project_id}/${container_name}/${object_name}"
```

Response

```
HTTP/1.1 201 Created
Server: CohesitySwift
Transfer-Encoding: chunked
Date: Thu, 02 Apr 2020 21:08:47 GMT
Connection: keep-alive
```

Edit a Swift only View

You can edit a Swift Only view and make changes to the view properties. You can update the following details in the View properties:

- View name
- Swift Ownership (Only project name and domain name can be updated.)
- Advanced options

To edit the Swift only view:

1. In the Cohesity Dashboard, select **SmartFiles > Views**
2. In the Views tab, click ⋮ and select **Edit**.
3. In the **Edit View** page, you can edit the following:
    - View Name
    - Swift Ownership
    - Description
    - Advanced Options
4. Click **Update View**.

Edit Keystone Config

You might need to edit the Keystone config details to update the authentication path, name of the Keystone config, admin credentials, or admin scope.

Points to consider when editing the Keystone config:

- When modifying the Cohesity Keystone config, the Cohesity Swift service verifies all the config contents specified in the Keystone and Swift by communicating with the Keystone Server. If the verification fails, the Swift configuration creation or modification fails.
- Similarly, when modifying a Keystone config shared by multiple organizations, all the roles of all the organizations sharing this Keystone config will be verified through this modified Keystone config. If verification fails, the Keystone config modification also fails.

To edit Keystone Server details:

1. In the Cohesity Dashboard **Settings > Access Management**.
2. Click the **Keystone** tab.
3. In the **Keystone** tab, click the Keystone config name link under the **Config Name**.
4. On the **Edit Keystone Server** page, you can edit and update the following:

- Keystone Configuration Name
- Authentication URL
- Details in the Admin Credentials

> **Note:** The Cohesity Dashboard does not display the current password. If you do not want to change the password, leave the field blank. If you want to change it, type a new password.

- Details in the **Admin Scope**.

5. Click **Save**.

Delete a Keystone Server

You can delete a Keystone config from the Cohesity cluster if you decide to stop using the Cohesity Swift service.

Once you delete the Keystone config, clients cannot access the Cohesity Swift service anymore, including the views and objects in the view.

To delete a Keystone server:

1. In the Cohesity Dashboard **Settings > Access Management**.
2. Click the **Keystone** tab.
3. Click the delete icon for the required **Keystone Server**.
4. In the **Delete Keystone Server** screen, provide the Keystone Admin Password.
5. Click **Delete**.

NAS Tiering

In primary storage devices such as NAS, there is unused or infrequently used data that remains relevant to the organization or must be retained for future reference or regulatory compliance reasons. As a result, the growing data takes more space in primary storage, which has a higher cost than secondary storage.

Cohesity supports the tiering of unused or infrequently used data from NAS primary storage to the Cohesity cluster. The process of moving unused data from the primary NAS storage to the Cohesity cluster is called downtiering. Based on a tiering policy the administrator sets while creating a data tiering task, the data from the NAS primary storage that is unused or infrequently used is periodically moved to the Cohesity cluster. When creating a data tiering task, the administrator determines how frequently the data tiering task runs and sets a tiering policy based on various parameters, such as:

- The last time the file was accessed
- The last time the file was modified
- The size of the file
- The file type to be allowed or denied for downtiering

> **Note:**
> - Cohesity does not support updating last access time (atime) on files belonging to a Cohesity view. Use caution when tiering data from a Cohesity View.
> - Data tiering is supported only for NAS storage that is registered as mount points of type SMB or NFS.
> - Mac clients are not supported to access files via symbolic links.
> - You must add all sources as **Generic NAS** only.

> **Important:** Uptiering is not supported. By default, orphan data deletion is disabled.

SmartFiles MTIME/CTIME Support for NAS Tiering

- **Tiering Policy Inputs**. NAS tiering jobs can use the last modified time of the file (MTIME), with other attributes such as file size and type, in the policy that determines the files to downtier from primary NAS to SmartFiles.
- **Change Detection**. SmartFiles file walkers and backup logic check both MTIME and CTIME to detect if the changes are made to file data or metadata. This ensures that updates affecting only attributes such as permissions or ownership are not missed in cases where MTIME remains unchanged.
- **Preservation of Tiered Copies**. During data tiering, the Cohesity filesystem (SnapFS) preserves the file metadata, including MTIME and CTIME, on the tiered copy. These timestamps are then used in the subsequent scans to avoid re-tiering or skipping modified data. Backups and tiering share this change-detection logic.

When data is tiered from the primary storage, a symbolic link is created for the original file. The name of this symbolic link is the same as that of the original file, ensuring that users and applications continue to access the data as before. On a Windows server, local-to-local symbolic links and local-to-remote symbolic links are enabled. However, remote-to-local symbolic links and remote-to-remote symbolic links are disabled by default. Therefore you must run the following commands on the Windows server to enable the symbolic link access:

```
C:\Windows\system32>fsutil behavior set SymlinkEvaluation R2R:1
```

```
C:\Windows\system32>fsutil behavior query SymlinkEvaluation
```

If you are permanently moving the unused data to the Cohesity cluster, you can choose not to create this symbolic link after tiering while creating a tiering task.

For NFS shares, for the user or applications to access the data that was tiered to the Cohesity View, the View needs to be mounted on the NFS client. The mount path where the Cohesity View will be mounted on the host must be provided when you create the data tiering task. This mount path cannot be changed later. All the NFS clients should use this mount path to mount the Cohesity View. For more information on mounting a view using NFS, see Mount a View Using NFS.

Once the data is tiered to the Cohesity cluster, the subsequent read and write operations of the tiered data are performed on the Cohesity server hosting the data.

A user with an admin role also can restore the tiered data back to the primary storage by mounting the Cohesity target (server hosting the data) view on a Windows client and then copying the data.

Supported Sources for NAS Downtiering

The following is the support matrix for NAS downtiering:

Source	Protocol	Without Stubbing	With Stubbing
Unix/Linux	NFSv3	Yes	Yes
Windows	SMB	Yes	Yes
	NFSv3	Yes	Yes
Cohesity	SMB	Yes	Yes
	S3	No	No
	MultiProtocol	No	No
Isilon	NFSv3	Yes	Yes
	SMB	Yes	Yes
	S3	No	No
	MultiProtocol	No	No

Source	Protocol	Without Stubbing	With Stubbing
NetApp	NFSv3	Yes	Yes
	SMB	Yes	No
	S3	No	No
	MultiProtocol	No	No

Backups and Access to Older Versions of Tiered Files

NAS downtiering is not a replacement for backing up primary NAS storage.

Changes made to files in the target are not visible to the primary storage. Consequently, Cohesity recommends taking snapshots of the target storage and replicating them, similar to the backup process for any SmartFiles content on Cohesity. You cannot restore these files from snapshots on the primary storage but can recover them by mounting the target view.

Prevent Unintentional Access to the Target View

For symlinks to work, there needs to be a connection from the user's systems to the target. To facilitate backup and restore, the target view is available in SmartFiles. You must choose a view name that indicates its purpose is solely for file recovery. If users mount the target and delete files, the symlinks on the source will become inactive.

Considerations

- Cohesity does not support downtiering of junction points present on the NAS source.
- If the downtiering task was created in a version prior to 7.0 and if the Cohesity cluster is upgraded to 7.0, the downtiering task is based on the old policy. You cannot edit the downtiering task. If you want to edit the downtiering task, Cohesity recommends doing the following:
    - Delete the previously created downtiering task
    - Create a new downtiering task
- If the downtiering task was created in a version prior to 7.0 and if the Cohesity cluster is upgraded to 7.0, the downtiering task is based on the old policy. You cannot edit the grouping of this downtiering task. If you want to edit the grouping, Cohesity recommends doing the following:

    - Delete the previously created downtiering task
    - Create a new downtiering task and edit the grouping
- If SMB permissions are modified on the original source, these changes are not fetched and applied to the data tiered on Cohesity.
- Avoid utilizing NAS downtiering for home directories or datasets with active user modifications. If used for such cases, consider the following:
    - When using **Save As** within certain applications after accessing the symlink, the new file is automatically written to the target instead of the source by default. Consequently, when searching for such files later, you must check the target location.
    - Installer programs that delete themselves upon completion may erase the file in the target, rather than removing the symbolic link on the primary. This results in a non-functional link remaining on the primary system.
    - Users with an older version of Windows 10 that does not support long path names might encounter difficulties when mass copying files from the downtiered target for tasks such as restoring older versions. To address this, a newer system that supports long path names may be necessary to easily copy the files back.
    - Deleting the symbolic link will not automatically free up space on Cohesity.

      Contact Cohesity Support to free up space.
    - Before 6.8.1_u7, commands existed to uptier and delete orphaned files — files whose symlinks were deleted or renamed by the user. If a user moves or renames one of the Cohesity-created symlinks, the `delete orphaned` command removes the file from Cohesity, and the `uptier` command restores the file under its old name, potentially overwriting a new file that reuses the original symlink name. To mitigate this risk, `uptier` and `delete orphaned` commands were removed in 6.8.1_u7 or later versions. If link alterations may have occurred, contact Cohesity Support to run diagnostics to identify and address such issues.
- Cohesity does not recommend downtiering of nested shares.
- Downtiering symlink source files without the stub option may result in dangling symlinks.
- NAS Tiering is not supported in multitenant environments, and organization administrators do not have the ability to configure NAS Tiering.
- You cannot modify the File Level DataLock on the tiering target view.
- The File Level DataLock configuration will be removed from the destination cluster's view when replicating a tiering target view.

- The **Tiering Audit** option is enabled by default starting from version 7.2.2, and audit logs for tiering jobs will be recorded. For clusters running version 7.2.1 or earlier, you must manually enable tiering audit logs.
- You cannot configure Lifecycle Management on a tiering target view.
- If the source is Dell EMC Isilon, orphan data deletion (ODD) is not supported.
- If a directory is deleted on the source and ODD is enabled, the corresponding directory on the target is deleted once it is empty. This process may take multiple runs to complete.
- Deletes and renames require multiple runs to complete.
- Cohesity atime is not updated each time a file is accessed; it is updated based on the file modification time.

NAS Tiering Dashboard

The **NAS Tiering** page consists of the following two sections:

- Capacity Trend
- Tiering by Source

Capacity Trend

The **Capacity Trend** section provides a graphical representation of the data on the source and the data down-tiered to the Cohesity cluster. The **Capacity Trend** histogram enables you to generate periodic analysis and predict the capacity needs. You can:

- Select a range—Past 12 Hours, Past 7 Days, Past 30 Days, or Past 90 Days.
- Filter the data by source if multiple sources are registered.
- Click on the legends to add or remove data displayed on the histogram. For example, if you click Hot, data pertaining to hot is removed from the histogram and if you click Hot again, the data is added back to the histogram.
- Hover over the bar graph to the Data on Source and Data on Cohesity for a particular day.
- Analyze the cost savings for the selected time range.

Tiering by Source Pane

The **Tiering by Source** pane lists the NAS sources that are registered with the Cohesity cluster for tiering. You can analyze the data available on the registered source by clicking **Analyze**. After analysis, click **Plan Tiering** to set a tiering policy, and based on this policy you can create a tiering task:

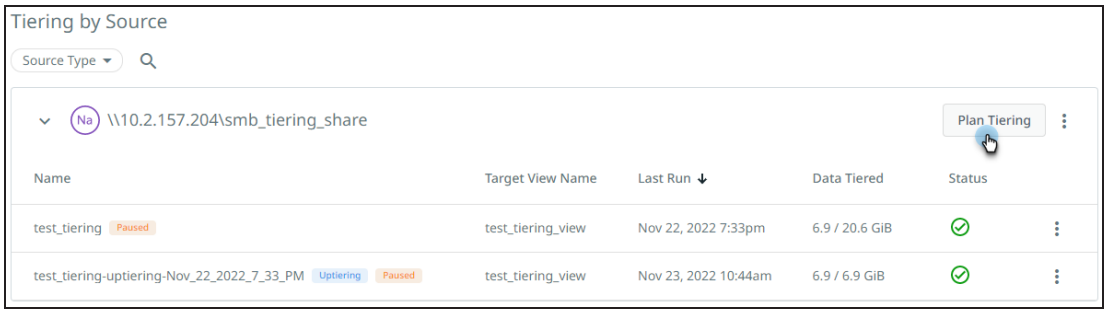

Once you execute the tiering of a registered source, you can click the arrow icon next to the registered source to view the following details:

- **Name**—The name of the tiering task. Click on the link to drill down and view more information about the tiering task such as runs, run details, and audit trail.
- **Target View Name**—The name of the View to which the data is tiered.
- **Last Run**—The last time the tiering task was run.
- **Data Tiered**—The total number of bytes tiered during the tiering task.
- **Status**—The status of the tiering task.
- Actions menu—To manage an existing tiering task, click the Actions menu (⋮) corresponding to the task and select one of the following options:
  - **Run Now**—Instantly initiates the tiering task.
  - **Pause Future Runs** or **Resume Future Runs**—Pause or resume the future tiering task runs.

- **Edit**—Configure the settings of a tiering task. For more information, see Downtier Data.
- **Delete**—Deletes the tiering task. The data that was tiered before deleting the tiering task will be stored in the Target View hosted on the Cohesity cluster. You can continue to access this data through symlinks created during the down-tiering task.

View Data Tiering Task Details

You can view the details of each tiering task by clicking the view name:

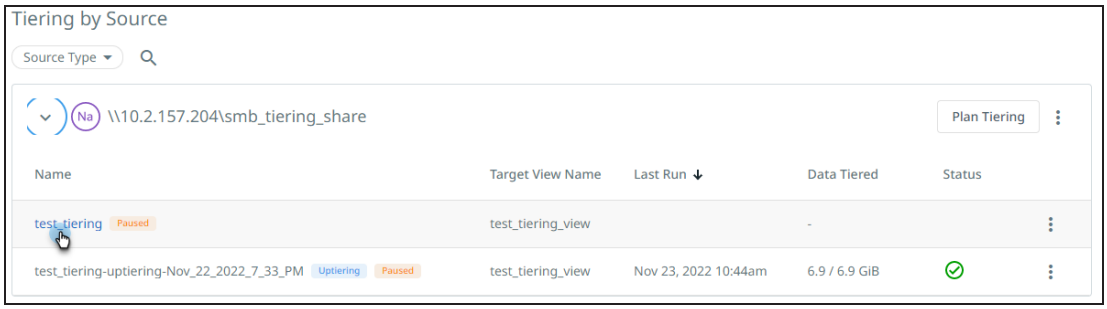

From the tiering details page, you can do the following to manage the data tiering run from the **Runs** tab:

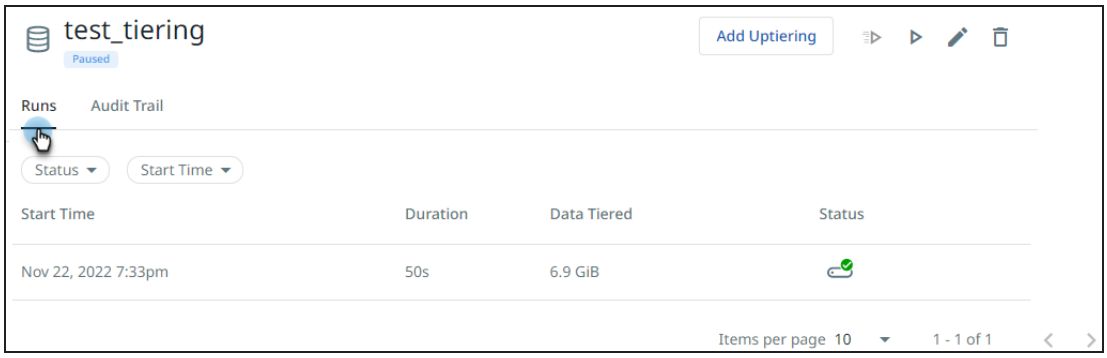

- **Run Now**—Instantly initiates the tiering task.
- **Pause Future Runs** or **Resume Future Runs**—Pause or resume the future runs of the tiering task.
- **Edit**—Edit the configuration settings of a tiering task. Follow the instructions provided in the Plan Tiering section.
- **Delete**—Deletes the tiering task. The data that was tiered before deleting the tiering task is stored in the Target View hosted on the Cohesity cluster. You can continue to access this data through symlinks created during the down-tiering task.

Click the data tiering run to navigate to the **Run Details** page. This page includes a summary bar, a filter option, and a table. The following is an illustration of the summary bar:

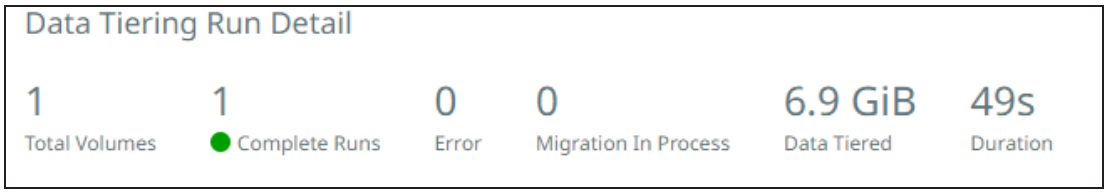

- **Mount Path**—The path used for mounting the NAS.
- **Target View Name**—The name of the View to which the data is tiered.
- **Start Time**—The date and time when the tiering run started.
- **Duration**—The time (in seconds) taken for the tiering run to complete.
- **Data Tiered**—The total number of bytes tiered during the tiering run.
- **Status**—The status of the tiering run.
- **Delete Run Snapshot**—After the tiering run is complete, you can delete the snapshot for a specific run of a tiering task without deleting the entire task.

Click the **Audit Trail** tab to view a sequential record of the events:

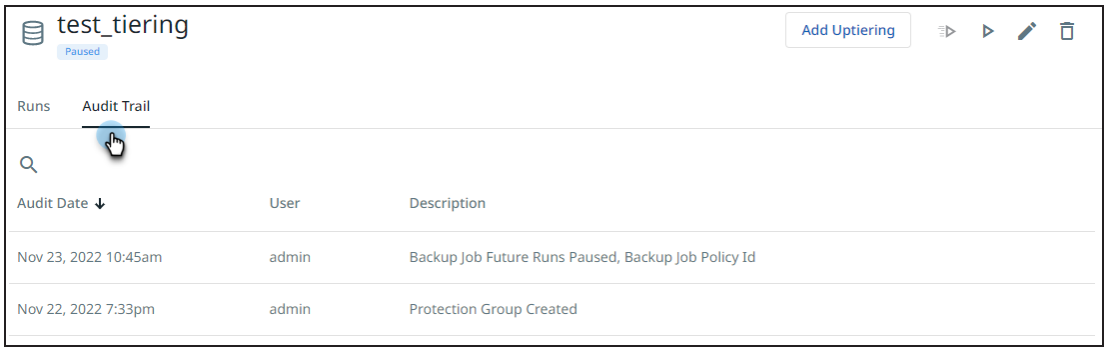

- **Audit Date**—Date and time of the event.
- **User**—Username of the person who initiated the event.
- **Description**—Description of the event.

Downtier Data

Downtiering involves the following steps:

1. Register the NAS Source
2. Analyze the Source
3. Plan Tiering
4. Configure a Down-Tiering Task

Register the NAS Source

To down-tier data from the NAS primary storage to the Cohesity cluster, you must first register the NAS as a source with the Cohesity cluster.

To register NAS:

1. Navigate to **SmartFiles > NAS Tiering**.
2. Click **Add** in the upper-right corner of the page and select **Register New Source**.

   The **Register NAS** dialog is displayed and **Generic NAS** is automatically selected in the **Type** field:

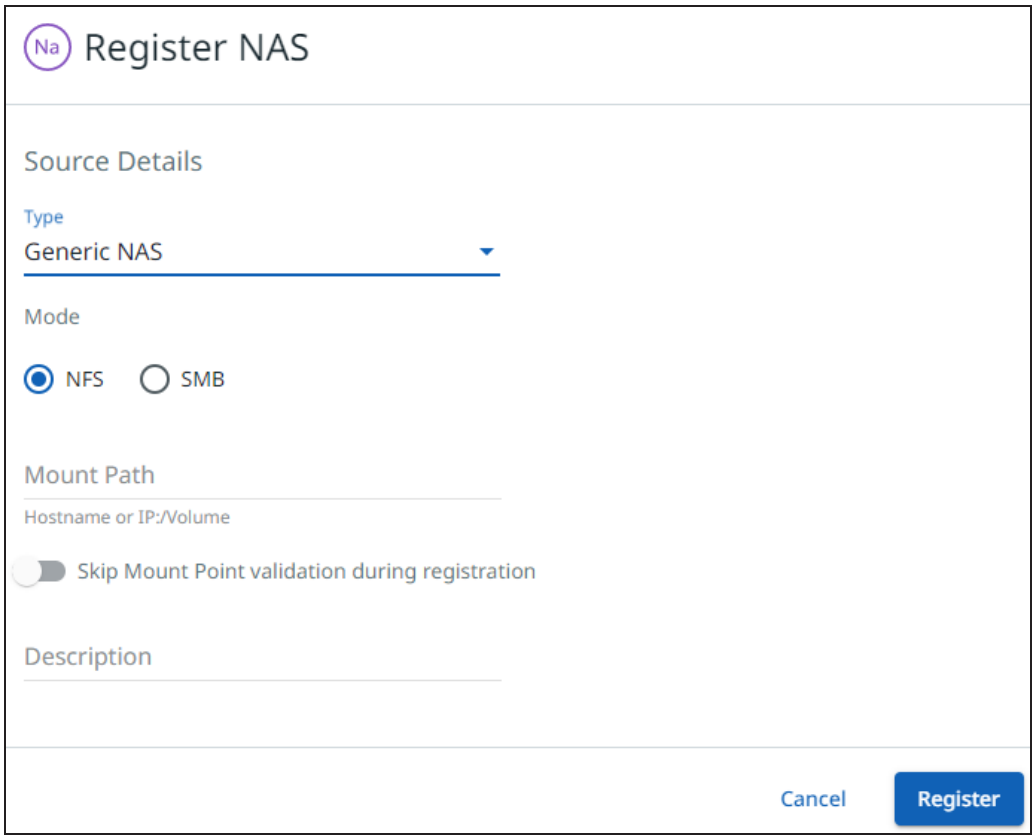

3. Select the Mode—**NFS** or **SMB**.
4. If you select **SMB**, provide the following details:
   - **Username**—SMB mount credentials that allow reading, writing, and backing up of files.
   - **Password**—SMB mount credentials that allow reading and backing up of files.

     > **Note:** For downtiering, the user must have read, write, and backup permissions.

5. **Mount Path**—The path for mounting the NAS. Use the following format:
   - **NFS**—Hostname or IP:/Volume
   - **SMB**—\\Hostname or IP\Share path

   > **Note:** Do not edit an existing mount path. If you need to change the mount path, register the mount path as a new source.

6. Optionally, you can skip mount point validation during registration.
7. **Description**—Optionally, you can enter a description for identifying the mount path.
8. Click **Register**.

   Once the registration is complete, the NAS source appears on the NAS Tiering page. You can now analyze the data on the source.

Analyze the Source

Once you register the source, you can analyze the data on the source. After the analysis is done, by default, the data on the source is classified on the basis of the following categories or parameters:

- **File Size**—Files are organized into groups based on the size of the file.
- **File Type**—Files are organized into groups based on the file type such as text, audio, and database.
- **Date of Last Access**—Files are organized into groups based on the last date on which the data was accessed.

When you plan to tier, the Cohesity cluster displays the histogram based on the groups of these categories. You can choose to view the histogram either based on the file type or file size, and the data is displayed against the date of the last access grouping (x-axis). To customize the buckets of these categories, see Edit Grouping.

To analyze the source, on the **NAS Tiering** page, click **Analyze** for the source you registered.

The analysis process will be initiated. The duration of the analysis may vary based on the number of files and directories in the source, and it may take some time to complete.

Refresh Analysis

To refresh the analysis done on the data of the source:

1. On the **NAS Tiering** page, go to the **Tiering by Source** pane.
2. Click the Actions menu (⋮) corresponding to the source you analyzed:

3. Click **Refresh Analysis**.

   The data on the source is analyzed again.

Edit Grouping

Once the initial analysis is complete, you can optionally customize the grouping within each category. By default, the data is grouped into **Hot**, **Warm**, or **Cold**. Once you customize the existing groups, a new analysis task is triggered. This analysis is performed on the basis of the new grouping. When you plan to tier, the Cohesity cluster will display the histogram based on the updated grouping.

To edit the grouping:

1. On the **NAS Tiering** page, go to the **Tiering by Source** pane.
2. Click the Actions menu (⋮) corresponding to the source you analyzed:
3. Click **Edit Grouping**.

   The **Edit Grouping** dialog is displayed:

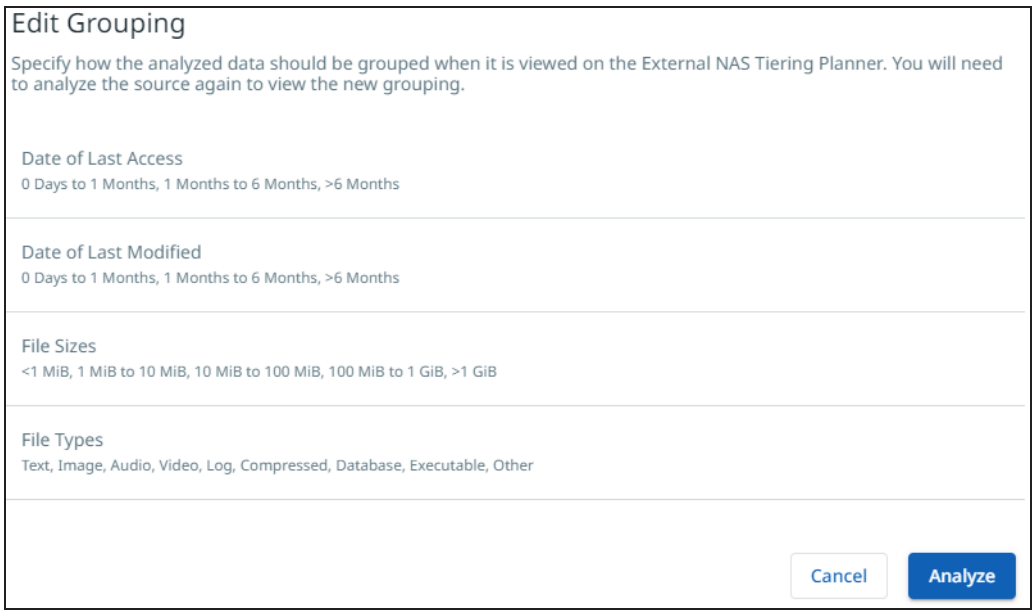

4. Hover over the group that you want to update, click the edit icon, and update grouping:

- Date of Last Access
- Date of Last Modified
- File Sizes
- File Types

5. Click **Analyze**.

Date of Last Access

The data accessed within the defined range are grouped and displayed on the **Select Data To Tier** page. To update a group, you can do one of the following:

- Click on the defined range and update the value
- Move the slider and update the value

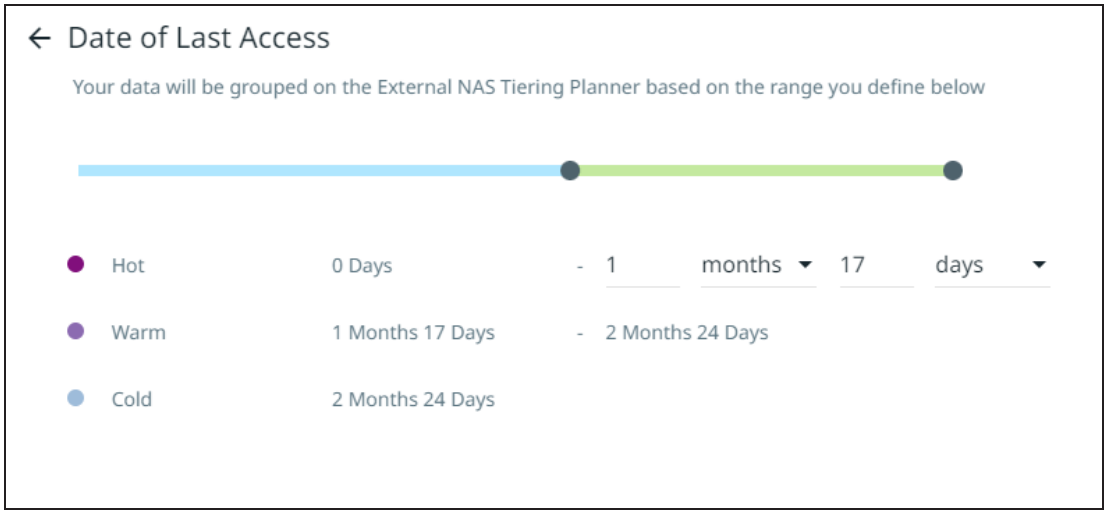

Date of Last Modified

The data modified within the defined range are grouped and displayed on the **Select Data To Tier** page. To update a group, you can do one of the following:

- Click on the defined range and update the value
- Move the slider and update the value

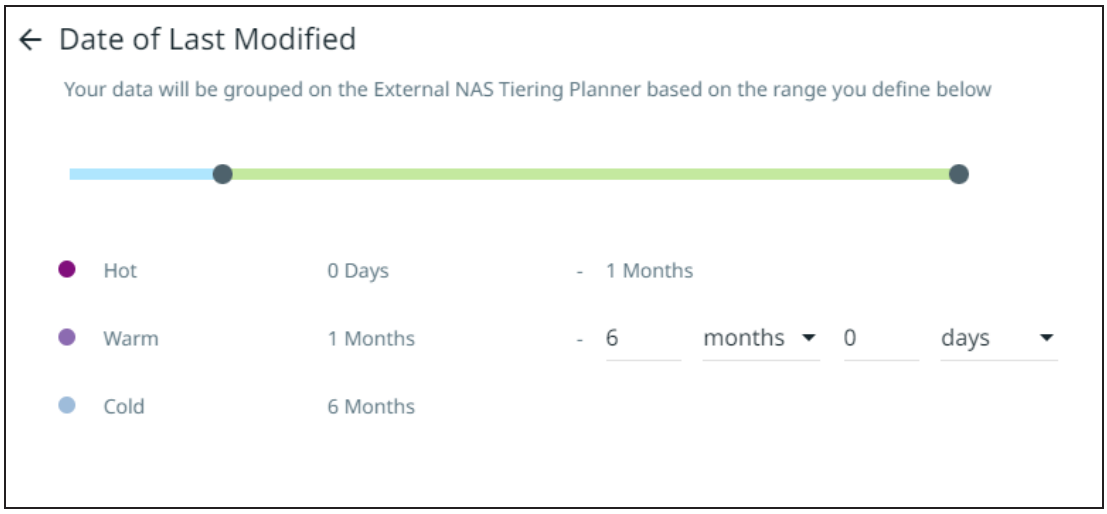

File Sizes

The files within this defined file size range are grouped and displayed on the **Select Data To Tier** page. To update a group, you can do one of the following:

- Click on the defined range and update the value
- Move the slider and update the value

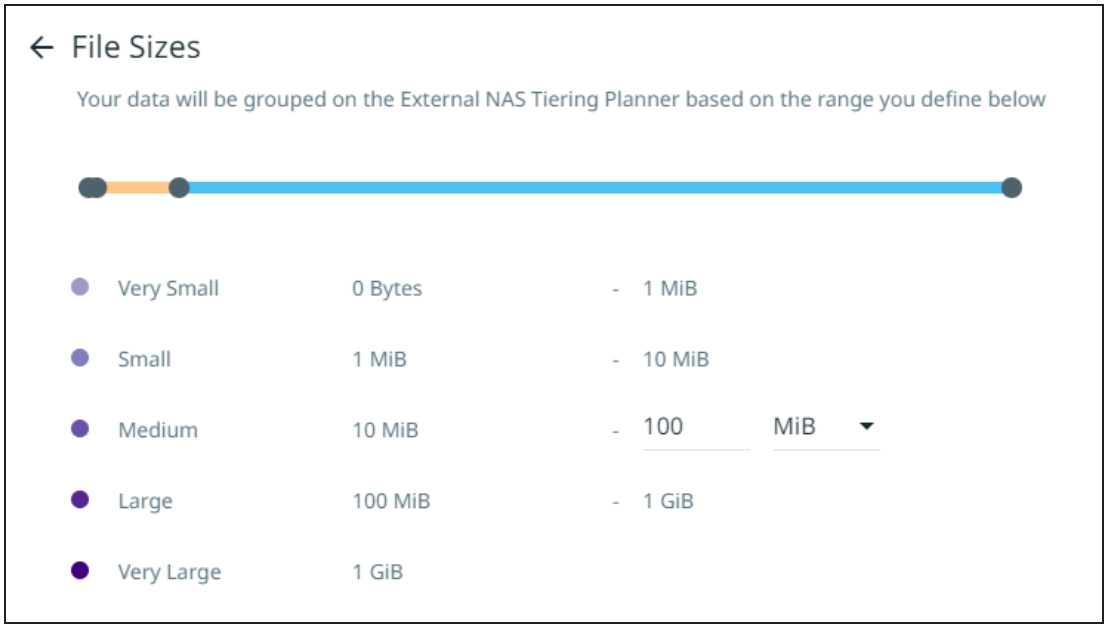

File Types

The files based on the category are grouped and displayed on the **Select Data To Tier** page. Click on a category to enter the extension of the file format you want to group under this category. You cannot enter an extension that is already a part of another file category.

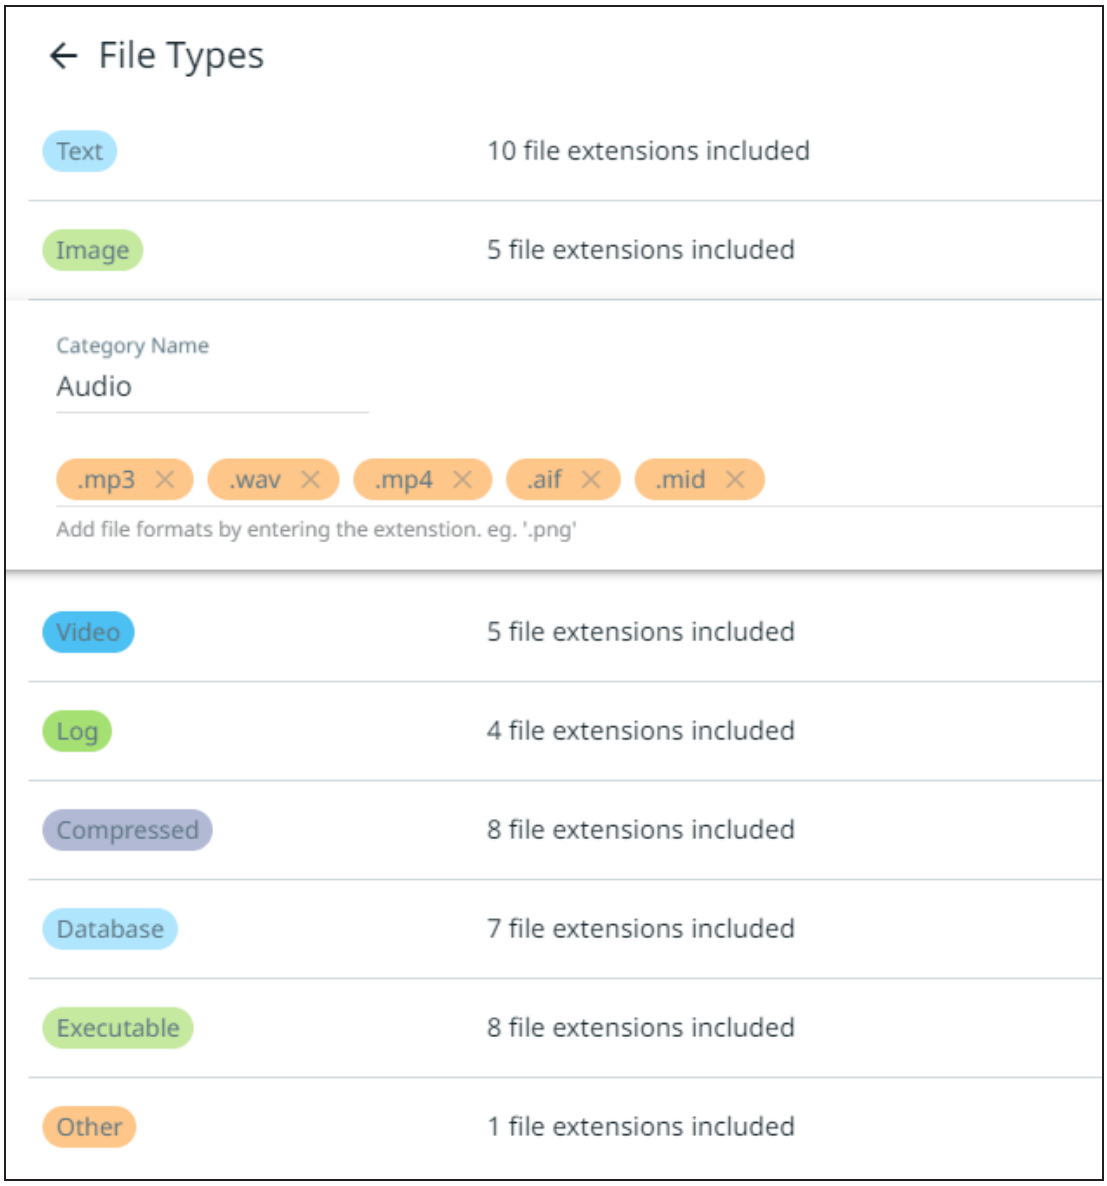

Delete Analysis

To delete the data analysis done on the source:

1. On the **NAS Tiering** page, go to the **Tiering by Source** pane.
2. Click the Actions menu (⋮) corresponding to the source you analyzed:
3. Click **Delete Analysis**.

   The **Delete Analysis** dialog is displayed.

4. Click **Delete** to confirm your action:

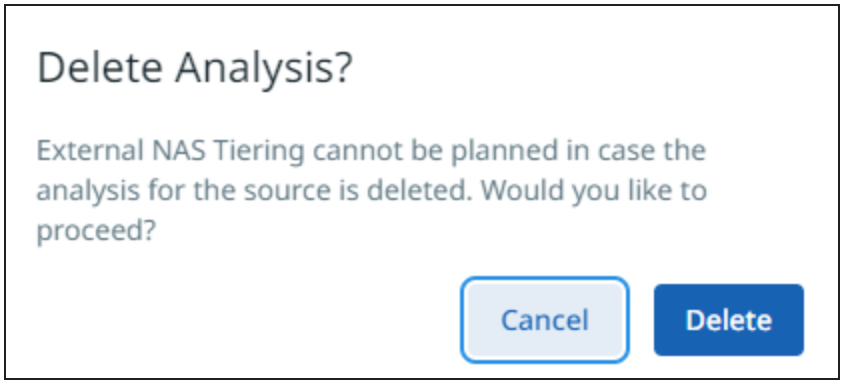

Edit Analysis Schedule

You can edit the schedule on an analysis run.

To edit the analysis schedule:

1. On the **NAS Tiering** page, go to the **Tiering by Source** pane.
2. Click the Actions menu (⋮) corresponding to the source you analyzed:
3. Click **Edit Analysis Schedule**.
4. The **Edit Analysis Schedule** dialog is displayed:

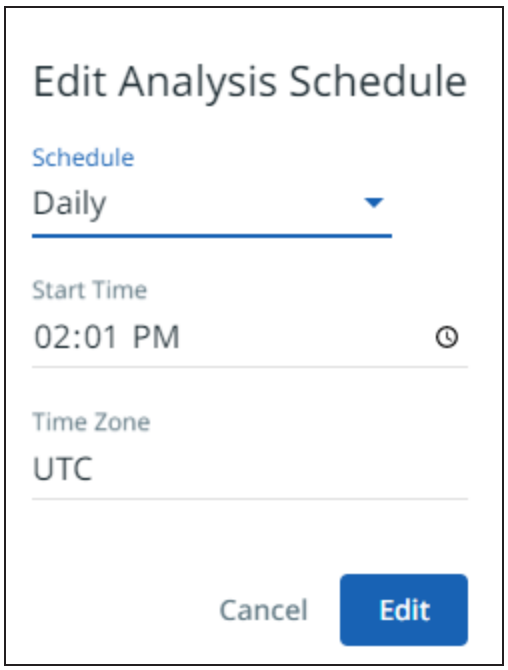

5. Modify the schedule as necessary and click **Edit** to save the changes.

Plan Tiering

Once the analysis is complete, based on the data analysis done, you can plan the tiering for your data on the NAS source. To plan data tiering to the Cohesity cluster, click the **Plan**

Tiering option that appears once you complete the analysis:

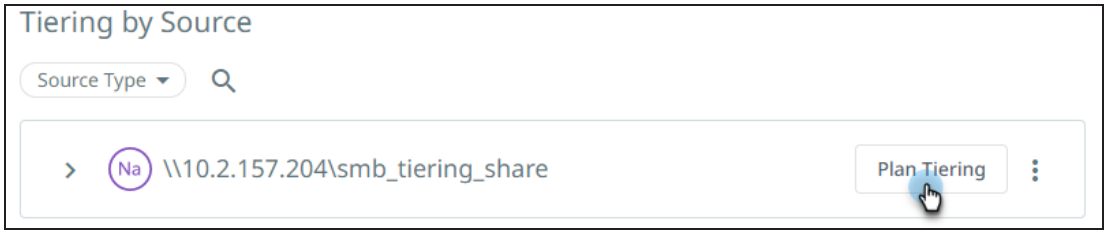

Plan tiering involves the following steps:

- Select Data To Tier
- Configure a Down-Tiering Task

Select Data To Tier

After you click **Plan Tiering**, the **Select Data To Tier** page is displayed. This page displays the name of the registered source and includes the following sections:

- Select Data To Tier
- Selection Summary
- Data Distribution Chart
- Data per Share

Select Data To Tier

In the **Select Data To Tier** section, you can select the files to tier. By default, the filters are set as follows:

- Scanned Objects—**All**
- File Accessed—**Warm** and **Cold**
- File Modified—**All**
- File Size—**All**
- File Type—**All**

You can modify the filters and select the data you want to tier. For example, if you select **File Accessed** as greater than six months, **File Size** as greater than 1 GiB, and **File Type** as Video, video files greater than 1 GiB that were accessed more than six months ago will be down-tiered:

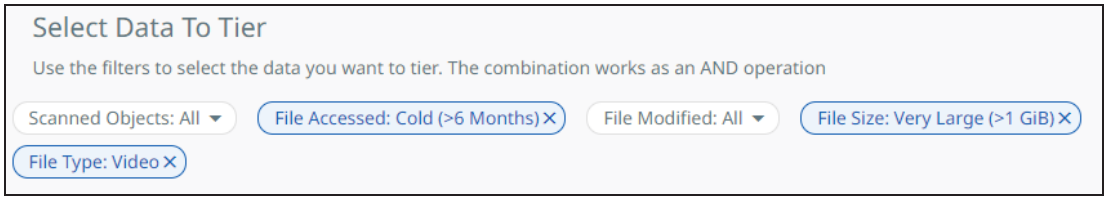

> **Important:** The combination works as an AND operation.

Selection Summary

The **Selection Summary** section displays the summary:

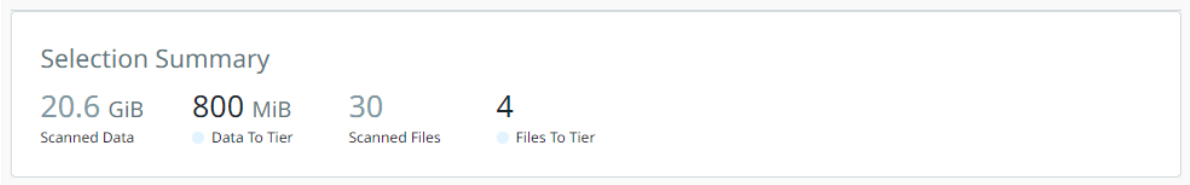

- Total size of scanned data
- Size of data that will be tiered
- Number of scanned files
- Number of files that will be tiered

Data Distribution Chart

The **Data Distribution** chart provides a graphical representation of the data (histogram) available on the source by distribution and group.

The following is an illustration of the data distribution by file accessed and grouped by file size. Hover over the chart to view granular details:

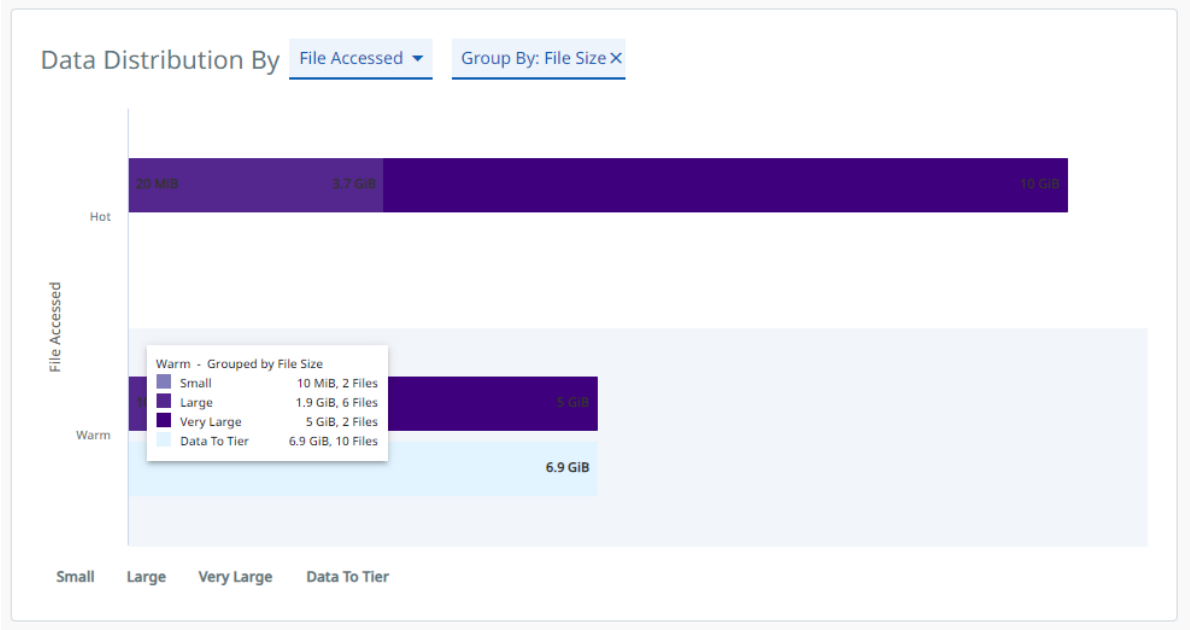

Data per Share

The **Data per Share** section displays the following data per share:

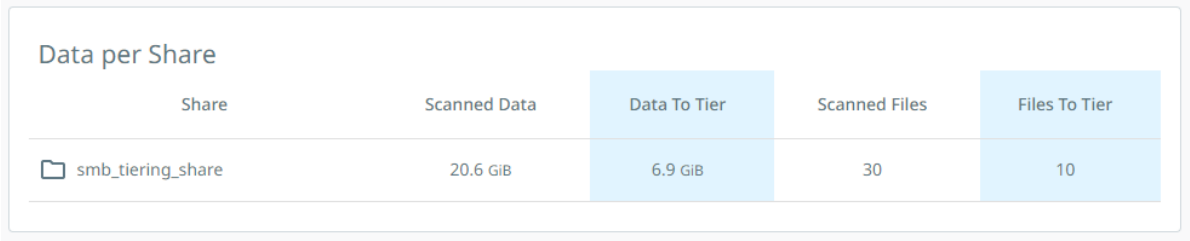

- Name of the share
- Scanned data
- Data to tier
- Scanned files
- Number of files to tier

> **Note:** Cohesity currently supports only one share per source. If you want to include another share, you must register the share as a new source.

Click **Continue** to navigate to the **Configure Tiering Task** page.

Configure a Down-Tiering Task

To configure a downtiering task, complete the steps described in the Select Data To Tier section and proceed:

1. In the **Task Name** section:
   - Enter a name for the tiering task
   - Optionally, enter a description.

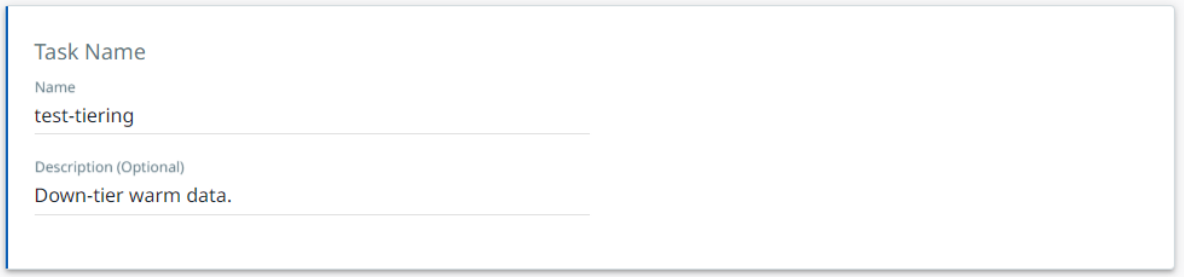

2. View the selected data to tier. The **Selected Data To Tier** section is read-only and displays the data that you selected in the **Select Data to Tier** section:

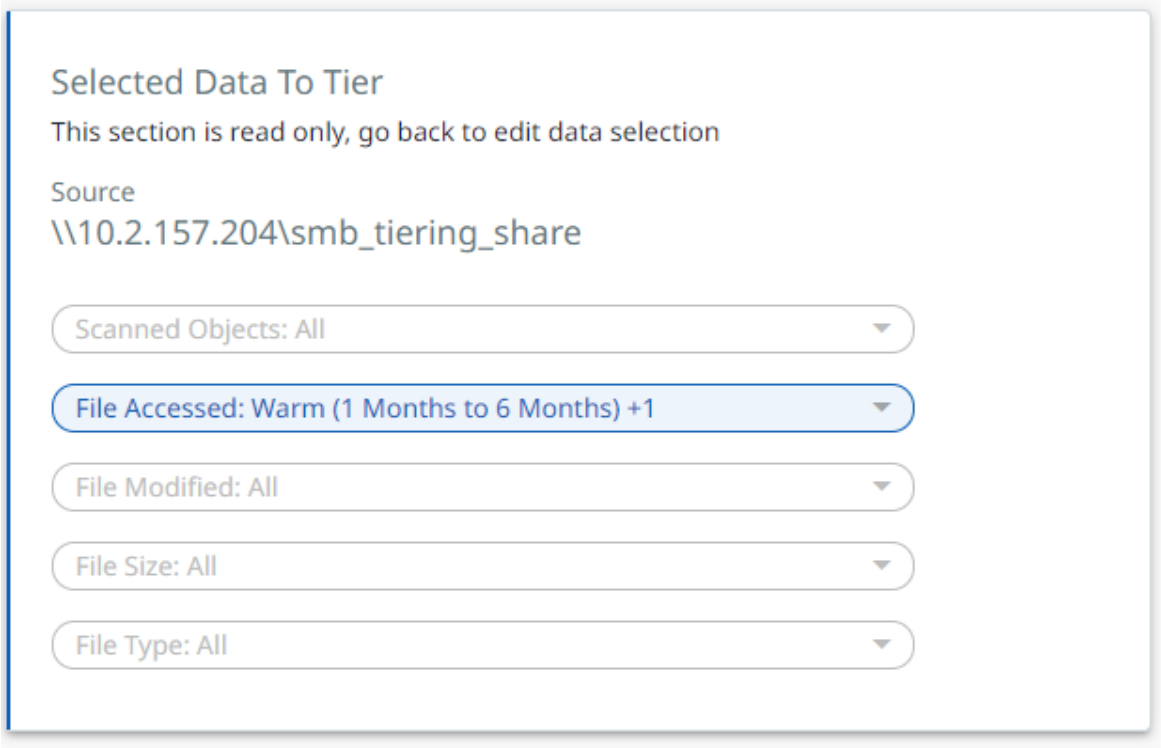

3. Select a storage domain where you would want the target View to be created.
4. Enter a name for the target View that will be created.
5. **Task Schedule**—Schedule the tiering task, by default it is set to **Daily**.
6. **Retain Snapshots**—Specify the number of days the snapshots are retained on the Cohesity cluster. At the end of each task run, a snapshot of the View is captured. This snapshot is retained for the number of days specified.
7. Select one of the following **Symbolic Links** options:
   - **Tier with symbolic links**—Select this option if you want to create symbolic links when data is tiered. When data is tiered from the primary storage, symbolic links are substituted in place of the original file. The name of this symbolic link is the same as that of the original file, ensuring that users and applications continue to access the data as before. Optionally, you can also delete data on the Cohesity target when the symbolic link is deleted.

     > **Note:** If you use the **Tier with symbolic links** option, you cannot use Mac clients to access data tiered with symbolic links.

   - **Tier without symbolic links**—Select this option if you do not want to create symbolic links when data is tiered. You can select this option if no user or

application will access the data on the source and you are permanently moving the unused data to the Cohesity cluster.

8. Click **More Options** to view additional settings.
9. **Audit Logging**—Enable this option if you want to log the information of the data down-tiered. You can also check the logs related to the symlinks created by the downtiering task and the files that are renamed or deleted. The audit log file is saved in the **tiering_audit** directory of the **COHESITY_AUDIT_VIEW** View. You can access this file by mounting the View as NFS share: *<cluster_ip>*:/COHESITY_AUDIT_VIEW.
10. **Tiering Goal**—Turn on the **Set Share Capacity Limit** toggle to set the amount of data you want to retain at the source. Depending on the value you set, Cohesity provides a recommendation for tiering data based on the last accessed date:

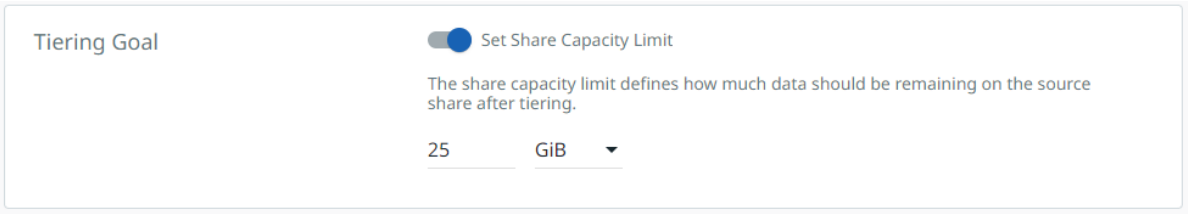

11. **QoS Policy**—Select an appropriate Quality of Service (QoS) Policy:
    - **Backup HDD**—The Cohesity cluster writes the data directly to an HDD drive for this tiering task.
    - **Backup SSD**—The Cohesity cluster writes the data directly to an SSD drive for this tiering task. Specify this policy only if you need fast ingest speed for a small number of tiering tasks.

    > **Note:** Cohesity recommends specifying a **Backup HDD**, which is the default.

12. **Indexing**—Enable this option to index the data of the tiering task. Indexing improves files and folder search. By default, a rule to include the root directory (/) is defined, and therefore by default, the entire file system is indexed. Indexing is recursive, so if a directory is indexed, all of the child subdirectories are also indexed.
13. **Alerts**—Select one or more of the following settings if you want alerts to be created for the following triggers: Success, Failure, or SLA Violation.

    Email Recipients—Click **Add** and enter an email address to the tiering task to notify the email recipients when alerts are triggered for the task.

    > **Note:** Emails are not sent when informational alerts are created.

14. Click **Create**.

    The summary dialog displays the details of the tiering task and the estimated capacity savings on source:

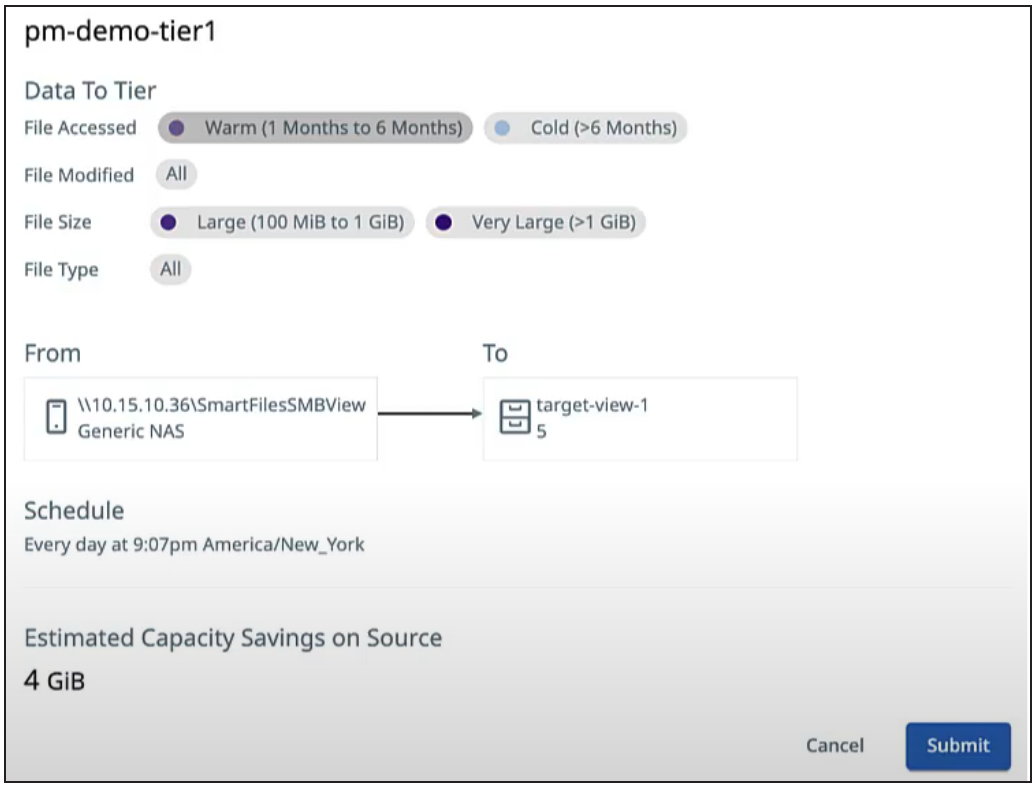

15. Click **Submit**.

    The downtiering task is created and will run at the specified scheduled time.

Tiering Audit

The **Tiering Audit** option at the cluster level and the **Audit Logging** option in the downtiering task record details about the downtiered data. This includes logs for created symlinks and files that have been renamed or deleted during the downtiering process. The audit log file is stored in the `tiering_audit` directory within the `COHESITY_AUDIT_VIEW` View. Mount the View as an NFS share at `<*cluster_ip*>:/COHESITY_AUDIT_VIEW` to access this file.

The following is an example audit log of a file that was down-tiered:

```
{"Timestamp" : "2024-07-23T06:10:04.314Z", "TieringAction" : "DownTier", "RemoteShare" :
"10.14.59.185:/source_view", "ViewName" : "target_view", "EntityPath" : "/err.txt", "EntityModifyTimestamp" :
"2024-07-23T06:09:50.491Z", "EntityAccessTimestamp" : "2024-07-23T06:09:50.535Z", "EntityMode" : 33152,
"EntityUserID" : 0, "EntityGroupID" : 0, "EntitySize" : 3, "EntityInode" : 154514, "JobName" : "tier"}
```

Enable Tiering Audit at the Cluster Level

To enable tiering audit logs at the cluster level:

1. Navigate to **System > Audit Logs > Log Settings**.
2. Turn on the **Tiering Audit** toggle:

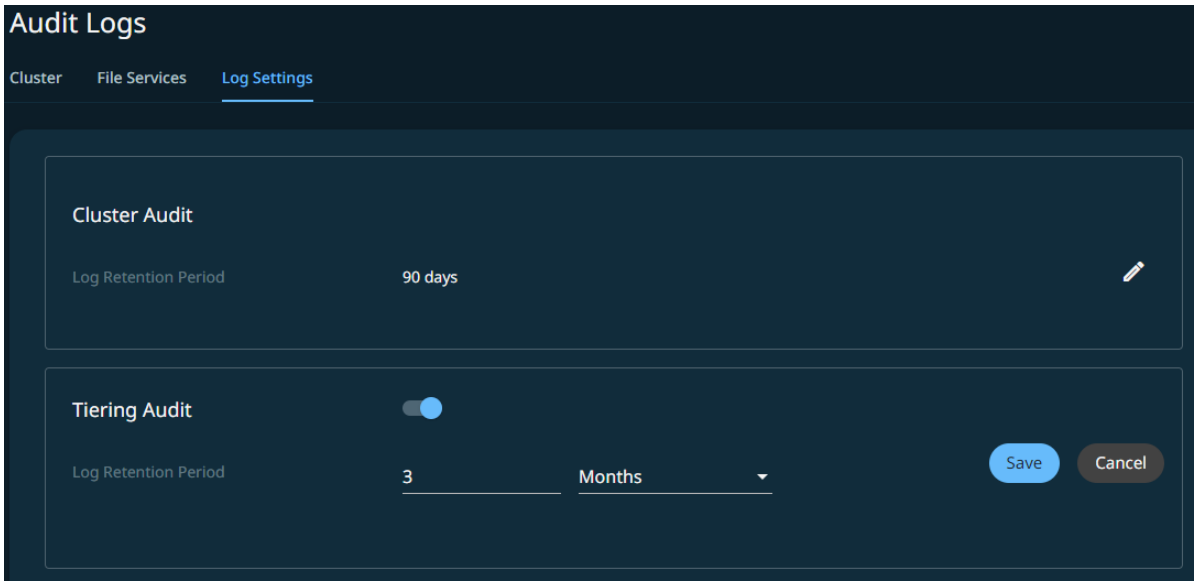

3. Update the log retention period as required.
4. Click **Save**.

Enable Audit Logs in the Downtiering Task

You can enable tiering audit logs for each downtiering task when creating or modifying a tiering job. Complete the following steps:

1. Navigate to **SmartFiles > NAS Tiering**.
2. In the **Tiering by Source** section, locate the registered source for which you want to enable audit logging:
   - To enable audit logging while creating a new tiering job, click **Plan Tiering** and **Select Data To Tier**.
   - To enable audit logging for an existing tiering job, hover over the job name and select **Edit**. On the **Select Data To Tier** page, click **Continue**.
3. Click **More Options** and turn on the **Audit Logging** option:

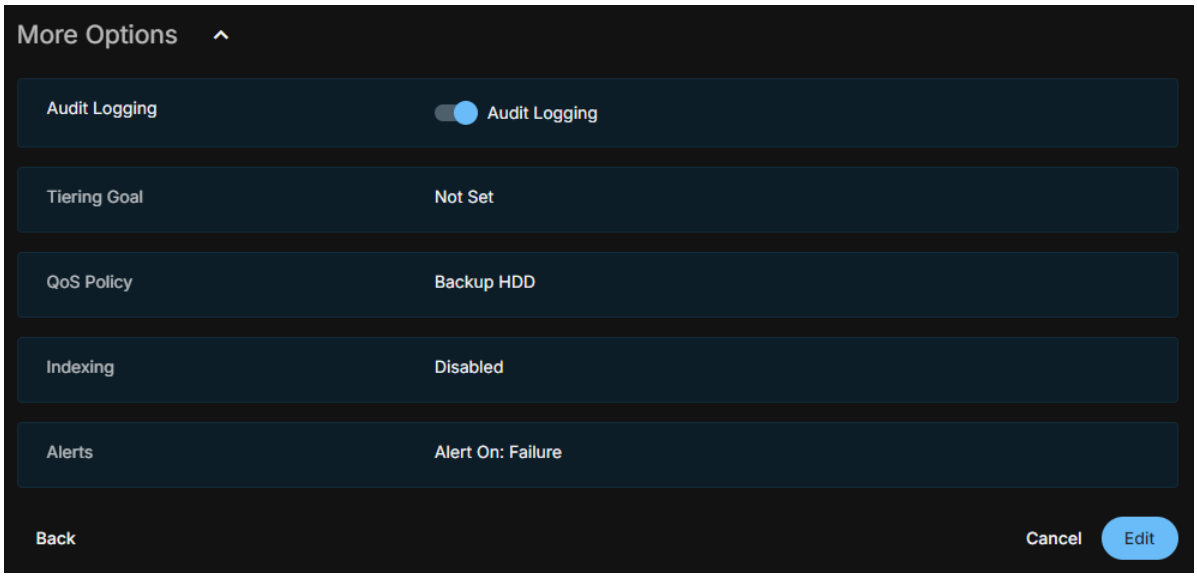

4. Click **Create** if you are creating a new job or **Edit** if you are modifying an existing job.

File Services Audit Logs

The File Services Audit Logs feature enables you to record and monitor events occurring on your SMB and NFS views.

Unlike the cluster-level audit logs, the logs for views are not displayed in the Cohesity GUI. The file services audit logs are saved in an internal view called ***filesystem_audit***. You can access these views by mounting the path ***COHESITY_AUDIT_VIEW\filesystem_audit***.

The **File Services** tab in the **Audit Log** page provides two mount paths for NFS and SMB. You can mount these paths to access and view the audit logs.

You can set permissions for the view and restrict access only to required subnets or users. Only the subnet's IP addresses in the allowlist can access and view the audit logs.

> **Note:**
>
> By default, when you create a view, the **Audit Log** option is disabled. You must manually enable it when creating the view or after creation.

Events Logged for NFS Views

Event	Description
Mount	To mount a view using NFS.
File Create	To create a file.
Link	To create a hard link to a file.
Folder Create	To create a folder.
Mknod	To create a special file used for resource like a device.
File Remove	To delete a file.
Folder Remove	To delete a folder.
Rename	To rename a file or folder.
Change Mode	To change the Unix mode bits of a file or folder for permission.
Change Ownership	To change the Unix UID or GID of a file or folder.
Modify Content	To modify the file's content.
Symlink	To create a soft link to a file or folder.

The following information is logged for NFS views. The information logged depends on the event type.

- Timestamp
- RecordID
- Protocol
- Client IP
- UserID
- ViewName
- Fsid
- RequestType
- Result
- EntityType

- EntityPath
- EntityInode
- EntityMode
- PrevEntityMode
- EntityUserID
- PrevEntityUserID
- EntityGroupID
- PrevEntityGroupID
- EntityModifyTimestamp
- DestinationEntityPath
- DestinationEntityInode
- InodeSizeBytes
- SourceParentEntityInode

Events Logged for SMB Views

Event	Description
Mount	To mount a view using SMB.
Logon	To establish an SMB session.
Logoff	To disconnect the current SMB session.
Close	To close a file or folder.
Create	To create a file or folder.
File Open	To open a file.
Delete	To delete a file or folder.
Rename	To rename a file or folder.
Modify Content	To modify the content of a file.
Change Ownership	To change the owner of a file or folder.

Event	Description
Change Permission	To add or remove the ACL permission of a file or folder.
File Change Protection	To enable or disable the Windows file protection of a file or folder.

The following information is logged for SMB views. The information logged depends on the event type.

- Timestamp
- RecordID
- Protocol
- Client IP
- Username
- UserID
- DomainName
- UserSID
- ViewName
- ViewAliasName
- Fsid
- RequestType
- Result
- EntityType
- EntityPath
- EntityInode
- EntityMode
- EntityUserID
- EntityGroupID
- SecurityDescriptor
- PrevSecurityDescriptor
- EntityModifyTimestamp
- EntityAccessMask
- DestinationEntityPath
- DestinationEntityInode
- DeleteOnClose

- InodeSizeBytes
- SourceParentEntityInode

Settings for File Services Audit Logs

Enable or Disable File Services Audit Logs

By default, the File Services Audit setting is disabled. You should enable the File Services Audit option for Cohesity to start logging operations on the cluster's views for auditing purposes.

You can disable or enable it back from the **Settings** tab of the **Audit Logs** page. When you disable the option, the old logs are retained until the retention period ends.

To enable or disable File Services Audit, complete the following:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **Log Settings** tab, turn on/off the **File Services Audit** toggle to enable or disable.

After enabling the **File Services Audit**, the **Log Retention Period** audit information option appears. Use this option to set the audit logs retention period.

Set Log Retention Period for File Services Audit Logs

You can set the retention period for file services audit logs on the Cohesity cluster.

> **Note:** The default retention period is 90 days. You can increase or decrease the renewal period. Once the retention period expires, the audit logs are deleted automatically.

Prerequisite:

- You must enable the File Services Audit option. For more information, see Settings for File Services Audit Logs.

To set a retention period for File Services Audit Logs, complete the following:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **File Services Audit** section of the **Log Settings** tab, click the icon for the **Log Retention Period** option, type the desired number, and select the retention period (Days, Weeks, Months, or Years).

   The Log Retention Period is displayed in days as Cohesity converts weeks, months, or years into days.
3. Click **Save**

Add IP Allowlist for File Services Audit Logs

By default, every client that is part of the Global Allowlist set at the cluster level can access the File Services audit logs. You can change access permission by the **Override Allowlist** option to allow or deny access to a subnet. Additionally, access can also be restricted using SMB Permissions on the view for AD or LOCAL principals.

As the file services audit logs are saved in an internal view, you can provide access only to required networks by adding them to the allowlist along with Allow Deny setting for NFS/SMB Clients in the allowlist.

To override the global IP allowlist set at the cluster level and add a subnet to the allowlist for file service audit:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **IP Allowlist** section of the **Log Settings** page, turn on the **Override Global IP Allowlist** toggle.
3. Click **Add**.
4. In the **Add Allowlist** dialog, complete the following:
    1. Type the Subnet IP address in CIDR format.

       Example: Subnet in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).
    2. Select **NFS Permissions**, **SMB Permissions**, or **S3 Permissions**. (**Read/Write**, **Read Only**, or **Disabled**).
    3. Optionally, type a description.
    4. Click **Add**.

    The added allowlist IP subnet details appear in the File Services Audit section of the page.

Edit IP Allowlist for File Services Audit Logs

You can edit the IP allowlist for File Services audit logs if you need to change the subnet or change the permission for access.

To provide access to a network by adding to the IP allowlist for file service audit:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **IP Allowlist** section of the **Log Settings** page, click the ✎ for the required subnet.
3. In the **Edit Allowlist** dialog, complete the following:
    1. Modify the Subnet IP address.

       Example: Subnet in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).

2. Modify **NFS Permissions**, **SMB Permissions**, or **S3 Permissions**. (**Read/Write**, **Read Only**, or **Disabled**).
3. Optionally, type a description.
4. Click **Add**.

The added allowlist IP subnet details appear in the File Services Audit section of the page.

Delete Allowlist for File Services Audit Logs

You can delete the allowlist IP subnet to remove access to a network. After deleting, clients in the deleted subnet cannot access the audit log files.

To delete an allowlist for File Services Audit logs:

1. Navigate to **System> Audit Logs>Log Settings**.
2. In the **IP Allowlist** section of the **Log Settings** page, click the delete icon for the required subnet.
3. In the **Delete Allowlist** dialog, enter **YES** to confirm, and click **Delete**.

The subnet disappears from the list and the audit logs view is restricted for the deleted subnet.

Add SMB Permissions for File Services Audit Logs

You can provide permission for SMB domain users and groups. The access permission is Read-Only and access type can be Allow or Deny.

To add SMB permissions for the File Services audit log:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **SMB Permissions** section of the **Log Settings** page, click the **Add**.
    1. In the **Add SMB Permission** dialog, complete the following:
    2. Select an **AD Domain** from the list.
    3. Select a **Principal** from the list.
    4. Select Type: **Allow** or **Deny**.
    5. Click **Add**.

    The new permission appears under the SMB Permissions section.

Delete SMB Permissions for File Services Audit Logs

You can delete an SMB permission to remove the set permission (Allow or Deny) for a principal (user or group).

To delete SMB permissions for file services audit log, complete the following:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **SMB Permissions** section of the **Log Settings** page, click the delete icon for the required SMB permission.
3. In the **Delete SMB Permission** screen, type **Yes** for confirmation.
4. Click **Delete**.

Access NFS Audit Logs using a Mount Path

Prerequisite:

- The subnet you are using should have read-only permission to the File services audit log view.

To access NFS audit logs, complete the following steps:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Log Location** section of the **File Services** tab, click the copy icon for NFS and mount using the following steps:
    1. Mount the path using the syntax given below:

        ```
        mount cluster_name:/COHESITY_AUDIT_VIEW/filesystem_audit /mnt/path
        ```

    2. Connect to a hidden internal view 'COHESITY_AUDIT_VIEW'.
    3. Connect to `mount cluster_name:/COHESITY_AUDIT_VIEW/filesystem_audit /mnt/path`.

Sample NFS Audit Log

```
{"Timestamp" : "2022-03-10T22:12:17.838Z", "RecordID" : "bd65:2afc4", "Pro-
tocol" : "NFS3", "ClientIP" : "10.2.223.127", "UserID" : 1000, "ViewName" :
"test", "Fsid" : 4377, "RequestType" : "Delete", "Result" : "NFS3_OK",
"EntityType" : "File", "EntityPath" : "bbb", "EntityInode" : 9422800,
"SourceParentEntityInode" : 9422329}
```

Access SMB Audit Logs using a Mount Path

Prerequisites:

- The subnet you are using should have read-only permission to the File services audit log View.
- As an SMB user, you must have permission to access the File services audit log View.

To access SMB audit logs, complete the following steps:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Log Location** section of the **File Services** tab, click the copy icon for SMB and mount using the following steps:
    1. Connect to a hidden internal view 'COHESITY_AUDIT_VIEW'.
    2. Access `\\cluster_name\COHESITY_AUDIT_VIEW\filesystem_audit`.

Sample SMB Audit Log


```
{"Timestamp" : "2022-03-04T00:01:40.503Z", "RecordID" : "30fe:216", "Pro-
tocol" : "SMB", "ClientIP" : "10.14.0.41", "UserName" : "Administrator",
"UserID" : 500, "DomainName" : "w2k16-0.qa.local", "UserSID" : "S-1-5-21-
2148495376-676708989-2532157484-500", "ViewName" : "audit_filer_actions",
"Fsid" : 6642, "RequestType" : "SetAttributes", "Result" : "Ok",
"EntityType" : "File", "EntityPath" : "audit-actions/d1/e1/f1.txt",
"EntityInode" : 6214}
```


Check Audit Logs Status for Views

You can enable or disable the status of the Audit Logs option for views.

You can check the audit status for views to know whether audit logs are being logged for required views. The GUI displays all views with the status and it allows you to filter the views to see the audit log enabled or disabled list of views.

To check the audit status for views, complete the following:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Audit Status of Views** section, click the **Audit Logs** drop-down menu and select **Enabled** or **Disabled**, and click **Apply**

Based on your selection, the list of views with **Audit Log enabled** or **disabled** status appears.

Enable or Disable Audit Logs for Views

You can check the audit status of a view to know whether audit logs are being logged for required the view and enable or disable the **Audit Log** option.

You can also enable or disable the Audit Logs option for a view from: **SmartFiles > Views**, click the actions menu ( ) and select **Edit** > **Audit Logs.**

To enable or disable audit logs status for a view, complete the following:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Audit Status of Views** section, click the **Audit Logs** drop-down menu and select **Enabled** or **Disabled**, and click **Apply**.

3. Based on your selection, the list of views with **Audit Log** enabled or disabled status appears.
4. To enable or disable audit logs, click the toggle corresponding to the view under the **Audit Logs** column.

Antivirus Software to Protect Cluster Data

You can enable antivirus software to protect your data on the Cohesity cluster from viruses. Once an antivirus software scans a file and detects a threat, the Cohesity cluster reports the threat on the Cohesity Dashboard. The administrators can either quarantine, unquarantine, reset, or delete the infected file.

Perform the following steps to enable antivirus scanning and manage threats:

1. Add an Antivirus Provider
2. Enable Antivirus Scan on a View
3. View and Manage Threats

Add an Antivirus Provider

To perform an antivirus scanning on your data in the Cohesity cluster, you must first add the antivirus software to the Cohesity cluster. You can add the software in one of the following ways:

- Register External Antivirus Software
- Install the Antivirus App from the Cohesity Marketplace

Register External Antivirus Software

1. Navigate to **SmartFiles > Antivirus** and click the **Providers** tab.
2. Click **Add External Provider** in the upper-right corner of the page.
3. In the **New External Provider** field, enter the name of the antivirus software provider of your choice.
4. (Optional) Describe the provider.
5. Enter the ICAP Server URL in the following format: `icap://<server>:<port>/<uri>`

Option	Description
icap	Indicates the required protocol identifier that enables communication with external servers hosting third-party antivirus software.

Option	Description
server	Specifies the domain name or IP address of the external server (ICAP server) hosting the third-party antivirus software.
port	Specifies the port that the external server monitors for incoming requests. If not specified, the default icap port, 1344 will be used.
uri	Specifies the path of the ICAP service.

6. (*Optional*) Click **Add** to register more than one external ICAP server.
7. Click **Check ICAP Server Connection Status** to check the status of the external antivirus software.

   If the connection is not successful, ensure that you have provided the correct ICAP URL and the external server is up and running.
8. Click **Add**.

   The antivirus software is registered on the Cohesity cluster.

Install the Antivirus App from the Cohesity Marketplace

1. Navigate to **SmartFiles > Antivirus** and click the **Providers** tab.
2. Click **Install Antivirus App** at the top right of the page.

   You will be directed to the Cohesity App Marketplace.
3. Click **GET APP** for the Antivirus App you want to install on the Cohesity cluster.

   You will be directed to the Helios login page.
4. Log in to Cohesity Helios, select the Cohesity cluster on which you want to install the Antivirus App, and click **Install**. You also have the option to install the App later.
5. Go back to the Cohesity cluster and go to **Marketplace > My Apps**.

   The page displays the list of Apps that are available on the Cohesity cluster.
6. Click **Run App** to create an instance of the Antivirus App.

   The instance created will be displayed as an Antivirus provider on the **Providers** tab.

Enable Antivirus Provider

1. Navigate to **SmartFiles > Antivirus** and click the **Providers** tab.

   The page displays the list of antivirus providers that can be enabled for antivirus scanning.
2. From the Action menu, toggle on the **Enable** option.

   You can enable only one antivirus provider group at a time.

Enable Antivirus Scan on a View

Once you install the antivirus software on the Cohesity cluster, you can enable the antivirus scan on a View.

To enable antivirus scan on a view:

1. Navigate to **SmartFiles > Antivirus** and click the **Secured Views** tab.
2. Click **Add a View** at the top right of the page.
3. From the View drop-down list, select a view for which you want to perform the scanning of files.

   Antivirus scanning is supported only on Views with SMB protocol. Therefore, the drop-down list displays the Views which support SMB protocol.
4. Provide the following details either by clicking **Edit All** or by clicking the **Edit** option for each scan settings option:

Option	Description
Scan on	Antivirus scans can be configured to send the files for scanning when they are closed, opened, or both. In each case, all the files to be scanned are added in a queue and sent to the antivirus provider(s) for scanning, as per the queue order. • **Close**: Files are scanned after they are closed (Default) With this configuration, you can access the files without any delay as file scanning is done after the user's file operation. This type of scanning is faster but less secure. • **Open**: Files are scanned before they are opened. With this configuration, you can access the files once they are scanned. You must wait for the file scan to complete. This type of scanning is slower but more secure. • **Open and Close**: Scans the files before they are opened and after they are closed. This type of scanning is the slowest and most secure, as the scanning queue includes all the live files (opened and closed). Cohesity recommends scanning the file after they are closed as this configuration ensures that the files are accessible even if the antivirus provider(s) or scanning network is down, unavailable or unreachable.

Option	Description
Files to Scan	• **All Files**: Select this option to scan all the files in the View. • **All Files Except**: Select this option to exclude files in the view from scanning. You can specify the file extensions (comma-separated) to be excluded in the **File extensions to skip** field. • **Only These Files**: Select this option to scan only the selected files. You can specify the file extensions (Comma-separated) to be included in the **File extensions to scan** field.
	Exclude files based on file size: Enable this option to exclude files based on the file size. In the **Max File Size** field, you can specify the file size limit for scanning. The default file size limit is 50 MB. The maximum file size limit is 125 MB.
Scan Timeout	Set the maximum amount of time (in seconds) that a scan can take before timing out.
If Scan Cannot Be Completed	Blocks file access if the antivirus servers are unavailable or are unable to scan the file due to any issue. This option is disabled by default to ensure that the files stay accessible even if the antivirus provider(s) or scanning network is down, unavailable, or unreachable.

5. Click **Add**.

You can also enable antivirus scan while creating a View or by editing an existing View.

View and Manage Threats

After you enable the antivirus scan on the View, you can view the scanned report from the **Threats** page and take action on the infected files.

To view the scanned report:

Navigate to **SmartFiles > Antivirus** to view the **Threats** page.

Alternatively, on the **Secured Views** page, click the view for which you want to view the reports.

The Threats page displays the following details:

- The total number of quarantined files
- The total number of unquarantined files
- The total number of views that are affected by the virus
- Option to filter the report based on the status of the file and View
- Information on the files that are affected:

Field	Description
Filename	The files that are infected, along with the file path. A lock icon will be displayed along with the file name if you have enabled the **File DataLock** option while creating the View.
View	The View name of the files. A lock icon will be displayed along with the View name if you have cloned the View as DataLock.
Threat	The threat type or virus signature along with the ICAP server URL that detected the threat.
Provider Name	The antivirus software provider that detected the threat.
Scan Time	The time at which the file is scanned for viruses.
Status	The status of the threat: • **Quarantine**: The file is determined to be unsafe. All access to the file is blocked until the administrator manually changes the file state. • **Unquarantine**: Allow file access. The file is not scanned again until the file state is reset.

Manage Threats

You can manage threats by performing the following actions from the **Threats** page:

Action	Description	How to
Unquarantine Files	Move the files from the quarantined state to allow file access. Unquarantined files are not scanned again.	Select the checkbox next to file names and then click **Unquarantine Files**.
Delete Files	Deletes the quarantined files. You cannot delete unquarantine files.	Select the checkbox next to file names and then click **Delete Files**.
Reset Files	Reset the unquarantine files. You cannot reset quarantine files	Select the checkbox next to file names and then click **Reset Files**.

Access Management

The following topics provide information about user and group access management:

Active Directory (AD)

Active Directory (AD) in Cohesity Access Management is a directory service implemented by Microsoft that provides several services such as LDAP (Lightweight Directory Access Protocol), Kerberos, and DNS. It is used for user and group identity management and authentication within a Cohesity cluster.

You can join the Cohesity cluster to one or more Active Directory (AD) domains. Joining to AD allows the following:

- Existing AD users and groups can be assigned Cohesity roles
- SMB shares created by the Cohesity cluster can be authenticated based on AD credentials

Cohesity Roles

Once the Cohesity cluster is joined to an AD domain, existing AD principals can be added assigned different roles so they can access the Cohesity Dashboard. For example, you can assign the View role to the AD group "ITGroup." This gives all group members read-only access to the Cohesity Dashboard. You can also assign the Admin role to the AD user "ITManager."

Cohesity SMB Share Access

The action steps in the following procedure are represented in the figure below with blue numbered time squares: T1 . The procedure assumes the Cohesity cluster is already joined to an Active Directory domain.

Client access to an Active Directory-joined Cohesity cluster SMB share consists of the following actions:

T1) A client requests access to a Cohesity cluster SMB share.

T2) The client attempts to perform an operation on the Cohesity cluster SMB share.

T3) Active Directory authenticates the client's membership within a group that can perform the requested operations on the Cohesity cluster SMB share.

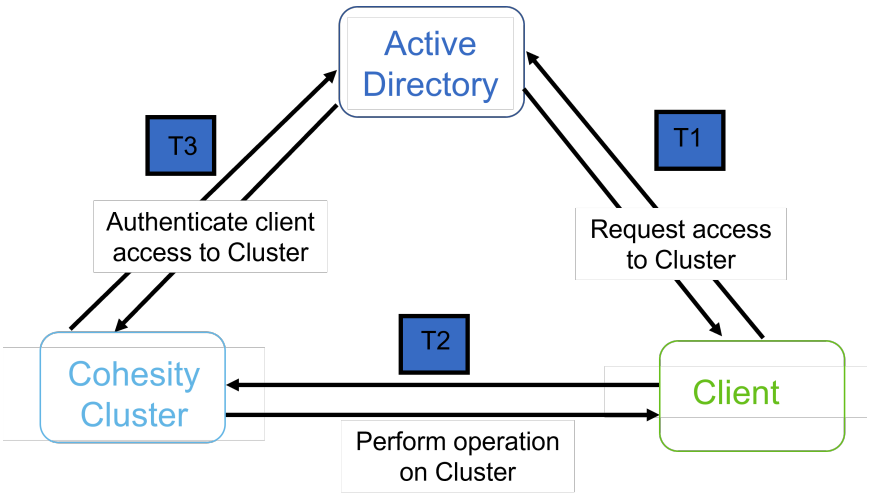

Join to Active Directory Domain

For detailed instructions, see Join to Active Directory.

Join Active Directory

This topic explains how to configure a Cohesity cluster to use Active Directory (AD) credentials for SMB access to views and to access the Cohesity Dashboard. (For information about protecting Active Directory domain controllers, see About Active Directory Data Protection.)

Joining the Cohesity cluster to an AD domain is optional. You can join the Cohesity cluster to one or more AD domains to enable:

- existing AD users can access Cohesity cluster SMB shares based on their AD credentials
- existing AD users and groups can be assigned Cohesity roles and access the Cohesity Dashboard

On Cohesity clusters that are **not** joined to AD, all views are accessible as SMB shares and there is no authentication. If the AD administrator blocks NTLM, the shares on the Cohesity cluster are not accessible.

All accounts and Cohesity cluster SMB share privileges for users in AD are managed through Active Directory or where the SMB shares are mounted. Non-admin AD users must be granted backup and restore privileges on Windows clients explicitly. This is in addition to backup and restore privileges for the Cohesity cluster.

Migrated AD Users and Groups with Historical SIDs

Cohesity supports migrated AD users and groups with historical SIDs and grants access accordingly in Active Directory domains. The SID history retains user and group membership and access rights during Active Directory domain migration.

Considerations

Review the following considerations of migrated AD users and groups with historical SIDs:

- Historical SIDs with multiprotocol access is not supported.
- Historical SIDs user quota accounting is not supported.
- Privilege assignment for Historical SIDs is not supported.

Prerequisites and Considerations

For information about the required firewall ports for AD, see Manage Firewall Ports.

Join the Cohesity cluster to an AD Domain

1. Go to **Settings > Access Management** and click the **Active Directory** tab.
2. Click **Add Active Directory**.

3. **Domain Name:** Provide the fully qualified domain name (FQDN) of the Active Directory domain to join.
4. **Username**: Provide the samAccountName.
5. **Password**: Provide the password for the username.
6. **Domain Controllers:** Add preferred domain controllers to the AD domain to ensure that the Cohesity cluster uses a local domain controller (when possible) to speed up AD queries. You can set preferred domain controllers for the following:
   - Domain controllers
   - Trusted domains of the domain controller

   > **Note:**
   > - The SRV records must exist for the domain (the list of domain controllers displayed in the Cohesity Dashboard is generated by querying the SRV records).
   > - You can add any number of preferred domain controllers.

   1. Click **Manage** to allow or deny domain controllers and change the order of allowed controllers:

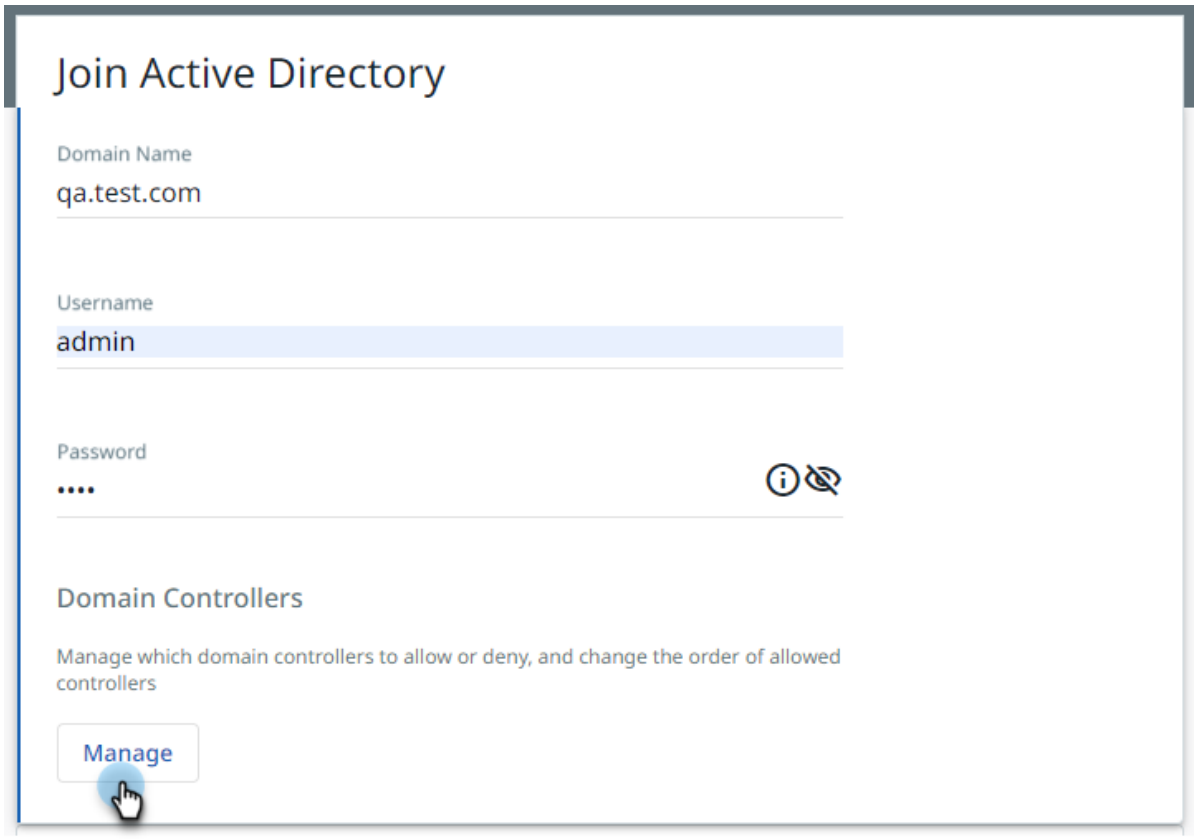

2. In the **Manage Domain Controllers for AD Domain** dialog, the domain controllers are listed in order of preference. You can move the allowed domain controller(s) up and down to put them in order of preference. To block a particular domain controller, hover on it and click **Deny:**

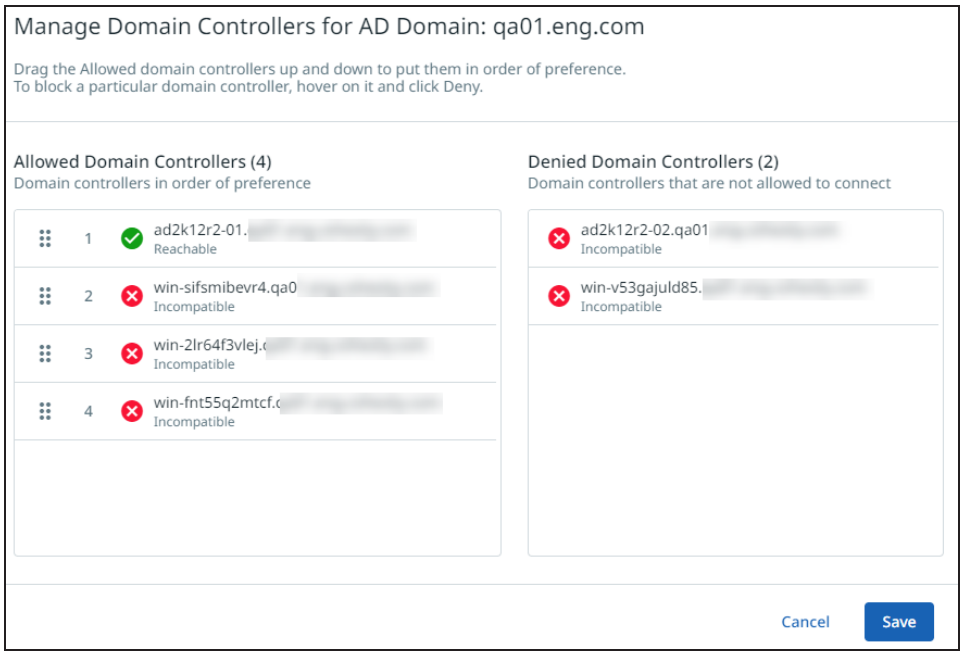

3. Click **Save**.

   The allowed domain controllers are listed in the **Domain Controllers** section:

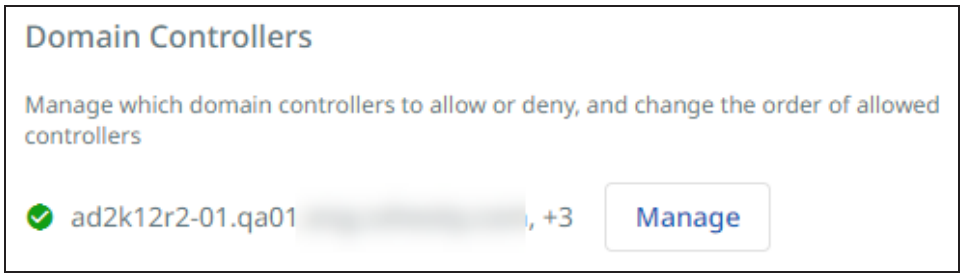

   The Cohesity cluster uses these domain controllers to join the AD domain and for all future communication with the AD domain.

7. **Machine Accounts:** Click **Add** and specify the following details:

   - **Machine Account**—Provide at least one machine account to identify the Cohesity cluster on the domain. For example, you can use the Cohesity cluster name.
   - **DNS Hostname**—Optionally, specify the DNS hostname of the machine account. For AD to work with the internal DNS server, apart from machine account names, you should add DNS hostnames that should be configured for each machine account. The DNS hostnames should have one of the In Bound DNS zone suffixes configured for the internal DNS server. For more information, see Internal DNS for Load Balancing.

> **Note:**
> - When you specify the DNS hostname, it is treated as a customized Service Principal Name (SPN), resulting in successful Kerberos authentication with this customized SPN.
> - If the DNS Hostname is modified when editing the machine accounts, all existing configuration of the machine account is overwritten on the Active Directory.

- **Encryption Types**—Optionally, select the encryption type(s) that must be used by the machine account. You can select multiple encryption types. The following encryption types are supported:
  - AES128-CTS-HMAC-SHA1-96
  - AES256-CTS-HMAC-SHA1-96
- **SPNs**—Service Principal Names (SPNs) are identifiers used to associate a service with a specific AD account. SPNs enable authentication and authorization for services. Specify the SPN in the following format:

  <service class>/<host>:<port>/<service name>

  The SPN syntax has four elements.
  - The following two elements are required:
    - **Service Class**: Specify the type of service.
    - **Host**: Specify the hostname or Fully Qualified Domain Name (FQDN) of the server on which the service is hosted.
  - If necessary, you can specify the following optional elements:
    - **Port**: Specify the port number on which the service is running.
    - **Service Name**: Specify the specific service instance on the host.

Note:

- All (or a subset) of the VIPs will resolve to the machine account. You can optionally add multiple machine accounts, which creates aliases for the same cluster. You can map these aliases to a subset (or all) of the VIPs to direct I/O to a subset (or all) of the nodes. Separate names with commas. For example: cluster1, cluster1-vip1, cluster1-vip2, cluster1-vip3, cluster1-vip4.
- If using VLANs—One of the machine accounts must be a DNS hostname for a virtual IP (VIP) of a VLAN configured on the cluster in order to recover to servers within that VLAN.
- After adding a Cohesity cluster to an AD, ensure the machine account used to register the Cohesity cluster in AD has the **Read membership** permissions across all users and groups.
- The same configuration must be done on every additional domain to which the Cohesity cluster is joined.
- The machine account must follow NetBIOS naming conventions of 15 maximum character length, those described in Naming conventions in Active Directory for computers, domains, sites, and OUs and the additional restriction that the ! @ ( ) characters are not supported.

8. **Use the above Machine Accounts even if they already exist in AD Domain:** Optionally toggle on to use the machine accounts that you specified above even if they already exist in the Active Directory domain.
9. **Mapped Provider:** Select the type of mapped provider to join the domain. By default, no mapped provider is selected. Optionally, select **LDAP** or **NIS**. If you select the **LDAP** option, you can select the provider from the **LDAP Provider** drop-down list. If you select the **NIS** option, you can select the NIS provider from the **NIS Provider** drop-down list.
10. **Organizational Unit:** Provide the organizational unit if applicable. Use the format: *OUName* or *OUName*/*SubOUName*. For example: Engineering or Sales/West or West Coast Sales/ North California.
11. **AD Workgroup / NetBIOS Name:** Provide the workgroup / NetBIOS name if applicable.
12. **Discover Trusted Domains:** Optionally, turn on this toggle to discover the trusted domains associated with an AD. You can exclude certain trusted domains after the

Cohesity cluster is joined to the AD domain. For more information, see Trusted Domains.

13. Click **Join**.
14. This step is required only for SMB access to views. On the DNS server that is integrated with the AD domain, add the A record for each machine account used to identify the Cohesity cluster to the AD domain. This is necessary to allow clients to access the Cohesity cluster as a resource. (Adding the A record is not generally necessary for Windows resources because Active Directory automatically adds it to the DNS server. However, the Cohesity cluster does not include the components that enable automatic addition to the DNS server.)
15. Optionally, repeat these steps to join the Cohesity cluster to another AD domain. For more information, see Guidelines for Joining Multiple AD Domains.
16. Optionally, you can add user ID mapping. For more information, see User ID Mapping for Multi Protocol Network Access.

Guidelines for Joining Multiple AD Domains

A cluster can join multiple Active Directory domains even if there are trust relationships between the AD domains.

After the Cohesity cluster joins the AD domain, you can add trusted domains to the allowlist or denylist. By default, all trusted domains are used for authentication unless added to the denylist. For more information, see Trusted Domains.

Only Kerberos authentication is supported for SMB shares when multiple AD domains are joined to the Cohesity cluster. NTLM authentication is disabled.

Leave Active Directory Domain

1. Go to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name and then click the leave domain icon ( ).
3. Provide the **AD Admin** and **Password** for the domain.
4. Click **Leave Domain**.

   When this succeeds, the machine accounts associated with the Cohesity cluster are removed from the Active Directory domain.

   > **Note:** If the Cohesity cluster is removed from a domain, ensure the corresponding DNS entries have been removed from the DNS server.

Enable Load Balancing

You can enable load balancing over multiple Cohesity cluster IPs by following these steps:

> **Note:** These example steps refer to the Windows DNS Manager. If you're using a different DNS Manager, ensure you perform step 3.

1. Open DNS Manager in Windows using the DNS Admin client: **Control Panel > Administrative Tools > DNS**.
2. Enable the Advanced view: **View > Advanced**.
3. Add multiple entries with the same Cohesity cluster name and do the following:
   - Give each entry a different **IP address** that corresponds to a VIP allocated to the Cohesity cluster in the DNS Zone where it will be added.
   - Generally, you can leave **Time to live (TTL)** at its default setting. You should only change this setting for specific use cases. Set TTL to a low value (such as 5 seconds or less) for backup use cases where a low number of clients are accessing the Cohesity cluster. Set TTL to a high value (such as 12 hours) for file services use cases with a high number of clients. This ensures that clients access the same Cohesity cluster nodes and benefit from caching.

Manage Domain Controllers From the Active Directory Details Page

You can also manage domain controllers from the Active Directory details page.

To manage domain controllers from the Active Directory details page:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the active directory for which you want to manage domain controllers.
   - In the **Details for AD Domain** section, click the edit icon corresponding to **Domain Controllers**:

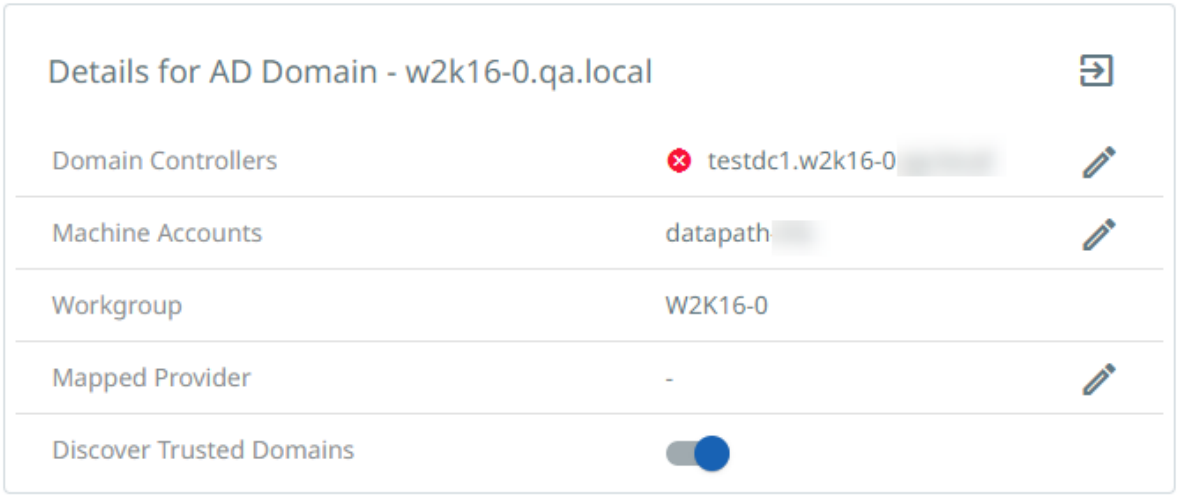

The **Manage Domain Controllers for AD Domain** dialog is displayed. You can manage the domain controllers for this AD domain.

- In the **Discovered Trusted Domains** section, the trusted domain controllers associated with the AD are discovered and listed. Click **Manage** corresponding to the domain name:

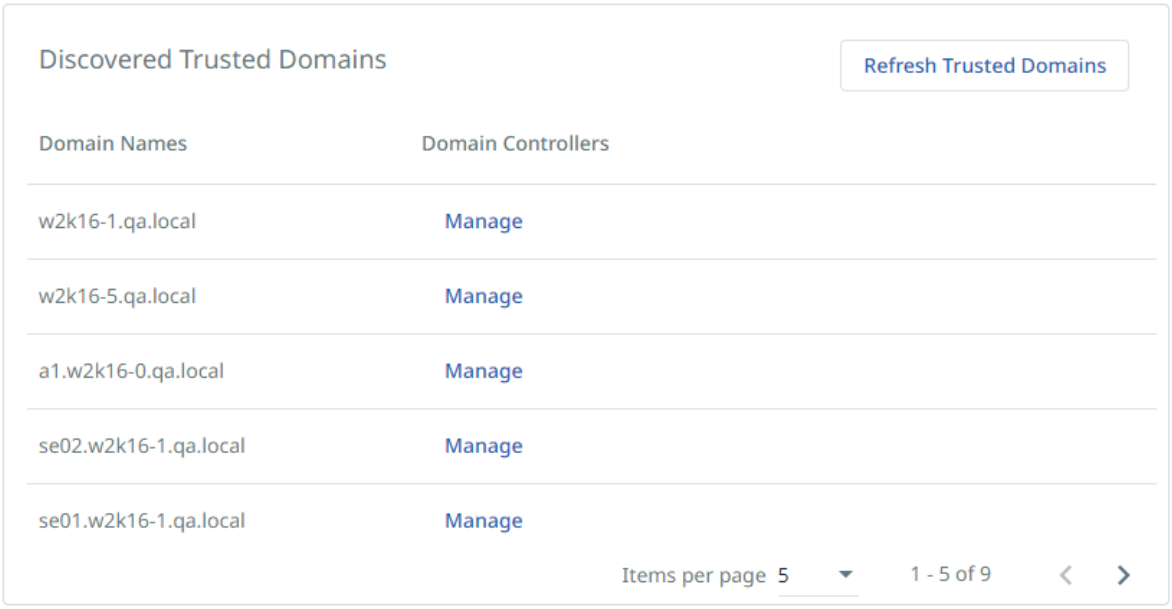

The **Manage Domain Controllers for AD Domain** dialog is displayed. You can manage the trusted domain controllers for the AD domain:

- **Allowed Domain Controllers**—All the domain controllers for the AD domain are listed in the **Allowed Domain Controllers** section. Drag a domain controller up or down to reorder the preference. To block a domain controller, hover over it and click **Deny**.
- **Denied Domain Controllers**—Domain controllers that are disallowed to connect the AD domain are listed in the **Denied Domain Controllers** section. To unblock a domain controller, hover over it and click **Allow**.

User ID Mapping for Multi Protocol Network Access

After the Cohesity cluster is joined to an Active Directory domain, you can add user ID mapping. The ID map database is stored in AD using industry-standard protocols. For more information, see ID Mapping for Multi-Protocol Views.

1. Go to **Settings > Access Management** and click the **Active Directory** tab.
2. Click **Add ID Mapping** or the edit icon ( ).
3. Select an **ID Mapping Type** and enter the required information:

- **Centrify**
  - **AD Domain:** Name of the Active Directory domain where the corresponding Centrify zone is located.
  - **Centrify Zone:** Selected zone that has the ID mappings.

  For Native security style Views, SMB will get Unix IDs from the Active Directory provider and NFS will get Unix IDs from the LDAP provider.

  For Unified or NTFS security style Views, SMB and NFS will get Unix IDs from the corresponding Centrify zone for the Active Directory provider.
- **Fixed**
  - **Unix UID to be applied to all files created:** Enter a user ID to apply to all created files.
  - **Unix GID to be applied to all files created:** Enter a group ID to apply to all created files.
- **LDAP Provider**

  The UID and GID corresponding to the Active Directory users are fetched from the configured LDAP provider. In LDAP, the most commonly used attributes to store the UID and GID are 'uidNumber' and 'gidNumber', respectively.

  All Views, regardless of their security style setting (Native, Unified or NTFS), will get Unix IDs from the LDAP provider and SIDs from the Active Directory provider.
- **RID (Relative ID of SID)**

  RID is a variable length number that is assigned to objects at creation and becomes part of the object's Security Identifier (SID) that uniquely identifies an account or group within a domain. For every Windows user or group, the RID of the corresponding SID is used as the UID or GID.

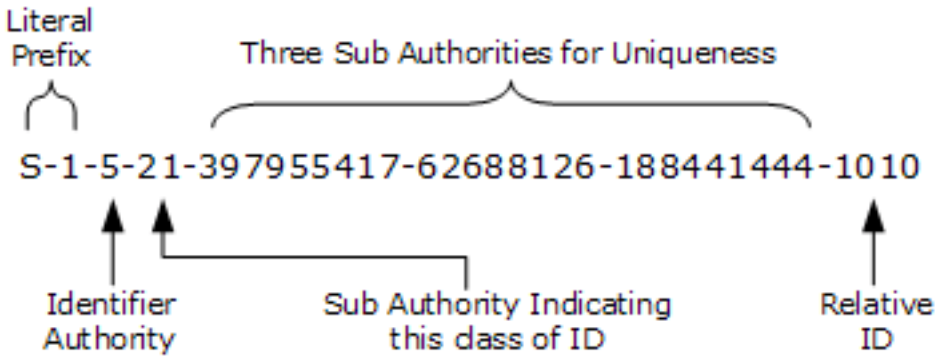

- **RFC2307**

  In RFC2307, the UID and GID are stored along with Active Directory objects in the following attributes: the UID for a user is stored in the attribute ‘uidNumber’ and GID for a group is stored in the attribute 'gidNumber'.

  For Native security style Views, SMB will get Unix IDs from the Active Directory provider and NFS will get Unix IDs from the LDAP provider.

  For Unified or NTFS security style Views, SMB and NFS will get Unix IDs from the Active Directory provider.

- **SFU (Services for Unix 3.0)**

  In SFU, the UID and GID are stored along with Active Directory objects in the following attributes: the UID for a user is stored in the attribute ‘uidNumber’ and GID for a group is stored in the attribute 'gidNumber'.

  For Native security style Views, SMB will get Unix IDs from the Active Directory provider and NFS will get Unix IDs from the LDAP provider.

  For Unified or NTFS security style Views, SMB and NFS will get Unix IDs from the Active Directory provider.

4. Select a **Fallback Option** to use if the selected ID Mapping Type is, for instance, not populated. This is available if the ID Mapping Type is not Fixed or RID.

   If you select Fixed for the Fallback Option, provide the following information:

   - **Unix UID to be applied to all files created:** Enter a user ID to apply to all created files.
   - **Unix GID to be applied to all files created:** Enter a group ID to apply to all created files.
5. Select the **AD User mapped to the Linux Root User** by selecting the AD domain and user.
6. Click **Save**.

Trusted Domains

After the Cohesity cluster joins the Active Directory domain, you can add trusted domains to the allowlist or denylist. By default, all trusted domains are used for authentication unless added to the denylist. You can add trusted domains to the allowlist so that only trusted domains added to the allowlist are used for authentication.

> **Note:** If the domain controller is not functioning correctly, the discovery of trusted domains may take a relatively long period of time or may provide inconsistent results. Therefore, a trusted domain might not be discovered at all. You must manually enter such domains in the allowlist or denylist.

A trusted domain forest can have multiple trusted domains. To provide optimized access, you can use a subset of trusted domains so that the Cohesity cluster does not send group membership queries to all trusted domains. You can choose a subset of trusted domains by using one of the following modes:

- **Denylist**—Ignores specified trusted domains. For example, if there are three trusted domains T1, T2, and T3 and if T2 is added to the denylist, T1 and T3 are the trusted domains used for authentication.
- **Allowlist**—Allows only specified trusted domains. For example, if there are three trusted domains T1, T2, and T3 and if T1 and T2 are added to the allowlist, T1 and T2 are the trusted domains used for authentication.

Add Trusted Domains to the Allowlist

To add a trusted domain to the allowlist:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page.
3. Click the edit icon in the **Denied Trusted Domains** section.
4. Click **Allowlist** and add the domain name in the **Allowed Trusted Domains** field. Click + to add another domain.

   Cohesity recommends adding at least one domain or switching to a denylist.
5. Click **Save**.

   The trusted domain is added to the allowlist.

Add Trusted Domains to the Denylist

If you want Cohesity to ignore specific trusted domains so that their principals are not assigned roles and denied access to Cohesity SMB shares, you can add the desired domains to the denylist.

To add a trusted domain to the denylist:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page.
3. Click the edit icon in the **Denied Trusted Domains** section.
4. Click **Denylist** and add the domain name in the **Denied Trusted Domains** field. Click + to add another domain.
5. Click **Save**.

   The trusted domain is added to the denylist.

Refresh Trusted Domains

If the **Discover Trusted Domains** option is enabled, the trusted domains associated with the AD are discovered and listed in the **Discovered Trusted Domains** section. After any

configuration change, trusted domains are not automatically refreshed in the **Discovered Trusted Domains** section. You must manually refresh the trusted domains.

To refresh the trusted domains list:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page:
3. In the **Discovered Trusted Domains** section, click **Refresh Trusted Domains**.

Add or Edit Preferred Domain Controllers

To add or edit preferred domain controllers of the discovered trusted domains:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the domain name to navigate to the AD domain details page:
3. In the **Discovered Trusted Domains** section, click **Select** corresponding to the domain name.
4. On the **Edit Preferred Domain Controllers** dialog, modify the domain controllers and click **Save**.

View Trust Relationships for Domains

To view the trust relationships between the domains:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the **View Trust Relationships** button. The topology of the trust relationships between domains is displayed.This page helps you analyze complex AD topologies with multiple domains and trusts.

   You can view the following details for the trust relationships:

   - **Domain Category**: Primary and Trusted domains. The domains are color-coded for easy differentiation.
   - **Relationship Direction**: One-way trust and Two-way trust.
   - **Filters**: Primary Domains, Trusted Domains, and Trusted Relationships.

     You can also drag the AD nodes to reorganize the topology.

One-Way Trust

A one-way trust is a unidirectional authentication path created between two AD domains. In a one-way trust, the trust relationship is only one way.

The following is an illustration of one-way trust and the direction of access:

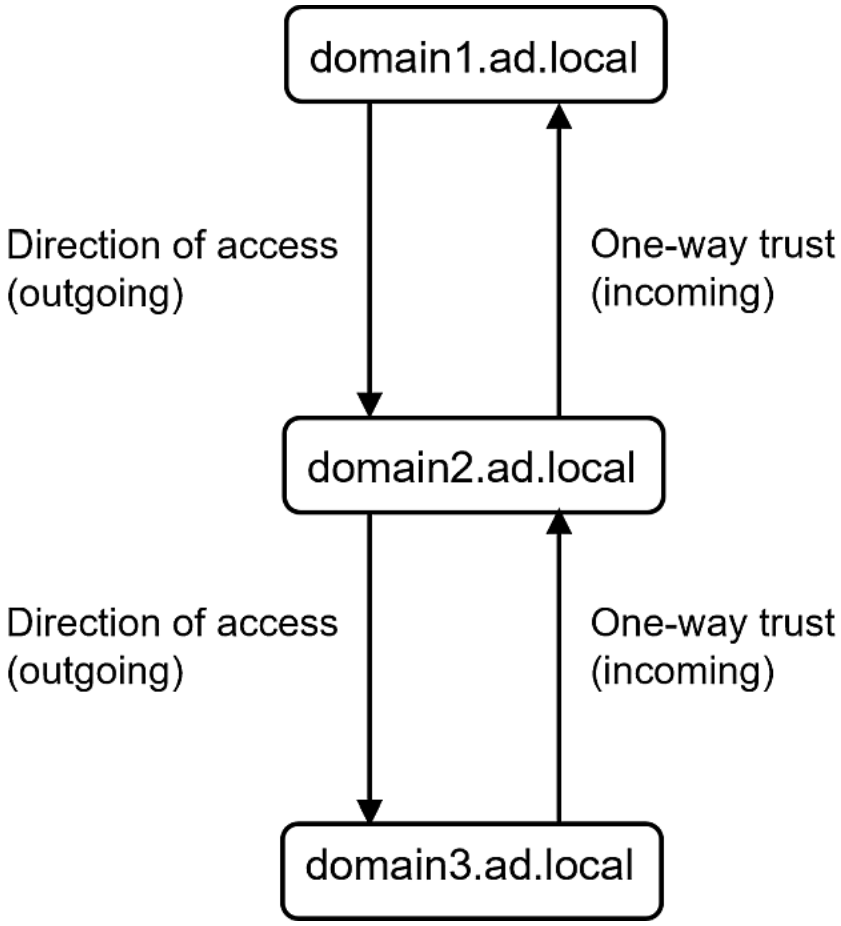

As shown in the above figure:

- domain2.ad.local has a one-way trust with domain1.ad.local. domain1.ad.local does not have a trust relationship with domain2.ad.local. Users in domain2.ad.local can access resources in domain1.ad.local. However, users in domain1.ad.local cannot access resources in domain2.ad.local.
- domain3.ad.local has a one-way trust with domain2.ad.local. domain2.ad.local does not have a trust relationship with domain3.ad.local. Users in domain3.ad.local can access resources in domain2.ad.local. However, users in domain2.ad.local cannot access resources in domain3.ad.local.

In a one-way trust, the direction of trust is the opposite of the direction of access. An outgoing trust allows incoming access and an incoming trust allows outgoing access.

If one-way trust is enabled on the AD domain, trusted domains are listed in the following way in the Cohesity cluster:

- If the Cohesity cluster joins domain2.ad.local and if trusted domains are enabled:
    - domain1.ad.local is displayed as a trusted domain because domain1.ad.local trusts domain2.ad.local

    - domain3.ad.local is not displayed as a trusted domain because domain3.ad.local does not trust domain2.ad.local
- If the Cohesity cluster joins domain3.ad.local, both domain2.ad.local and domain1.ad.local are displayed as trusted domains because domain2.ad.local trusts domain3.ad.local and domain1.ad.local trusts domain2.ad.local.
- If the Cohesity cluster joins domain1.ad.local, no domain is displayed as a trusted domain because domain1.ad.local does not have a trust relationship with domain2.ad.local or domain3.ad.local.

View Active Directory Join and Leave Status

You can add one or more Active Directory (AD) domains to a Cohesity cluster. Joining the AD domain to a Cohesity cluster is a multi-step process and involves the following steps:

- Connecting to the DNS server
- Getting the list of domain controllers from the DNS server
- Connecting to the DC server
- Adding machine account(s)
- Setting the encryption type

You can view the progress of each subtask. For example, while joining an AD domain, you can check whether the machine account was added successfully. This log helps in troubleshooting and identifying the exact step at which the task failed.

View Active Directory Join Status

To view the AD join status:

1. Follow the steps described on the Join Active Directory page to join the Cohesity cluster to the AD domain.
2. After you click **Join**, the following error message is displayed if the process fails:

   ```
   Joining Active Directory <ad-domain-name> failed. Show Subtasks
   ```

3. Click **Show Subtasks**.

   A dialog is displayed and the subtasks are listed.

   You can view the subtask at which the process failed.

You can also view the status when you edit an AD domain. For example, if you add or modify a machine account after the Cohesity cluster joins the AD domain and if the process fails, you can view the status by following the steps outlined below:

1. Navigate to **Settings > Access Management** and click the **Active Directory** tab.
2. Click the AD domain that you want to edit.

   The AD Domain details page is displayed.

3. Click the edit icon corresponding to **Machine Accounts**.
4. Modify the machine account details as necessary.
5. Click **Save**.

   The following error message is displayed if the process fails:

   ```
   Failed to join Active Directory. Show Subtasks
   ```

6. Click **Show Subtasks** to view the subtask at which the process failed.

View Active Directory Leave Status

Leaving an AD domain is a multi-step process. You can view the progress of each subtask. For example, while leaving an AD domain, you can check whether the machine account was removed successfully. This log helps in troubleshooting and identifying the exact step at which the task failed.

To view the AD leave status:

1. Follow the steps described in the Leave Active Directory Domain section to remove the Cohesity cluster from an AD domain.
2. Click **Leave Domain**.

   The following error message is displayed if the process fails:

   ```
   Failed to leave Active Directory. Show Subtasks
   ```

3. Click **Show Subtasks** to view the subtask at which the process failed.

Users and Groups

Cohesity allows adding users with different privileges by assigning them various roles within the cluster. These users and groups can manage data and view statistics based on their assigned roles.

Cohesity supports the following types of users and groups. Except for the **support** user, all user types can log in to the Cohesity Dashboard to manage data on the Cohesity cluster and view statistics, depending on their role.

- **Local users:** Local users exist on the Cohesity cluster only.
- **Local groups:** You can define groups by adding local users and Active Directory principals.
- **SSO users:** If the Cohesity cluster is configured to use an identity provider for single sign-on access, you can add SSO users and assign them roles. For more information, see SSO Users.
- **Active Directory principals:** If the Cohesity cluster is joined to an AD domain, you can assign roles to principals in the AD or the AD's trusted domains, which determine the Cohesity cluster access. For more information, see Active Directory (AD).

- **Support user:** Used to access the bash shell of the Cohesity cluster using SSH. The support user account is disabled by default. For more information, see Manage Support User.

Types of Users and Groups

Local Users

Local users are exclusive to the Cohesity cluster. They can be created and managed directly within the system.

Add or Edit a Local User

1. Navigate to **Settings > Access Management**.

   In the **Users & Groups** tab, from the **Add AD Users & Groups** drop-down list, select **Add Local User**, or click the action menu and select **Edit** to edit a local user.

2. Enter the following information:

Add Local User	Enter the user name. Do not enter whitespace characters such as spaces. Do not specify any of these characters: `/ < > : #`
Email Address	Enter the email address associated with this user.
Password	Enter a password that is at least eight characters.
Confirm Password	Enter the same password again to confirm.
Groups	Add the user to one or more user groups by selecting the groups from the drop-down list
Roles	(Optional) Select the appropriate roles. For a detailed description of each role, see Manage Roles.
Effective Date	Specify the date that the user can begin using the account.
Expiration Date	Specify the date when the user's access will be revoked.

Restrict access to specific Objects	Optionally restrict access to specific Objects and Views only by toggling this option on. By default, the allowed actions apply to all Source Objects registered with the Cohesity cluster and all Views. After restricting access, you must specifically add Objects and Views through **Add Registered Object(s)** and **Add Cohesity View(s)**. The user or group does not have access to any Objects or Views until you add them. When a user has access to an Object, that user also has access to all previous Snapshots of that Object and any Objects added to the selected Source folder in the future. For example, if a user has access to an ESX host, they will also have access to VMs added to the host later. Users that have access restricted to specific Objects cannot modify users or Sources even if their role grants them these privileges. Object permissions are replicated for Virtual Machine Objects but not currently for physical Server and View Objects. The Object permissions for a local user or group are not replicated to the remote Cohesity cluster. Local users inherit the permissions of the Restricted Objects provided to the Restricted Groups to which the users belong. Users that have access restricted to specific objects cannot update the Protection Group created by an admin user even if the restricted user has access to all the objects in the Protection Group.
Set Primary User Group for File Access	Enable this option for the user to access file services. From the **Primary Group** drop-down list, select a group as a Primary group. The user can be a part of one or more groups and the SMB/NFS protocol requires a primary group to access the files.

3. Click **Add** or **Save Changes** if editing a user.

Local Groups

Local groups consist of local users and Active Directory (AD) principals. These groups can be defined and managed within the Cohesity cluster.

Add or Edit a Local Group

1. Navigate to **Settings > Access Management**.
2. In the **Users & Groups** tab, from the **Add Users & Groups** drop-down list, select **Add Local Group**, or click the action menu and select **Edit** for a local group.

3. Enter the following information:

Group Name	Enter the group name. Do not enter whitespace characters such as spaces. Only letters, numbers and the special characters given below are allowed: _ - .
Local User	Add one or more users to the group by selecting the users from the drop-down list.
Roles	Select the appropriate roles from the drop-down list. For more information on roles, see Manage Roles.
Effective date	Specify the date when the user group can begin using the account
Restrict access to specific Object	Optionally restrict access to specific Objects and Views only by toggling this option on. By default, the allowed actions apply to all Source Objects registered with the Cohesity cluster and all Views. After restricting access, you must specifically add Objects and Views through **Add Registered Object(s)** and **Add Cohesity View(s)**. The user or group does not have access to any Objects or Views until you add them. When a user has access to an Object, that user also has access to all previous Snapshots of that Object and any Objects added to the selected Source folder in the future. For example, if a user has access to an ESX host, they will also have access to VMs added to the host later. Users that have access restricted to specific Objects cannot modify users or Sources even if their role grants them these privileges. Object permissions are replicated for Virtual Machine Objects but not currently for physical Server and View Objects.
Set Primary User Group for File Access	Enable this option for the user to access file services. From the **Primary Group** drop-down list, select a group as a Primary group. The user can be a part of one or more groups and the SMB/NFS protocol requires a primary group to access the files.

4. Click **Add**, or **Save Changes** if editing a group.

Add a Local Administrator Group Using the CLI

You can use local administrator groups to grant SMB access to views and/or administrative access to the Cohesity cluster.

The Cohesity cluster has a built-in local administrators group called Administrators. By default, the group has only SMB Security privileges and the standard Windows administrator Security ID (SID) S-1-5-32-544. You can use and modify the Administrators group, or create your own local administrator groups and assign roles and SIDs by using the Cohesity CLI.

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Use the following command syntax to add a local administrator group.

   ```
   group add name=<group-name> smb-only=true roles=<cohesity_roles>
   ```

   (You can use the `role list` command to list the roles defined in the Cohesity cluster.)

   For example:

   ```
   admin@node1> group add name=admin2 smb-only=true roles=cohesity_admin
   GROUP NAME                    : admin2
   DOMAIN                        : LOCAL
   ROLES                         : COHESITY_ADMIN
   SMB PRINCIPAL ENABLED         : true
   ```

Use the following command syntax to add a user or group.

```
group add-smb-principal name=<group_name> type=<user or group> domain=<AD_
domain> principal-name=<user or group> sid=<SID>
```

For example:

```
admin@node1> group add-smb-principal name=admin2 type=group
domain=myorg.ad.domain.com principal-name=PIMS sid=S-1-5-21-4152169185-
4221176486-4244366766-1001
GROUP NAME                    : admin2
DOMAIN                        : LOCAL
ROLES                         : COHESITY_ADMIN
SMB PRINCIPAL ENABLED         : true
SMB PRINCIPAL                 :
NAME                          : PIMS
DOMAIN NAME                   : myorg.ad.domain.com
SID                           : S-1-5-21-4152169185-4221176486-4244366766-
1001
TYPE                          : group
```

> **Tip:** Use the CLI Help by running the following command: `group add-smb-principal help`.

CLI Option	Example	Description
`domain=`	`org.ad.domain.com`	Optional. The AD domain name that includes the principal you are adding.
`name=`	`admin2`	Required. Your choice of a name for the local administrator group.
`principal-name=`	`group1` `jsmith`	Optional. The name of the SMB principal user or group. The name must exist in the AD domain.
`sid=`	`S-1-100-21-8372`	Optional. A unique Security ID for the principal.
`type=`	`group` or `user`	Optional. Specify `user` or `group` to indicate the type of principal you are adding.

Enable User Account Lockout

The admin user can enable account lockout on the user account. User accounts can get locked after a threshold number of failed login attempts. The user cannot log in to the Cohesity Dashboard after three consecutive invalid login attempts for 15 minutes.

> **Note:** You can enable account lockout only for non-admin local users.

Perform the following steps to enable user account lockout:

1. Navigate to **Settings > Access Management**.
2. Select the **Users & Groups** tab.
3. Select the user, click the actions menu (⋮) and select **Lock**.

Disable User Account Lockout

The admin user can disable the account lockout on a user account.

Perform the following steps to disable user account lockout:

1. Navigate to **Settings > Access Management**.
2. Select the **Users & Groups** tab.
3. Select the user, click the actions menu (⋮) and select **Unlock**.

Delete Users and Groups

You can delete users and groups.

> **Note:** You cannot delete default users and groups.

Perform the following steps to delete user and group:

1. Navigate to **Settings > Access Management**.
2. Select the **Users & Groups** tab.
3. Select the user, click the actions menu (⋮) and select **Delete**.
4. Click **Delete**.

Change Local User Password

Using the Cohesity Dashboard, you can change your local user password. By default, the new password must replace at least eight characters with the current password.

Using the following Cohesity CLI command, you can define the number of characters you have to replace in your existing local user password:

```
iris_cli security-config update num-diff-chars=<number_of_characters>
```

For example:

If you want to restrict those six characters in your existing local user password must be replaced while changing the password, then use the following command:

```
iris_cli security-config update num-diff-chars=6
```

> **Note:** The restriction is not applicable if the admin user wants to change the password of other local users in the Cohesity cluster.

To change the local user password:

1. Navigate to **Settings** > **Access Management**.
2. Click the actions menu (⋮) next to the user and click **Edit**.
3. In the **Edit Local User: <user>** page:
    1. Enable **Change Password**.
    2. Enter the **Current Password** and then enter and confirm the new **Password**.
    3. Click **Save Changes**.

Manage LDAP Providers

A Unix user belonging to a primary group can also be a member of many other groups defined in the database. Because of an NFS protocol limitation, NFS users who belong to

more than 16 groups have access issues.

You can add an LDAP provider to the Cohesity cluster to provide access to NFS exports for Unix users belonging to more than 16 groups. The LDAP server provides a central repository to store user identities in a network security system. The Unix group information (GIDs) is stored on the LDAP server. When a Unix user accesses data through NFS on Cohesity, the GIDs corresponding to the users are fetched from the configured LDAP server using the credentials configured by the administrator from the Cohesity GUI.

> **Note:** Cohesity does not support user authentication to the Cohesity Dashboard through LDAP.

If the Cohesity cluster is joined to an Active Directory domain, you can map the AD domain to an LDAP provider across SMB and NFS protocols so users can access the same files on the Cohesity cluster from both Windows and Linux systems. Mapping happens based on the user name in Windows and the user ID in Unix.

> **Note:** Cohesity supports SASL over LDAP by default. If you are seeking stronger security, we recommend you to refer Secure LDAP for LDAP over SSL.

Prerequisites

- Cohesity supports LDAP version 3.
- The port that the LDAP provider listens on must be open. The default port is 389 (outgoing, TCP/UDP).
- Only administrators or users who have the **Manage AD and LDAP** role assigned can add or modify the LDAP provider.

Add an LDAP Provider to the Cohesity cluster

1. Navigate to **Settings > Access Management** and select the **LDAP** tab.
2. Click **Add LDAP Provider**.
3. Specify the name of the LDAP provider. You can change the name later.
4. **LDAP Servers**—Identify the LDAP server to use.
    - **IP or FQDN**—The IP addresses or fully qualified domain names for the LDAP server.
    - **Domain Name**—The domain name to use when querying LDAP from DNS.
    - **Port**—The port on which the LDAP server is listening. The default port is 389.
        - **Base Distinguished Name**—The base distinguished name to use for the LDAP server.
5. **Authentication Type**—Select the authentication type to use:

- **Simple**—The Bind Distinguished Name and bind password to use for authentication to the LDAP server. Indicate whether to use SSL for LDAP connections.
- **Anonymous**—There is no username and password provided to bind to an LDAP server. The bind session has no user associated with it.

6. **Mapped AD Domain**—Select the Active Directory domain to which you want to map the LDAP server.
7. Click **More Options**.
8. **Additional Attributes**—Default values are populated for the following attributes. If required, you can modify the attributes:
   - Common Name Attribute
   - UID Attribute
   - GID Attribute
   - Member of Attribute
   - Username Attribute
9. **Group Query**—A default value is populated. If required, you can modify the **Group Object Class**.
10. **User Query**—A default value is populated. If required, you can modify the **User Object Class**.
11. Click **Add**.

    The LDAP provider is added and the status is displayed on the **Access Management > LDAP** page.

Edit an LDAP Provider

To edit an LDAP Provider:

1. Navigate to **Settings > Access Management** and select the **LDAP** tab.
2. Hover over the LDAP provider that you want to edit, click ⋮, and select **Edit**.
3. On the **Edit LDAP Provider** page, modify the settings as required.
4. Click **Save**.

Delete an LDAP Provider

To delete an LDAP Provider:

1. Navigate to **Settings > Access Management** and select the **LDAP** tab.
2. Hover over the LDAP provider that you want to delete, click ⋮, and select **Delete**.
3. In the **Delete LDAP Provider** dialog, enter **YES** to confirm.
4. Click **Delete**.

Secure LDAP

Lightweight Directory Access Protocol (LDAP) is a protocol used for interacting with Active Directory servers. By default, LDAP traffic is not encrypted. The Cohesity cluster uses the LDAP service to:

- Search users and groups
- Get the group membership of users and groups
- Get the SID of a user or group

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

This document describes how to enable LDAPS after the Cohesity cluster is joined to AD, including certificate installation and LDAP bind credential configuration.

Ports for LDAP

Port 389 uses LDAP with Simple Authentication and Security Layer (SASL), which encrypts user and group data within LDAP messages. For security policies, this provides protection comparable to Lightweight Directory Access Protocol over SSL/TLS (LDAPS).

Port 636 is the standard network port used for LDAPS. Ensure to enable LDAPS on Cohesity only when your security policy specifically requires LDAP over SSL on port 636.

Method	Port	Role
LDAP with SASL	389	Always required. AD and the Cohesity cluster use this port during normal operation.
LDAPS	636	Uses TLS from the start of the connection. After LDAPS is enabled, most Cohesity user and group lookups use this port.

Considerations

- Port 389 must remain open at all times between the Cohesity cluster and the AD domain controllers. Do not block port 389 and open it only for occasional AD changes.
- Windows AD and Cohesity continue to use port 389 during normal operation. If port 389 is closed, Cohesity cannot bypass the restrictions enforced by Windows AD.
- After you enable LDAPS, most Cohesity LDAP lookups (like SMB session setup) use port 636. Some LDAP traffic continues to use port 389. Ensure not to force all LDAP traffic to port 636 in firewalls or AD, since it can break the functionality.

- AD operations like joining AD, changing AD settings, and leaving AD requires port 389. Other ongoing cluster and Windows AD communication also uses port 389.
- LDAPS uses simple bind with LDAP bind credentials. Kerberos bind and NTLM bind are not supported on port 636.
- LDAPS port settings are global per cluster. You cannot configure one AD to use port 389 and another AD to use port 636.
- Ensure to install an SSL certificate on every cluster node.

LDAP Bind Credentials

Bind operations authenticate the Cohesity cluster to AD when opening an LDAPS connection. LDAP bind credentials are stored in the Advanced Encryption Standard (AES) format in the Cohesity cluster. They are used when the cluster needs to create a new LDAP connection to AD. LDAP bind credentials must be the same for all trusted domains.

Use the following CLI to update bind credentials periodically:

```
iris_cli ad update-ldap-bind-credentials
domain-name=<example>
user-name=<test>
password=<test123>
```

After you update the bind credentials, you must restart the I/O Operations service using the following command:

```
admin@198.51.100.12>cluster restart
service-names=bridge
```

Use the following CLI to clear LDAP bind credentials:

```
iris_cli ad update-ldap-bind-credentials
domain-name=<example> clear-credentials={}
```

Install SSL Certificates

To enable LDAPS, you must install an SSL certificate on every node in the Cohesity cluster.

> **Important:** Use the Host Shell to perform the steps to install the SSL certificate on every node in the Cohesity cluster.

> **Note:** Ensure to upload the root CA certificate for your Active Directory domain. Do not use individual Domain Controller (DC) certificates. Using the root CA certificate ensures that Secure LDAP continues to work when domain controllers are added, removed, or replaced.

To install the certificate:

1. Create a `.pem` file from the certificate file imported from the Active Directory domain controller (not an individual DC certificate). You can perform this step on any Linux system on which the OpenSSL tool is available. The filename must be in lowercase letters and be the same as the AD domain FQDN:

   ```
   $ openssl x509 -text -in <certificate>.cer -out <example>.pem
   ```

2. If the following directories do not exist, create these directories on the Cohesity cluster node(s):

   ```
   /home/cohesity/data/bridge/openldap/certs
   /home/cohesity/data/bridge/openldap/cacerts
   ```

3. Copy the `.pem` file to the `/home/cohesity/data/bridge/openldap/certs` directory.
4. Generate the certificate hash from the `.pem` file:

   ```
   $ openssl x509 -in <example>.pem -hash -noout
   ```

   For example, the certificate hash can be `75b8ef63`.

5. Use the certificate hash to create symbolic links in the `cacerts` directory and point it to the `certs` directory:

   ```
   $ ln -s /home/cohesity/data/bridge/openldap/certs/<example>.pem
   /home/cohesity/data/bridge/openldap/cacerts/75b8ef63.0
   ```

   > **Note:** .0 indicates the version number of the certificate. Each hash is identified by the subject hash. You can install multiple certificates from the same or different AD. The OpenLDAP library picks the correct certificate file for verification using this hash.

Chapter 10: Cloud Platforms

This chapter includes the following topics:

- AWS 1549
- Troubleshooting 1617
- Execute the CloudFormation Template 1630
- Amazon S3 Entity Hierarchy 1714
- What are Materialized Views 1775
- Why Refresh is Needed After Recovery 1775
- Materialized View Refresh Options 1775
- Best Practices 1778
- What are Materialized Views 1826
- Why Refresh is Needed After Recovery 1826
- Materialized View Refresh Options 1826
- Best Practices 1829
- Requirements 1930
- Considerations 1930
- Register SQL Server on Linux OS 1931
- Azure 2093
- GCP 2341

AWS

Cohesity provides a simple and fast backup, recovery, and data management solution for Amazon Web Service (AWS) workloads.

Cohesity supports the following solutions on AWS:

- Cloud Edition

  Deploy Cohesity clusters in AWS to backup workloads to AWS.

- Backup and Recovery

  Back up and recover Amazon services, such as Amazon EC2, Amazon S3, and Amazon RDS, using Cohesity clusters.

- Disaster Recovery

  Disaster recovery failover and failback of VMware VMs backed up using the on-prem Cohesity cluster to AWS using Cloud Edition.

- Replication

  Replicate data from on-premise Cohesity cluster to Cloud Edition in AWS.

- Clone VMs Using CloudSpin

  Convert VMware VMs backed up using on-premises Cohesity cluster to Amazon EC2 instances.

- CloudArchive

  Archive data backed up using Cohesity cluster to Amazon S3 for long-term retention.

This chapter includes the following topics:

Supported AWS Regions

Cohesity supports the backup and recovery of AWS workloads (such as Amazon EC2, S3, RDS, and Aurora) across all AWS regions, including AWS GovCloud. For the list of supported regions. For more details on these regions, see Regions, Availability Zones, and Local Zones.

AWS Permissions

This topic provides information on the permissions required to backup, recover and to perform other supported workflows for AWS workloads:

- Minimum Permissions for Individual Workloads and Workflows
- Comprehensive Permissions List for AWS Source Protection
- Assign AWS Permissions

Minimum Permissions for Individual Workloads and Workflows

This topic lists the minimum permissions required for protecting individual workload. If you intend to protect only a specific workload, then you can choose to remove the permissions related to other workloads in the policy.

Permissions for Amazon EC2 Protection

Cloud Snapshot Manager-based Protection

Open the JSON (or right-click and save as a file).

AWS Resource type	Permissions	Description
EC2	AttachVolume CopySnapshot CreateSnapshot CreateTags CreateVolume DeleteSnapshot DeleteVolume DeregisterImage DescribeAvailabilityZones DescribeImages DescribeInstanceAttribute DescribeInstanceCreditSpecifications DescribeInstanceStatus DescribeInstances DescribeKeyPairs DescribeRegions DescribeReservedInstances DescribeReservedInstancesListings DescribeReservedInstancesModifications DescribeReservedInstancesOfferings DescribeSecurityGroupReferences DescribeSecurityGroups DescribeSnapshots DescribeSubnets DescribeVolumeAttribute DescribeVolumeStatus DescribeVolumes DescribeVpcAttribute DescribeVpcClassicLink DescribeVpcClassicLinkDnsSupport DescribeVpcEndpointConnectionNotifications DescribeVpcEndpointConnections DescribeVpcEndpointServiceConfigurations DescribeVpcEndpointServicePermissions DescribeVpcEndpointServices DescribeVpcEndpoints DescribeVpcPeeringConnections DescribeVpcs	Basic EC2 permissions for Backup and Recovery

AWS Resource type	Permissions	Description
	DetachVolume ModifySnapshotAttribute RegisterImage RunInstances StartInstances StopInstances	
IAM	AddRoleToInstanceProfile GetAccessKeyLastUsed GetInstanceProfile GetUser GetUserPolicy ListAccessKeys PassRole	Cohesity uses the secret access key and access key provided to validate the register the source on the cluster. While IAM user registration the credentials for the user needs to be fetched and validated against the credentails provided. The policy of the user needs to be fetched after registration. IAM permissions are required for these actions. "iam:AddRoleToInstanceProfile" is required to register source using IAM role.
KMS	CreateGrant Decrypt DescribeKey Encrypt GenerateDataKey ListAliases ReEncryptFrom ReEncryptTo	KMS permissions are required while backing up and recovering an encrypted AWS EC2 instance.

Cloud Native Backup and Recovery

Open the JSON (or right-click and save as a file).

AWS Resource Type	Permissions	Description
IAM	AddRoleToInstanceProfile AttachRolePolicy CreateRole DeleteRole DeleteRolePolicy GetAccessKeyLastUsed GetInstanceProfile GetRole GetRolePolicy GetUser GetUserPolicy ListAccessKeys ListRolePolicies ListRoleTags ListRoles PassRole PutRolePolicy TagRole UntagRole UpdateRole UpdateUser	Cohesity uses the secret access key and access key provided to validate the register the source on the cluster. During IAM user registration, the credentials for the user are fetched and validated against the credentails provided. The policy of the user needs to be fetched after registration. IAM permissions are required for these actions. Cohesity creates an S3 bucket and attaches a policy and role to it while recovering the VM back to AWS. The boot disk of the VM is put, as an object, into the bucket which is used for importing the VM disk as an EC2 Image. The AWS EC2 instance is recovered by launching this EC2 Image. IAM permissions are required to create this role and policy and attach it to the bucket. Delete Role and Delete Role policy are required to delete the role and policy created and attached to the S3 bucket after the task is complete. "iam:AddRoleToInstanceProfile" is required to register source using IAM role.
S3	CreateBucket DeleteObject GetBucketLocation GetBucketTagging GetObject ListBucket ListBucketVersions PutBucketTagging PutBucketPublicAccessBlock PutObject	Cohesity creates an S3 bucket and attaches a policy and role to it while recovering the VM back to AWS. This boot disk of the VM is put, as an object, into this bucket which is used for importing the VM disk as an EC2 Image. The AWS EC2 instance is recovered by launching this EC2 Image. S3 permissions are required to create this bucket and put the VM disk as an object into this bucket.

AWS Resource Type	Permissions	Description
KMS*	CreateGrant Decrypt DescribeKey Encrypt GenerateDataKeyWithoutPlaintext ListAliases ReEncryptFrom ReEncryptTo	KMS permissions are required while backing up and recovering an encrypted AWS EC2 instance.

AWS Resource Type	Permissions	Description
EC2	AllocateAddress AssociateSubnetCidrBlock AssociateVpcCidrBlock AttachInternetGateway AttachNetworkInterface AttachVolume AuthorizeSecurityGroupEgress AuthorizeSecurityGroupIngress CopySnapshot CreateImage CreateSecurityGroup CreateSnapshot CreateTags CreateVolume DeleteSnapshot DeleteVolume DescribeAvailabilityZones DescribeImportSnapshotTasks DescribeInstanceAttribute DescribeInstanceStatus DescribeInstances DescribeKeyPairs DescribeRegions DescribeReservedInstances DescribeReservedInstancesListings DescribeReservedInstancesModifications DescribeReservedInstancesOfferings DescribeSecurityGroupReferences DescribeSecurityGroups DescribeSnapshotAttribute DescribeSnapshots DescribeSubnets DescribeVolumeAttribute DescribeVolumeStatus DescribeVolumes DescribeVpcAttribute DescribeVpcClassicLink	These permissions are required to register the AWS IAM user on the Cohesity cluster. Once the IAM user is registered on the cluster, describe permissions on every EC2 service are required for the user to successfully select from all the available options for each EC2 services from the drop down while performing any of the AWS workflows. Cohesity creates proxy EC2 instances for doing back up and recovery of AWS EC2 instance. Cohesity creates snaphots of the EC2 volumes while backing up and stores the different instance attributes and tags. While recovering the AWS EC2 instance, Cohesity creates volumes of original disk size and imports ab image of the boot disk from the AWS S3 bucket. It also attaches the original tags and corresponding network and security groups as part of the recovery. Cohesity requires the delete volume and snapshots permissions to delete the intermediate volumes and snapshots it creates during the backup and recovery from the proxy EC2 instances.

AWS Resource Type	Permissions	Description
	DescribeVpcClassicLinkDnsSupport DescribeVpcEndpointConnectionNotifications DescribeVpcEndpointConnections DescribeVpcEndpointServiceConfigurations DescribeVpcEndpointServicePermissions DescribeVpcEndpointServices DescribeVpcEndpoints DescribeVpcPeeringConnections DescribeVpcs DetachVolume ImportImage ImportKeyPair ImportSnapshot ModifyInstanceAttribute ModifySnapshotAttribute ModifyVolume RegisterImage DeregisterImage RunInstances StartInstances StopInstances TerminateInstances	
EBS	CompleteSnapshot GetSnapshotBlock ListChangedBlocks ListSnapshotBlocks PutSnapshotBlock StartSnapshot	These permissions are additionaly required for AWS Diff API changes along with previous AWS Native backup and restore permissions.

*If you want to use a KMS key belonging to a different AWS account, then perform the steps described in the AWS documentation.

CloudSpin On Demand

Open the JSON (or right-click and save as a file).

AWS Resource Type	Permissions	Description
IAM	AddRoleToInstanceProfile AttachRolePolicy CreateInstanceProfile CreateRole GetInstanceProfile GetRole GetRolePolicy GetUser PassRole PutRolePolicy TagRole	Cohesity creates a VM import policy which is assigned to a role in AWS IAM. The custom created role is assigned to the S3 bucket where the corresponding VM disks are stored which are later used to import into EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS during deployment. "iam:GetUser" and "iam:GetRole" are used to fetch the user information and validate the credentials while registering the AWS IAM user on the cluster. "iam:AddRoleToInstanceProfile" is required to register source using IAM role.
S3 Bucket	AbortMultipartUpload CreateBucket DeleteObject DeleteObjectTagging DeleteObjectVersion DeleteObjectVersionTagging GetBucketAcl GetBucketTagging GetObject GetObjectAcl GetObjectTagging GetObjectTorrent GetObjectVersion GetObjectVersionAcl GetObjectVersionTagging GetObjectVersionTorrent ListMultipartUploadParts PutBucketAcl PutBucketTagging PutBucketPublicAccessBlock PutObject PutObjectAcl PutObjectTagging PutObjectVersionAcl PutObjectVersionTagging RestoreObject	Cohesity creates a S3 bucket in AWS. The VM disk is copied and stored as an object in this S3 bucket. A policy is attached to this bucket, which is later used to import into EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS.

AWS Resource Type	Permissions	Description
EC2	AttachNetworkInterface AttachVolume CancelImportTask CreateImage CreateNetworkInterface CreateSecurityGroup CreateSnapshot CreateTags CreateVolume DeleteNetworkInterface DeleteSnapshot DeleteTags DeleteVolume DeregisterImage DescribeAvailabilityZones DescribeImages DescribeImportImageTasks DescribeImportSnapshotTasks DescribeInstanceStatus DescribeInstanceTypes DescribeInstances DescribeKeyPairs DescribeNetworkInterfaces DescribeRegions DescribeReservedInstancesOfferings DescribeSecurityGroups DescribeSnapshots DescribeSubnets DescribeVolumes DescribeVpcs DetachNetworkInterface DetachVolume ImportImage ImportSnapshot ModifyInstanceAttribute ModifyNetworkInterfaceAttribute ModifyVolumeAttribute RegisterImage	These permissions are required to register the AWS IAM user on the Cohesity cluster. Once the IAM user is registered on the cluster, describe permissions on every EC2 service is required for the user to successfully select from all the avaible options for each EC2 services from the drop down while performing any of the AWS workflows. Cohesity imports the VM disk stored in the S3 bucket into EC2 images and snapshots. These snapshots and images are used to create instance in AWS. While creating the instance, Cohesity creates network interfaces, disks, and volumes. Tags are created and added to these artifacts created in AWS. To tear down this deployment VM, delete permissions are required on all artifacts created in AWS. Once the convert and deploy task is complete for all the volumes, snapshots created for this task are cleaned up. The instance is torn down on demand.

AWS Resource Type	Permissions	Description
	RunInstances StartInstances StopInstances TerminateInstances	
KMS	CreateGrant ReEncryptFrom ReEncryptTo RetireGrant	These permissions are required to Cloudspin and deploying an instance in AWS region with encryption support. These permissions are required while importing images and snapshots in EC2 for AWS region with encryption enabled.
SSM	GetCommandInvocation ListCommandInvocations SendCommand	Cohesity creates a converter instance on AWS to run some commands on the boot disk of VM in order to boot it up as instance on AWS. SSM is used to run these commands on the converter instance.

CloudSpin Policy Based Full Backup

Open the JSON (or right-click and save as a file).

AWS Resource Type	Permissions	Description
IAM	AddRoleToInstanceProfile AttachRolePolicy CreatePolicy CreateRole GetInstanceProfile GetPolicy GetRole GetRolePolicy GetUser GetUserPolicy ListAccessKeys ListPolicies ListRolePolicies PassRole PutRolePolicy TagRole	Cohesity creates a VM import policy which is assigned to a role in AWS IAM. The custom created role is assigned to the S3 bucket where the corresponding VM disks are stored and later used to import into the EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS. "iam:GetUser", "iam:ListAccessKeys" and "iam:GetUserPolicy" are used to fetch the user information and validate the credentials while registering the AWS IAM user on the cluster. "iam:AddRoleToInstanceProfile" is required to register source using IAM role.
S3 Bucket	CreateBucket DeleteObject GetBucketAcl GetBucketTagging GetBucketPolicy GetObject ListBucket PutBucketPolicy PutBucketTagging PutBucketPublicAccessBlock PutObject	Cohesity creates a S3 bucket in AWS. The VM disk is copied and stored as an object in this S3 bucket. A policy is attached to this bucket, which is later used to import into EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS.

AWS Resource Type	Permissions	Description
EC2	AttachVolume AuthorizeSecurityGroupIngress CancelImportTask CreateFleet CreateImage CreateNetworkInterface CreateSecurityGroup CreateSnapshot CreateTags CreateVolume DeleteSnapshot DeleteTags DeleteVolume DeregisterImage DescribeAvailabilityZones DescribeFleetInstances DescribeImages DescribeImportImageTasks DescribeImportSnapshotTasks DescribeInstanceAttribute DescribeInstanceCreditSpecifications DescribeInstanceStatus DescribeInstances DescribeKeyPairs DescribeNetworkInterfaces DescribeRegions DescribeReservedInstances DescribeReservedInstancesListings DescribeReservedInstancesModifications DescribeReservedInstancesOfferings DescribeSecurityGroupReferences DescribeSecurityGroups DescribeSnapshotAttribute DescribeSnapshots DescribeSubnets DescribeTags DescribeVolumeAttribute	These permissions are required to register the AWS IAM user on the Cohesity cluster. Once the IAM user is registered on the cluster, describe permissions on every EC2 service is required for the user to successfully select from all the avaible options for each EC2 services from the drop down while performing any of the AWS workflow. Cohesity imports the vm disk stored in the S3 bucket into EC2 image and snapshots. These snapshots and Images are used to create instance in AWS. While creating the instance, we create network interfaces, disks, and volumes. Tags are created and added to these artifacts creates in AWS. To tear down this deployment VM, delete permissions are required on all artifacts created in AWS. Once the retention peroid of the cloud spin task is over, the correcponding EC2 AMI and snapshots are deleted which requires delete permissions as well.

AWS Resource Type	Permissions	Description
	DescribeVolumeStatus DescribeVolumes DescribeVpcAttribute DescribeVpcClassicLink DescribeVpcClassicLinkDnsSupport DescribeVpcEndpointConnections DescribeVpcEndpointServiceConfigurations DescribeVpcEndpointServicePermissions DescribeVpcEndpointServices DescribeVpcEndpoints DescribeVpcPeeringConnections DescribeVpcs DetachVolume ImportImage ImportInstance ImportSnapshot ImportVolume ModifyInstanceAttribute RunInstances StartInstances StopInstances TerminateInstances	
KMS	CreateGrant ReEncryptFrom ReEncryptTo RetireGrant	These permissions are required to Cloudspin and deploying an instance in AWS region with encryption support. These permissions are required while importing images and snapshots in EC2 for AWS region with encryption enabled.
SSM	GetCommandInvocation ListCommandInvocations SendCommand	Cohesity creates a converter instance on AWS to run some commands on the boot disk of VM in order to boot it up as instance on AWS. SSM is used to run these commands on the converter instance.

CloudSpin Policy Based Incremental Backup

Open the JSON (or right-click and save as a file).

AWS Resource Type	Permissions	Description
IAM	AddRoleToInstanceProfile AttachRolePolicy CreatePolicy CreateRole GetInstanceProfile GetPolicy GetRole GetRolePolicy GetUser GetUserPolicy ListAccessKeys ListPolicies ListRolePolicies PassRole PutRolePolicy TagRole	Cohesity creates a VM import policy which is assigned to a role in AWS IAM. The custom created role is assigned to the S3 bucket where the corresponding VM disks are stored, for the first full data copy run of incremental Cloud Spin, which are later used to import into EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS. "iam:GetUser" and "iam.GetRole" are used to fetch the user information and validate the credentials while registering the AWS IAM user on the cluster. "iam:AddRoleToInstanceProfile" is required to register source using IAM role.
S3 Bucket	CreateBucket DeleteObject GetBucketAcl GetBucketTagging GetBucketPolicy GetObject ListBucket PutBucketPolicy PutBucketTagging PutBucketPublicAccessBlock PutObject	Cohesity creates a S3 bucket in AWS. The VM disk is copied and stored as an object in this S3 bucket. A policy is attached to this bucket, which is later used to import into EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS.

AWS Resource Type	Permissions	Description
EC2	AttachVolume AuthorizeSecurityGroupIngress CancelImportTask CreateFleet CreateImage CreateNetworkInterface CreateSecurityGroup CreateSnapshot CreateTags CreateVolume DeleteSnapshot DeleteTags DeleteVolume DeregisterImage DescribeAvailabilityZones DescribeFleetInstances DescribeImages DescribeImportImageTasks DescribeImportSnapshotTasks DescribeInstanceAttribute DescribeInstanceCreditSpecifications DescribeInstanceStatus DescribeInstances DescribeKeyPairs DescribeNetworkInterfaces DescribeRegions DescribeReservedInstances DescribeReservedInstancesListings DescribeReservedInstancesModifications DescribeReservedInstancesOfferings DescribeSecurityGroupReferences DescribeSecurityGroups DescribeSnapshotAttribute DescribeSnapshots DescribeSubnets DescribeTags DescribeVolumeAttribute	These permissions are required to register the AWS IAM user on the Cohesity cluster. Once the IAM user is registered on the cluster, describe permissions on every EC2 service is required for the user to successfully select from all the avaible options for each EC2 services from the drop down while performing any of the AWS workflow. Cloud spin of a VM with Incremental Data copy requires the user to provide a VPC and Subnet while policy creation along with the Region, this is required because this workflow creates a temporary instance for copying the difference in data to AWS. Cohesity imports the vm disk stored in the S3 bucket into EC2 image and snapshots. These snapshots and Images are used to create instance in AWS. While creating the instance, we create network interfaces, disks, and volumes. Tags are created and added to these artifacts creates in AWS. To tear down this deployment VM, delete permissions are required on all artifacts created in AWS. Once the retention peroid of the cloud spin task is over, the correcponding EC2 AMI and snapshots are deleted which requires delete permissions as well.

AWS Resource Type	Permissions	Description
	DescribeVolumeStatus DescribeVolumes DescribeVpcAttribute DescribeVpcClassicLink DescribeVpcClassicLinkDnsSupport DescribeVpcEndpointConnections DescribeVpcEndpointServiceConfigurations DescribeVpcEndpointServicePermissions DescribeVpcEndpointServices DescribeVpcEndpoints DescribeVpcPeeringConnections DescribeVpcs DetachVolume ImportImage ImportInstance ImportSnapshot ImportVolume ModifyInstanceAttribute RunInstances StartInstances StopInstances TerminateInstances	
KMS	CreateGrant ReEncryptFrom ReEncryptTo RetireGrant	These permissions are required to Cloudspin and deploying an instance in AWS region with encryption support. These permissions are required while importing images and snapshots in EC2 for AWS region with encryption enabled.
SSM	GetCommandInvocation ListCommandInvocations SendCommand	Cohesity creates a converter instance on AWS to run some commands on the boot disk of VM in order to boot it up as instance on AWS. SSM is used to run these commands on the converter instance.

Disaster Recovery Failback and Failover

Open the JSON (or right-click and save as a file).

AWS Resource Type	Permissions	Description
IAM	AddRoleToInstanceProfile CreateRole GetInstanceProfile GetRole GetRolePolicy GetUser PutRolePolicy	Failover to AWS workflow involves Cloud Spining the replicated VM on demand to AWS. For this workflow of Cloud Spin on demand, Cohesity creates a VM import policy which is assigned to a role in AWS IAM. The custom created role is assigned to the S3 bucket where the corresponding VM disks are stored and later used to import into EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS during deployment. "iam:GetUser" and "iam:GetRole" are used to fetch the user information and validate the credentials while registering the AWS IAM user on the cluster. "iam:AddRoleToInstanceProfile" is required to register source using IAM role.

AWS Resource Type	Permissions	Description
S3	AbortMultipartUpload CreateBucket DeleteObject DeleteObjectTagging DeleteObjectVersion DeleteObjectVersionTagging GetBucketAcl GetBucketTagging GetObject GetObjectAcl GetObjectTagging GetObjectTorrent GetObjectVersion GetObjectVersionAcl GetObjectVersionTagging GetObjectVersionTorrent HeadObject ListMultipartUploadParts PutBucketAcl PutBucketTagging PutBucketPublicAccessBlock PutObject PutObjectAcl PutObjectTagging PutObjectVersionAcl PutObjectVersionTagging RestoreObject	Failover to AWS workflow involves Cloud Spining the replicated VM on demand to AWS. For this workflow of Cloud Spin on demand, Cohesity creates a S3 bucket in AWS. The VM disk is copied and stored as an object in this S3 bucket. A policy is attched to this bucket, which is later used to import into EC2 images and snapshots. These EC2 images and snapshots are used for launching the instance in AWS.

AWS Resource Type	Permissions	Description
EC2	AttachNetworkInterface AttachVolume CancelExportTask CancelImportTask CreateImage CreateInstanceExportTask CreateNetworkInterface CreateTags CreateVolume DeleteNetworkInterface DeleteSnapshot DeleteTags DeleteVolume DeregisterImage DescribeAvailabilityZones DescribeExportTasks DescribeImages DescribeImportImageTasks DescribeImportSnapshotTasks DescribeInstanceStatus DescribeInstances DescribeKeyPairs DescribeNetworkInterfaces DescribeRegions DescribeReservedInstancesOfferings DescribeSecurityGroups DescribeSnapshots DescribeSubnets DescribeVolumeAttribute DescribeVolumes DescribeVpcs DetachNetworkInterface DetachVolume ImportImage ImportSnapshot ModifyInstanceAttribute ModifyNetworkInterfaceAttribute ModifyVolumeAttribute	Failover to AWS requires Cloud Spin of VMs on demand to AWS and Failback involves taking a backup of the AWS VM onto Cohesity, therefore permissions are required for both workflows. EC2 permissions are required to register the AWS IAM user on the Cohesity cluster. Once the IAM user is registered on the cluster, describe permissions on every EC2 service is required for the user to successfully select from all the available options for each EC2 services from the drop down while performing any of the AWS workflows. Cohesity imports the VM disk stored in the S3 bucket into the EC2 image and snapshots. These snapshots and images are used to create the instance in AWS. While creating the instance, Cohesity creates network interfaces, disks, and volumes. Tags are created and added to these artifacts created in AWS. Once the convert and deploy task is complete for all the volumes, snapshots created for this task are cleaned up. Cohesity creates proxy EC2 instances for doing back up of AWS EC2 instance. Cohesity creates snaphots of the EC2 volumes while backing up and stores the different instance attributes and tags.

AWS Resource Type	Permissions	Description
	RegisterImage RunInstances StartInstances StopInstances TerminateInstances	
SSM	GetCommandInvocation ListCommandInvocations SendCommand	Failover to AWS workflow involves Cloud Spining the replicated VM on demand to AWS. For this workflow of Cloud Spin on demand Cohesity creates a converter instance in AWS. There are some commands that need to be run on the boot disk of VM in order to boot it up as instance on AWS. For these commands to run, SSM is used.

Permissions for EBS Volume Protection

Permission Category	Permissions
EC2	ec2:DescribeSnapshots ec2:CreateSnapshot ec2:DeleteSnapshot ec2:CopySnapshot ec2:DescribeVolumes ec2:DescribeInstances ec2:CreateTags ec2:DeleteTags ec2:DescribeTags

Permission Category	Permissions
EBS	ebs:ListSnapshotBlocks ebs:StartSnapshot ebs:PutSnapshotBlock ebs:GetSnapshotBlock ebs:ListChangedBlocks ebs:CompleteSnapshot
KMS	kms:DescribeKey kms:Decrypt kms:Encrypt kms:GenerateDataKey kms:CreateGrant

Permissions for Amazon S3 Data Protection

Resource	Permissions	Reason
S3	s3:GetBucketLocation s3:GetBucketNotification s3:GetBucketOwnershipControls s3:GetBucketTagging s3:GetBucketVersioning s3:GetInventoryConfiguration s3:GetObject s3:GetObjectAcl s3:GetObjectTagging s3:GetObjectVersion s3:GetObjectVersionAcl s3:GetObjectVersionTagging s3:ListAllMyBuckets s3:ListBucket s3:PutBucketNotification s3:PutInventoryConfiguration s3:PutObject s3:PutObjectAcl s3:PutObjectTagging s3:PutObjectVersionAcl s3:PutObjectVersionTagging	These permissions are required for the backup and recovery of S3 objects.
iam	iam:SimulatePrincipalPolicy	SimulatePrincipalPolicy is needed to ensure that the required actions are allowed on the IAM role we created as part of the Cloud Formation template.
kms*	kms:CreateGrant kms:DescribeKey kms:ListAliases kms:GenerateDataKey	KMS permissions are needed to read data of an encrypted database at the time of backup, as well as write encrypted data to the recovered database. Describe permissions are needed so we can list & identify keys associated with database instances.

Resource	Permissions	Reason
Events	events:DeleteRule events:PutTargets events:RemoveTargets	These permissions are required for capturing the incremental changes on the S3 buckets.
Glue	glue:CreateJob glue:DeleteJob glue:GetJobRun glue:StartJobRun glue:TagResource glue:UpdateJob	These permissions are required for sorting the inventory report. The sorted inventory report is then used to back up the S3 objects to the Cohesity Platform.
SQS	sqs:CreateQueue sqs:DeleteMessage sqs:DeleteQueue sqs:GetQueueUrl sqs:GetQueueAttributes sqs:PurgeQueue sqs:ReceiveMessage sqs:SetQueueAttributes sqs:ListQueues sqs:SetQueueAttributes sqs:TagQueue	These permissions are required for capturing the incremental changes on the S3 buckets.

Permissions for Amazon RDS for MySQL Protection

Permissions for AWS Snapshot-based Protection

Resource	Permissions	Purpose

rds	`rds:AddTagsToResource`
	`rds:CopyDBClusterSnapshot`
	`rds:CopyDBSnapshot`
	`rds:CreateDBClusterSnapshot`
	`rds:CreateDBInstance`
	`rds:CreateDBSnapshot`
	`rds:DeleteDBClusterSnapshot`
	`rds:DeleteDBSnapshot`
	`rds:DescribeDBClusterSnapshots`
	`rds:DescribeDBClusters`
	`rds:DescribeDBInstances`
	`rds:DescribeDBParameterGroups`
	`rds:DescribeDBSnapshots`
	`rds:DescribeDBSubnetGroups`
	`rds:DescribeOptionGroups`
	`rds:DescribeSourceRegions`
	`rds:ModifyDBClusterSnapshotAttribute`
	`rds:ModifyDBSnapshotAttribute`
	`rds:RestoreDBClusterFromSnapshot`
	`rds:RestoreDBInstanceFromDBSnapshot`
	`rds:RestoreDBInstanceToPointInTime`
		These permissions are required to register the AWS account on Cohesity with the IAM role which got created by the Cloud Formation template.
		Once the source is registered with Cohesity, described permissions are needed so Cohesity can identify resources present in the account, which will be used for backups as well as at the time of recovery we use this information to provide a list of options to choose from.
		Cohesity creates snapshots of the Amazon RDS instances while backing up and storing the different database instance attributes and tags.
		While recovering the database instance, Cohesity creates DB instance/cluster, it also attaches the original tags.
		Cohesity requires the delete snapshots permissions to delete the expired/old snapshots it creates during the backup.
		The modify snapshot attributes permissions are required for sharing the snapshot across accounts if cross-account recovery is attempted.
		`RestoreDBInstanceToPointInTime` and `RestoreDBClusterToPointInTime` is needed to do the point in time recoveries.

Permissions for Ingest-based Protection

Resource	Permission	Purpose
rds	`rds:DescribeDBInstances`	List RDS instances.

Permissions for Amazon Aurora MySQL Protection

Permissions for AWS Snapshot-based Protection

Resource	Permissions	Purpose

rds	`rds:AddTagsToResource`
	`rds:CopyDBClusterSnapshot`
	`rds:CopyDBSnapshot`
	`rds:CreateDBClusterSnapshot`
	`rds:CreateDBInstance`
	`rds:CreateDBSnapshot`
	`rds:DeleteDBClusterSnapshot`
	`rds:DeleteDBSnapshot`
	`rds:DescribeDBClusterSnapshots`
	`rds:DescribeDBClusters`
	`rds:DescribeDBInstances`
	`rds:DescribeDBParameterGroups`
	`rds:DescribeDBSnapshots`
	`rds:DescribeDBSubnetGroups`
	`rds:DescribeOptionGroups`
	`rds:DescribeSourceRegions`
	`rds:ModifyDBClusterSnapshotAttribute`
	`rds:ModifyDBSnapshotAttribute`
	`rds:RestoreDBClusterFromSnapshot`
	`rds:RestoreDBInstanceFromDBSnapshot`
	`rds:RestoreDBInstanceToPointInTime`
		These permissions are required to register the AWS account on Cohesity with the IAM role which got created by the Cloud Formation template.
		Once the source is registered with Cohesity, described permissions are needed so Cohesity can identify resources present in the account, which will be used for backups as well as at the time of recovery we use this information to provide a list of options to choose from.
		Cohesity creates snapshots of the Amazon Aurora instances while backing up and storing the different database instance attributes and tags.
		While recovering the database instance, Cohesity creates DB instance/cluster, it also attaches the original tags.
		Cohesity requires the delete snapshots permissions to delete the expired/old snapshots it creates during the backup.
		The modify snapshot attributes permissions are required for sharing the snapshot across accounts if cross-account recovery is attempted.
		`RestoreDBInstanceToPointInTime` and `RestoreDBClusterToPointInTime` is needed to do the point in time recoveries.

Permissions for Ingest-based Protection

The following table lists the minimum required permissions included in the CloudFormation template for protecting Amazon Aurora MySQL at the database level (Ingest-based protection).

> **Note:** You do not need to add these permissions manually, as they are automatically added when you run the CFT.

Resource	Permission	Purpose
rds	`rds:DescribeDBClusters`	List the Aurora clusters.

Permissions for Amazon Aurora PostgreSQL Protection

Permissions for AWS Snapshot-based Protection

Resource	Permissions	Purpose

rds	`rds:AddTagsToResource`
	`rds:CopyDBClusterSnapshot`
	`rds:CreateDBClusterSnapshot`
	`rds:DescribeDBInstances`
	`rds:CreateDBInstance`
	`rds:DeleteDBClusterSnapshot`
	`rds:DescribeDBClusterSnapshots`
	`rds:DescribeDBClusters`
	`rds:DescribeDBParameterGroups`
	`rds:DescribeDBSubnetGroups`
	`rds:DescribeSourceRegions`
	`rds:ModifyDBClusterSnapshotAttribute`
	`rds:RestoreDBClusterFromSnapshot`
		These permissions are required to register the AWS account on Cohesity with the IAM role which got created by the Cloud Formation template.
		Once the source is registered with Cohesity, described permissions are needed so Cohesity can identify resources present in the account, which will be used for backups as well as at the time of recovery we use this information to provide a list of options to choose from.
		Cohesity creates snapshots of the Amazon Aurora instances while backing up and storing the different database instance attributes and tags.
		While recovering the database instance, Cohesity creates DB instance/cluster, it also attaches the original tags.
		Cohesity requires the delete snapshots permissions to delete the expired/old snapshots it creates during the backup.
		The modify snapshot attributes permissions are required for sharing the snapshot across accounts if cross-account recovery is attempted.

Permissions for Ingest-based Protection

Resource	Permission	Purpose
rds	`rds:DescribeDBClusters`	List RDS instances.

Permissions for Amazon RDS for PostgreSQL Protection

Permissions for AWS Snapshot-based Protection

Resource	Permissions	Purpose

rds	`rds:AddTagsToResource`
	`rds:CopyDBClusterSnapshot`
	`rds:CopyDBSnapshot`
	`rds:CreateDBClusterSnapshot`
	`rds:CreateDBInstance`
	`rds:CreateDBSnapshot`
	`rds:DeleteDBClusterSnapshot`
	`rds:DeleteDBSnapshot`
	`rds:DescribeDBClusterSnapshots`
	`rds:DescribeDBClusters`
	`rds:DescribeDBInstances`
	`rds:DescribeDBParameterGroups`
	`rds:DescribeDBSnapshots`
	`rds:DescribeDBSubnetGroups`
	`rds:DescribeOptionGroups`
	`rds:DescribeSourceRegions`
	`rds:ModifyDBClusterSnapshotAttribute`
	`rds:ModifyDBSnapshotAttribute`
	`rds:RestoreDBClusterFromSnapshot`
	`rds:RestoreDBInstanceFromDBSnapshot`
	`rds:RestoreDBInstanceToPointInTime`
		These permissions are required to register the AWS account on Cohesity with the IAM role which got created by the Cloud Formation template.
		Once the source is registered with Cohesity, described permissions are needed so Cohesity can identify resources present in the account, which will be used for backups as well as at the time of recovery we use this information to provide a list of options to choose from.
		Cohesity creates snapshots of the Amazon RDS instances while backing up and storing the different database instance attributes and tags.
		While recovering the database instance, Cohesity creates DB instance/cluster, it also attaches the original tags.
		Cohesity requires the delete snapshots permissions to delete the expired/old snapshots it creates during the backup.
		The modify snapshot attributes permissions are required for sharing the snapshot across accounts if cross-account recovery is attempted.
		`RestoreDBInstanceToPointInTime` and `RestoreDBClusterToPointInTime` is needed to do the point in time recoveries.

Permissions for Ingest-based Protection

Resource	Permission	Purpose
rds	`rds:DescribeDBInstances`	List RDS instances.

Permissions for Amazon RDS for Oracle Protection

Permissions for AWS Snapshot-based Protection

Resource	Permissions	Purpose
rds	`rds:AddTagsToResource`	To add tags to RDS resources.

Resource	Permissions	Purpose
	`rds:CopyDBClusterSnapshot`	To copy an RDS DB cluster snapshot to another region or account.
	`rds:CopyDBSnapshot`	To copy an RDS DB snapshot.
	`rds:CreateDBClusterSnapshot`	To create a manual DB cluster snapshot. Used when protecting Aurora-based or cluster-based Oracle deployments (if applicable).
	`rds:CreateDBSnapshot`	To create a manual DB snapshot. This is the core permission for Oracle snapshot backups.
	`rds:DeleteDBClusterSnapshot`	To delete a DB cluster snapshot.
	`rds:DeleteDBSnapshot`	To delete a DB snapshot.
	`rds:ModifyDBClusterSnapshotAttribute`	To update attributes of a DB cluster snapshot.
	`rds:ModifyDBSnapshotAttribute`	To modify attributes of a DB snapshot.
	`rds:CreateDBInstance`	To create a new RDS DB instance. Required when restoring a DB instance from a snapshot.
	`rds:RestoreDBClusterFromSnapshot`	To restore an RDS DB cluster from a snapshot and create a new cluster.
	`rds:RestoreDBInstanceFromDBSnapshot`	To restore a standalone RDS DB instance from a snapshot.
	`rds:DescribeDBClusterSnapshots`	To list all DB cluster snapshots for discovery and inventory.
	`rds:DescribeDBClusters`	To retrieve information about DB clusters (engine type, settings, etc.).
	`rds:DescribeDBInstances`	To list DB instances and identify Oracle databases configured for backup.
	`rds:DescribeDBParameterGroups`	To read parameter groups associated with RDS instances for configuration awareness.

Resource	Permissions	Purpose
	`rds:DescribeDBSnapshots`	To list DB snapshots for visibility and validation of backups.
	`rds:DescribeDBSubnetGroups`	To retrieve DB subnet groups to ensure compatibility when restoring instances.
	`rds:DescribeOptionGroups`	To list option groups and validate Oracle-specific configuration during restore.
	`rds:DescribeSourceRegions`	To retrieve supported source regions for cross-region snapshot copy operations.

Permissions for Ingest-based Protection

Resource	Permissions	Purpose
rds	`rds:DescribeDBClusters`	To list RDS DB clusters for discovery.
	`rds:DescribeDBInstances`	To list RDS DB instances for discovery.
	`rds:DescribeTenantDatabases`	To discover tenant databases.
	`rds:DescribeDBProxies`	To list all RDS proxies.
	`rds:DescribeDBProxyTargets`	To list RDS DB instances associated with a particular proxy.
	`rds:DescribeDBProxyEndpoints`	To list proxy endpoints for a given RDS proxy.
	`rds:ListTagsForResource`	To list tags associated with RDS resources.
	`rds:DescribeOptionGroupOptions`	To list Oracle option group options and verify if export to EFS or S3 is enabled.
	`rds:ModifyOptionGroup`	To modify an option group to add S3 export capability.

Resource	Permissions	Purpose
ec2	`ec2:DescribeInstanceTypes`	To retrieve available instance types. Used to determine parallelism required for backup and restore operations.
	`ec2:DescribeNetworkInterfaces`	To list network interfaces required for EFS connectivity.
	`ec2:DescribeSubnets`	To list subnets required for EFS connectivity.
	`ec2:CreateNetworkInterface`	To create network interfaces needed for EFS access.
s3	`s3:CreateBucket`	To create an S3 bucket required for backups.
	`s3:ListBucket`	To check whether an S3 bucket already exists.
	`s3:GetObjectAcl`	To retrieve ACLs of S3 objects created during backup.
	`s3:GetObject`	To retrieve contents of S3 objects used in backup or restore workflows.
	`s3:DeleteObject`	To delete objects from the S3 bucket during backup operations.
	`s3:PutObject`	To upload data into the S3 bucket during backup workflows.
	`s3:TagResource`	Allows tagging of the S3 staging bucket.

Resource	Permissions	Purpose
elasticfilesystem	`elasticfilesystem:DescribeFileSystems`	To list EFS file systems.
	`elasticfilesystem:CreateFileSystem`	To create an EFS file system.
	`elasticfilesystem:DeleteFileSystem`	To delete an EFS file system.
	`elasticfilesystem:DescribeMountTargets`	To list mount targets for the EFS file system.
	`elasticfilesystem:CreateMountTarget`	To create mount targets for EFS.
	`elasticfilesystem:DeleteMountTarget`	To delete EFS mount targets.
	`elasticfilesystem:ClientMount`	To mount the EFS file system on the client.
	`elasticfilesystem:ClientWrite`	To write data to EFS.
	`elasticfilesystem:ClientRootAccess`	To enable root-level access over NFS for EFS operations.
	`AmazonRDSReadOnlyAccess`	To read data placed in RDS directories so it can be exported to EFS or S3.

Permissions for Amazon RDS for MS SQL Server Protection

Resource	Permissions	Purpose
rds	`rds:DescribeDBClusters`	Lists RDS database clusters for discovery
	`rds:DescribeDBInstances`	Lists RDS database instances for discovery.
ec2	`ec2:DescribeInstanceTypes`	Retrieves details of instance types available in a specific location. Used to determine parallelism during RDS backup and restore operations.

Resource	Permissions	Purpose
S3	`s3:CreateBucket`	Creates an S3 bucket required during RDS restore operations.
	`s3:ListBucket`	Checks whether an S3 bucket already exists
	`s3:GetObjectAcl`	Retrieves ACLs for an S3 object stored in a bucket.
	`s3:GetObject`	Retrieves the contents of an S3 object stored in a bucket.
	`s3:PutObject`	Uploads data to an S3 bucket.
	`s3:DeleteObject`	Deletes an object from an S3 bucket.
	`s3:TagResource`	Allows tagging of the S3 staging bucket

Permissions Required for DocumentDB Protection

Resource	Permissions	Purpose
iam	`iam:GetUser`	Required to validate the IAM identity used for account registration.
rds	`rds:DescribeDBClusters`	Required to discover Amazon DocumentDB instance-based clusters.
ec2	`ec2:DescribeRegions`	Required for region discovery.
docdb-elastic	`docdb-elastic:ListClusters`	Required to discover Amazon DocumentDB elastic clusters.

Permissions for DynamoDB Protection

Resource	Permissions	Purpose
IAM	`iam:PassRole`	Allow passing an IAM role to AWS services (e.g., EC2, Lambda) for resource access.

Resource	Permissions	Purpose

kms	`kms:CreateGrant`
	`kms:DescribeKey`
	`kms:Decrypt`
	`kms:Encrypt`
	`kms:ListAliases`
	`kms:ListKeys`
		Enable reading encrypted data during backup and writing encrypted data during recovery. Also allows listing and identifying keys associated with database instances.

dynamodb	`dynamodb:BatchWriteItem`
	`dynamodb:DeleteItem`
	`dynamodb:GetItem`
	`dynamodb:PutItem`
	`dynamodb:Query`
	`dynamodb:RestoreTableToPointInTime`
	`dynamodb:Scan dynamodb:UpdateItem`
	`dynamodb:CreateTable`
	`dynamodb:UpdateTable`
	`dynamodb:DescribeContinuousBackups`
	`dynamodb:DescribeExport`
	`dynamodb:DescribeImport`
	`dynamodb:DescribeTable`
	`dynamodb:ExportTableToPointInTime`
	`dynamodb:ImportTable`
	`dynamodb:ListTables`
	`dynamodb:ListTagsOfResource`
	`dynamodb:TagResource`
	`dynamodb:UpdateContinuousBackups`
		Required for backing up and recovering Amazon DynamoDB tables, including reading, writing, restoring, and managing table metadata and tags.

Resource	Permissions	Purpose

s3	`s3:CreateBucket`
	`s3:GetBucketLocation`
	`s3:PutBucketTagging`
	`s3:DeleteBucket`
	`s3:AbortMultipartUpload`
	`s3:DeleteObject`
	`s3:GetBucketVersioning`
	`s3:GetObject`
	`s3:GetObjectTagging`
	`s3:GetObjectVersion`
	`s3:ListBucket`
	`s3:ListBucketMultipartUploads`
	`s3:ListMultipartUploadParts`
	`s3:PutObject`
	`s3:PutObjectTagging`
	`s3:TagResource`
		Manage S3 buckets and objects for storing, tagging, and retrieving backup/exported data during DynamoDB protection operations.

glue	`glue:CreateJob`
	`glue:TagResource`
	`glue:DeleteJob`
	`glue:StartJobRun`
	`glue:GetJobRun`
		Manage AWS Glue jobs for data transformation and movement as part of backup/recovery workflows.

CloudWatch	`logs:DescribeLogGroups`
	`logs:CreateLogGroup`
	`logs:CreateLogStream`
	`logs:DescribeLogStreams`
	`logs:GetLogEvents`
	`logs:PutLogEvents`
		Create and manage log groups and streams, retrieve and publish log events for auditing and monitoring backup/recovery operations.

Permissions for Amazon Redshift Cluster Protection

Resource	Permissions	Purpose
IAM	`iam:GetRole`	Access IAM roles
redshift	`redshift:ListDatabases`	To list databases of a Redshift cluster. Retrieve information about database names and their metadata. This permission is for cluster level.
	`redshift:GetClusterCredentialsWithIAM`	To connect to Redshift cluster database using IAM.
	`redshift:DescribeClusters`	For getting properties of Redshift clusters.
	`redshift:GetClusterCredentials`	To connect to Redshift cluster database.
redshift-data	`redshift-data:ExecuteStatement`	To run a query in a Redshift cluster database.
	`redshift-data:ListDatabases`	To list databases of a Redshift cluster via redshift-data API which is a different endpoint than redshift API endpoint. This permission is required for redshift without a server.
	`redshift-data:GetStatementResult`	To fetch temporarily cached result of an SQL statement executed on Redshift cluster databases.
	`redshift-data:CancelStatement`	For canceling a query executed on Redshift cluster database used during Protection Group cancellation.
	`redshift-data:DescribeStatement`	Required to get the details about a specific instance when a query is run by the Amazon Redshift Data API.

Resource	Permissions	Purpose
S3	`s3:CreateBucket`	To create an S3 bucket for storing backup or exported data.
	`s3:ListBucket`	To list objects within an S3 bucket.
	`s3:GetBucketLocation`	To get location of the S3 bucket.
	`s3:GetObjectAcl`	To retrieve the ACL of S3 objects used during backup or export workflows.
	`s3:GetObject`	To retrieve objects from an S3 bucket (e.g., exported data).
	`s3:PutObject`	To upload objects to an S3 bucket.
	`s3:DeleteObject`	To delete objects from the S3 bucket to free up storage and manage lifecycle costs.
	`s3:TagResource`	To apply tags to the S3 staging bucket.

Permissions for Amazon Redshift Serverless Protection

Resource	Permissions	Purpose
IAM	`iam:GetRole`	Access IAM roles

Resource	Permissions	Purpose
redshift-serverless	`redshift-serverless:ListNamespaces`	List all Redshift Serverless namespaces in the account.
	`redshift-serverless:ListWorkgroups`	List all Redshift Serverless workgroups in the account.
	`redshift-serverless:GetNamespace`	Get details of a specific Redshift Serverless namespace.
	`redshift-serverless:GetWorkgroup`	Get details of a specific Redshift Serverless workgroup.
	`redshift-serverless:ListTagsForResource`	List tags associated with Redshift Serverless resources.
	`redshift-serverless:GetCredentials`	Get temporary credentials to connect to a Redshift Serverless workgroup.
redshift-data	`redshift-data:ExecuteStatement`	Run a query in a Redshift Serverless namespace database.
	`redshift-data:ListDatabases`	List databases via redshift-data API.
	`redshift-data:GetStatementResult`	Fetch cached result of an SQL statement .
	`redshift-data:CancelStatement`	Cancel a query during Protection Group cancellation.
	`redshift-data:DescribeStatement`	Get details of a query run by the Redshift Data API
secretsmanager	`secretsmanager:GetSecretValue`	Retrieve secret values for Redshift Serverless authentication.

Resource	Permissions	Purpose
S3	`s3:CreateBucket`	To create an S3 bucket for storing backup or exported data.
	`s3:ListBucket`	To list objects within an S3 bucket.
	`s3:GetBucketLocation`	To get location of the S3 bucket.
	`s3:GetObjectAcl`	To retrieve the ACL of S3 objects used during backup or export workflows.
	`s3:GetObject`	To retrieve objects from an S3 bucket (e.g., exported data).
	`s3:PutObject`	To upload objects to an S3 bucket.
	`s3:DeleteObject`	To delete objects from the S3 bucket to free up storage and manage lifecycle costs.
	`s3:TagResource`	To apply tags to the S3 staging bucket.

Comprehensive Permissions List for AWS Source Protection

Permission Required for Delegated Administrator IAM Role

The following table lists the permission required by the delegated administrator IAM role.

For Org-level and account-level CFT-based (IAM role) registration, these permissions are available in the CFT for discovering the member accounts and OUs of your organization. You do not need to manually add these permissions.

For Org-level preconfigured role-based registration, you must add these permissions to the IAM policy that will be created for attaching to the delegated administrator IAM role.

AWS Resource	Permissions
organizations	organizations:DescribeOrganization organizations:ListRoots organizations:ListAccounts organizations:DescribeAccount organizations:ListOrganizationalUnitsForParent organizations:ListChildren organizations:DescribeOrganizationalUnit organizations:ListParents
sts	sts:AssumeRole

Sample

```
{
   "Version": "2012-10-17",
   "Statement": [
       {
           "Effect": "Allow",
           "Action": [
               "organizations:DescribeOrganization",
               "organizations:ListRoots",
               "organizations:ListAccounts",
               "organizations:DescribeAccount",
               "organizations:ListOrganizationalUnitsForParent",
               "organizations:ListChildren",
               "organizations:DescribeOrganizationalUnit",
               "organizations:ListParents",
               "sts:AssumeRole"
           ],
           "Resource": "*"
       }
   ]
}
```

Comprehensive Permissions list for AWS Source Protection

The following table lists the permissions required for protecting AWS source.

For Org-level and account-level CFT-based (IAM role) registration, these permissions are available in the CFT for the backup and recovery of AWS sources. You do not need to manually add these permissions.

For account-level IAM-user based registration, you must add these permissions to the IAM policy that will be created for attaching to the IAM user.

For Org-level preconfigured role-based registration, you must add these permissions to the IAM policy that will be created for attaching to the **Member account IAM role**

AWS Resource Type	Permissions

EBS	`ebs:GetSnapshotBlock`
	`ebs:ListChangedBlocks`
	`ebs:ListSnapshotBlocks`

AWS Resource Type	Permissions
EC2	ec2:AllocateAddress ec2:AssociateSubnetCidrBlock ec2:AssociateVpcCidrBlock ec2:AttachInternetGateway ec2:AttachNetworkInterface ec2:AttachVolume ec2:AttachVpnGateway ec2:AuthorizeSecurityGroupEgress ec2:AuthorizeSecurityGroupIngress ec2:CancelImportTask ec2:CopySnapshot ec2:CreateFleet ec2:CreateImage ec2:CreateInstanceExportTask ec2:CreateKeyPair ec2:CreateNetworkInterface ec2:CreateSecurityGroup ec2:CreateSnapshot ec2:CreateTags ec2:CreateVolume ec2:DeleteNetworkInterface ec2:DeleteSnapshot ec2:DeleteTags ec2:DeleteVolume ec2:DeregisterImage ec2:DescribeAvailabilityZones ec2:DescribeExportTasks ec2:DescribeFleetHistory ec2:DescribeFleetInstances ec2:DescribeFleets ec2:DescribeImages ec2:DescribeImportImageTasks ec2:DescribeImportSnapshotTasks ec2:DescribeInstanceAttribute ec2:DescribeInstanceCreditSpecifications ec2:DescribeInstances ec2:DescribeInstanceStatus ec2:DescribeKeyPairs

AWS Resource Type	Permissions
	ec2:DescribeNetworkInterfaces ec2:DescribeRegions ec2:DescribeReservedInstances ec2:DescribeReservedInstancesListings ec2:DescribeReservedInstancesModifications ec2:DescribeReservedInstancesOfferings ec2:DescribeSecurityGroupReferences ec2:DescribeSecurityGroups ec2:DescribeSnapshotAttribute ec2:DescribeSnapshots ec2:DescribeSubnets ec2:DescribeTags ec2:DescribeVolumeAttribute ec2:DescribeVolumes ec2:DescribeVolumeStatus ec2:DescribeVpcAttribute ec2:DescribeVpcClassicLink ec2:DescribeVpcClassicLinkDnsSupport ec2:DescribeVpcEndpointConnectionNotifications ec2:DescribeVpcEndpointConnections ec2:DescribeVpcEndpoints ec2:DescribeVpcEndpointServiceConfigurations ec2:DescribeVpcEndpointServicePermissions ec2:DescribeVpcEndpointServices ec2:DescribeVpcPeeringConnections ec2:DescribeVpcs ec2:DetachNetworkInterface ec2:DetachVolume ec2:ImportImage ec2:ImportInstance ec2:ImportKeyPair ec2:ImportSnapshot ec2:ImportVolume ec2:ModifyFleet ec2:ModifyInstanceAttribute ec2:ModifyNetworkInterfaceAttribute ec2:ModifySnapshotAttribute ec2:ModifyVolume ec2:ModifyVolumeAttribute ec2:RegisterImage

AWS Resource Type	Permissions
IAM	iam:AddRoleToInstanceProfile iam:AttachRolePolicy iam:CreatePolicy iam:CreateRole iam:DeleteRole iam:DeleteRolePolicy iam:GetAccessKeyLastUsed iam:GetInstanceProfile iam:GetPolicy iam:GetRole iam:GetRolePolicy iam:GetUser iam:GetUserPolicy iam:ListAccessKeys iam:ListPolicies iam:ListRolePolicies iam:ListRoles iam:ListRoleTags iam:PassRole iam:PutRolePolicy iam:TagRole iam:UntagRole iam:UpdateRole iam:UpdateUser
KMS	kms:CreateGrant kms:Decrypt kms:DescribeKey kms:Encrypt kms:GenerateDataKey kms:GenerateDataKeyWithoutPlaintext kms:ListAliases kms:ListKeys kms:ReEncryptFrom kms:ReEncryptTo

AWS Resource Type	Permissions
RDS	rds:AddTagsToResource rds:CopyDBSnapshot rds:CreateDBSnapshot rds:DeleteDBSnapshot rds:DescribeDBClusters rds:DescribeDBInstances rds:DescribeDBParameterGroups rds:DescribeDBSnapshots rds:DescribeDBSubnetGroups rds:DescribeOptionGroups rds:DescribeSourceRegions rds:ModifyDBSnapshotAttribute rds:RestoreDBInstanceFromDBSnapshot rds:RestoreDBInstanceToPointInTime rds:CopyDBClusterSnapshot rds:CreateDBClusterSnapshot rds:CreateDBInstance rds:DeleteDBClusterSnapshot rds:DescribeDBClusterSnapshots rds:DescribeTenantDatabases rds:ModifyDBClusterSnapshotAttribute rds:RestoreDBClusterFromSnapshot

AWS Resource Type	Permissions
S3	s3:AbortMultipartUpload s3:CreateBucket s3:DeleteObject s3:DeleteObjectTagging s3:DeleteObjectVersion s3:DeleteObjectVersionTagging s3:GetBucketAcl s3:GetBucketLocation s3:GetBucketNotification s3:GetBucketObjectLockConfiguration s3:GetBucketOwnershipControls s3:GetBucketPolicy s3:GetBucketTagging s3:GetBucketVersioning s3:GetInventoryConfiguration s3:GetLifecycleConfiguration s3:GetObject s3:GetObjectAcl s3:GetObjectAttributes s3:GetObjectTagging s3:GetObjectTorrent s3:GetObjectVersion s3:GetObjectVersionAcl s3:GetObjectVersionAttributes s3:GetObjectVersionTagging s3:GetObjectVersionTorrent s3:HeadObject s3:ListAllMyBuckets s3:ListBucket s3:ListBucketMultipartUploads s3:ListMultipartUploadParts s3:PutBucketAcl s3:PutBucketNotification s3:PutBucketPolicy s3:PutBucketPublicAccessBlock s3:PutBucketTagging s3:PutInventoryConfiguration s3:PutObject s3:PutObjectAcl s3:PutObjectRetention

AWS Resource Type	Permissions
	s3:PutObjectTagging s3:PutObjectVersionAcl s3:PutObjectVersionTagging s3:RestoreObject s3:TagResource
SSM	ssm:GetCommandInvocation ssm:SendCommand ssm:UpdateInstanceInformation ssmmessages:*
DynamoDB	dynamodb:BatchWriteItem dynamodb:CreateTable dynamodb:DeleteItem dynamodb:DescribeContinuousBackups dynamodb:DescribeExport dynamodb:DescribeImport dynamodb:DescribeTable dynamodb:ExportTableToPointInTime dynamodb:GetItem dynamodb:ImportTable dynamodb:ListTables dynamodb:ListTagsOfResource dynamodb:PutItem dynamodb:Query dynamodb:RestoreTableToPointInTime dynamodb:Scan dynamodb:TagResource dynamodb:UpdateContinuousBackups dynamodb:UpdateItem dynamodb:UpdateTable

AWS Resource Type	Permissions
Glue	glue:CreateJob glue:DeleteJob glue:GetJobRun glue:StartJobRun glue:TagResource
CloudWatch	logs:CreateLogGroup logs:CreateLogStream logs:DescribeLogStreams logs:GetLogEvents logs:PutLogEvents logs:PutRetentionPolicy
DocDB Elastic	docdb-elastic:ListClusters
Redshift	redshift:DescribeClusters redshift:GetClusterCredentials redshift:GetClusterCredentialsWithIAM redshift:ListDatabases
Redshift Data API	redshift-data:CancelStatement redshift-data:DescribeStatement redshift-data:ExecuteStatement redshift-data:GetStatementResult redshift-data:ListDatabases
secretsmanager	secretsmanager:GetSecretValue

Assign AWS Permissions

AWS permissions are required to use Cohesity cloud services such as CloudArchive and CloudSpin. Cohesity recommends that you create an AWS policy with permissions and then attach it to IAM users. For example, you can create a policy that grants permissions for back up and recovery and attach it to users.

The permissions are available as JSON in Minimum Permissions for Individual Workloads and Workflows. You can copy the JSON and paste it into the policy, as described below.

Create a Policy with Permissions

1. Log in to the Amazon AWS console using your AWS account with admin privileges.
2. In the top left of the menu bar, select **Services**.
3. Under **Security, Identity & Compliance**, select **IAM**.
4. In the left frame, click **Policies**.
5. Click **Create policy**.
6. In the **Create policy** page, select the **JSON** tab.
7. Create the JSON code with the required permissions and paste it in the **JSON** tab.
8. Click **Next: Tags**. Adding tags is optional.

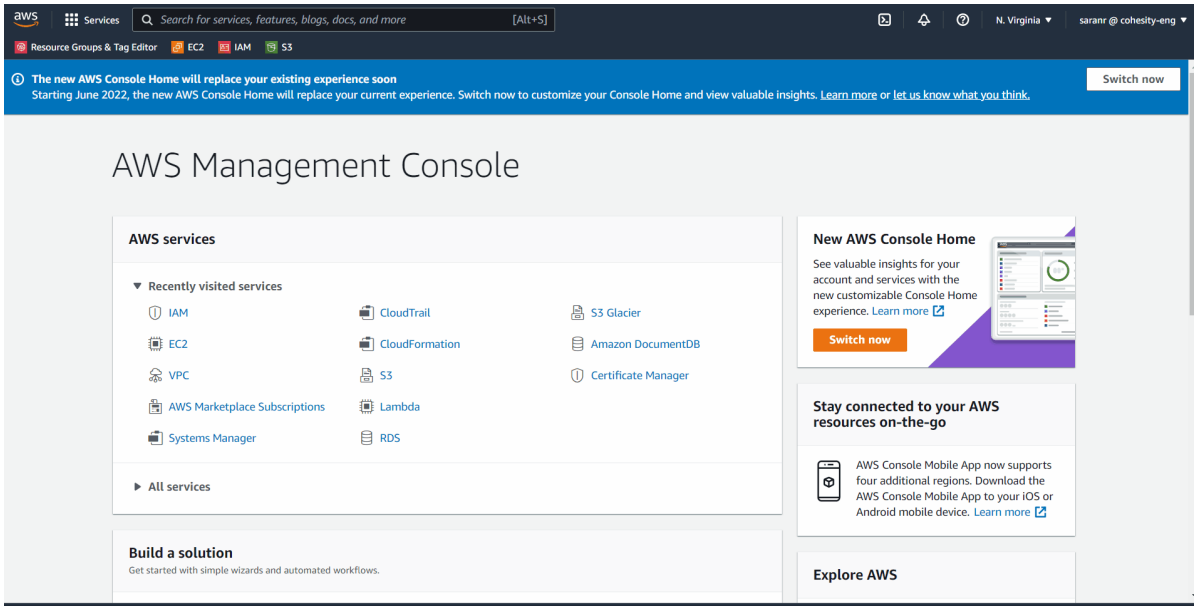

9. Click **Next: Review**.
10. Enter a name for the policy and an optional description and click **Create Policy**.

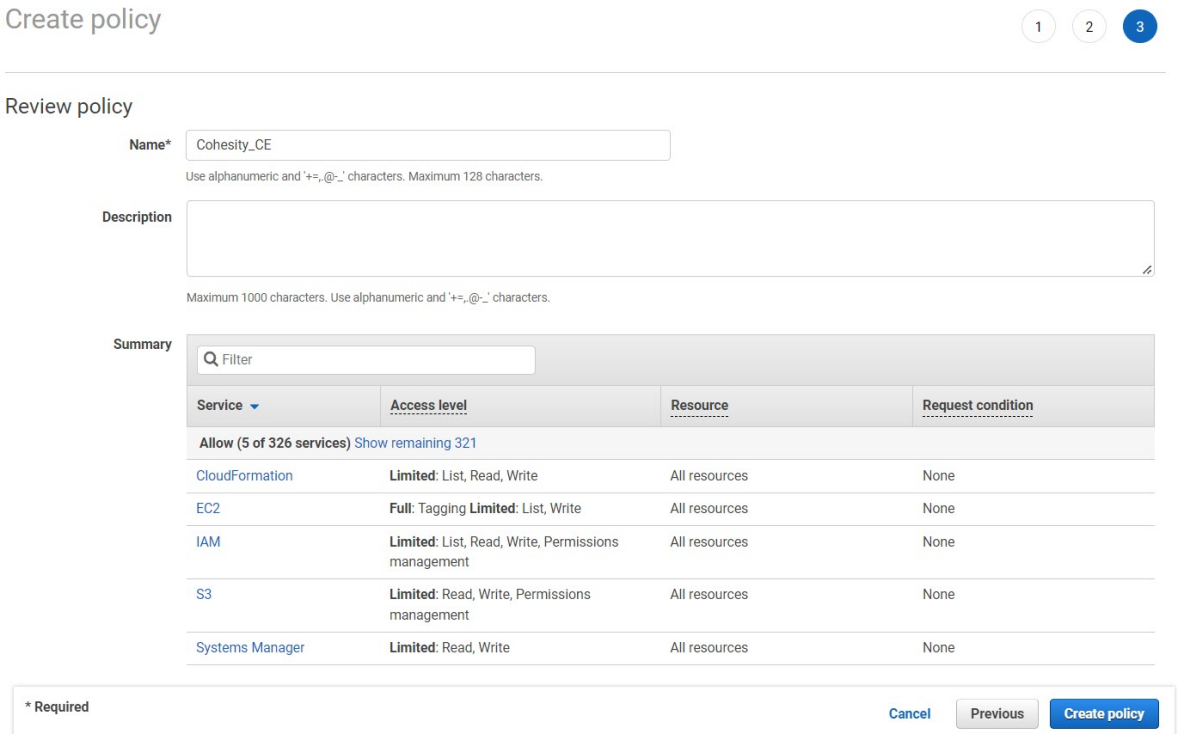

Now you can attach the policy to an IAM user.

Create an IAM User

1. Log in to the Amazon AWS console using an account with administrative privileges.

1. Navigate to **IAM**.
2. In the left navigation pane, click **Users**.

   Click **Create user**.
3. Click **Add users**.
4. In the **User name** field, enter a user name such as `CohesityCE`.
5. For **Access type**, leave the default settings, and thn click **Next**.

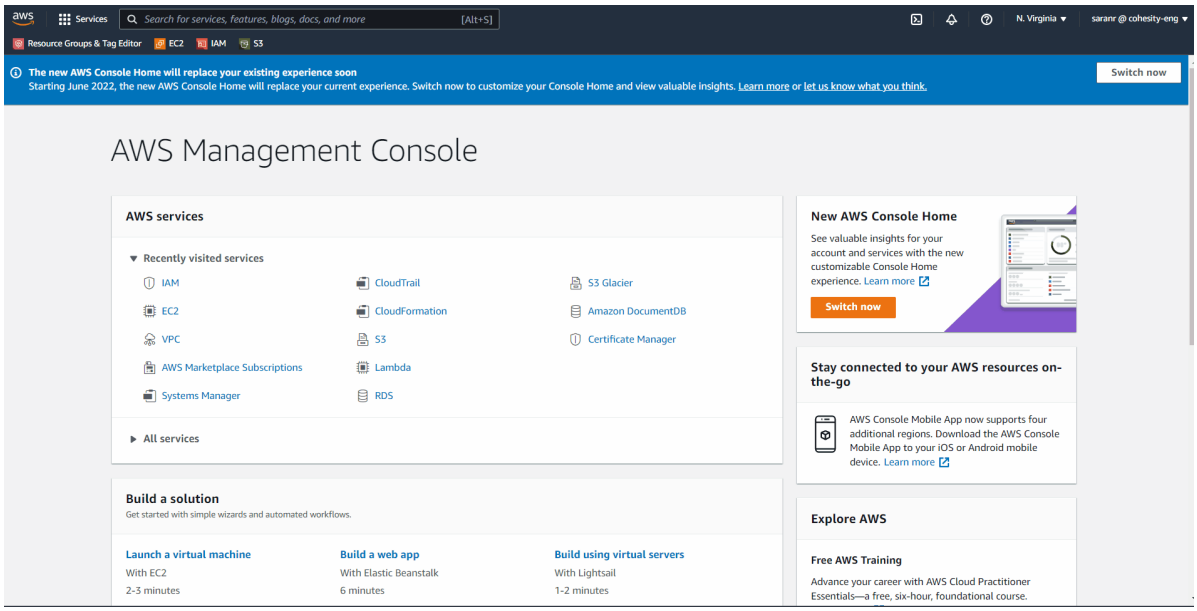

6. Click **Attach existing policies directly**, then search for the name of the policy you created previously and select that policy.
7. Click **Next: Tags**.
8. Optionally, enter tags for the policy and then click **Next: Review**.
9. Review your choices and then click **Create User**.
10. In the **Secret access key** column, click **Show**.

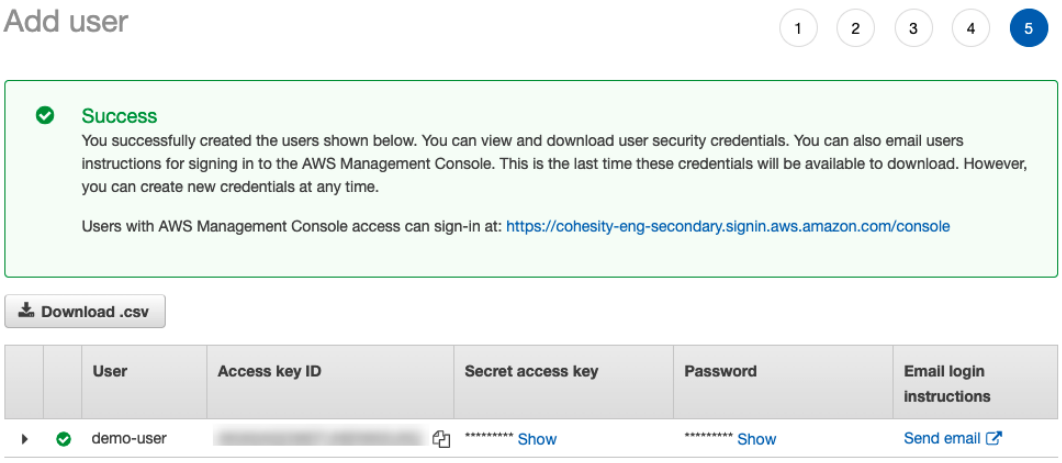

11. Copy and paste the **Access key ID** and the **Secret access key** in to a text file. You'll need them later when you register the AWS Cloud source.
12. Click **Close**.

AWS Permission Updates for Cluster Upgrade

If you have upgraded your cluster from an earlier version and already have an AWS account registered, you must add required permissions to protect any newly supported AWS workloads.

AWS Account Registered Using IAM User Method

Update the IAM policy associated with the IAM user and add the required permissions to protect the newly supported AWS workloads:

Steps to update IAM Policy Permissions

1. From the AWS Management Console, select **Users**.
2. Open the **Permissions** tab and locate the attached policy.
3. Create new policy or update existing policy with required permissions. For more information, see Minimum Permissions for Individual Workloads and Workflows.

AWS Account Registered Using IAM Role Method

Update the existing *CohesitySourceRegistrationRole* by adding the required permissions. For more information, see Minimum Permissions for Individual Workloads and Workflows.

Steps to Update Role Permissions

1. Navigate to **IAM** > **Roles** in the AWS Management Console.
2. Search for and open *CohesitySourceRegistrationRole*.
3. Open the **Permissions** tab and locate the inline policy *source-regis-access*.
4. Click **Edit**.
5. Add the permissions to the first statement's Action array.
6. Click **Next**, then save changes.

1. Assign a policy name (example, *CohesityDynamoDBPolicy*).
2. Click **Create policy**.

Register AWS Source

To protect your AWS workloads, you must register AWS as a source on the Cohesity cluster. Cohesity supports registering AWS environments at both the Organization level and the account level, providing flexibility in how you onboard and manage AWS resources for protection.

You can register an AWS Organization and its member accounts as separate AWS sources in the same Cohesity cluster.

Organization-Level Registration

With organization-level registration, you can register an AWS Organization to enable centralized onboarding and management of multiple AWS accounts.

> **Note:** Organization-level registration is only supported on Cohesity Cloud Edition on AWS.

Account-Level Registration

With account-level registration, you can register and manage individual AWS accounts independently.

AWS Organization-Level Registration

Use organization-level registration to onboard and manage multiple AWS accounts through AWS Organizations. This method enables centralized management of AWS environments from a single registration workflow.

During registration, you can optionally specify Organizational Unit (OU) IDs to limit the discovery of AWS accounts within those OUs. This allows you to control the scope of resource discovery and protection.

Organization-level registration is useful when:

- You want centralized onboarding of multiple AWS accounts
- Your environment is structured using AWS Organizations

AWS Organization Registration Methods

To protect AWS source in your organization with Cohesity, you must register your AWS Organization as a source in the Cohesity cluster and perform backup and recovery operations.

You can register AWS Organization using one of the following options:

- Register AWS Organization using Preconfigured IAM roles
- Register AWS Organization Using CloudFormation Template

Comparison At a Glance

Registration Method	Description
Using preconfigured IAM roles	Registers an AWS Organization using the delegated admin IAM role and IAM member role that are already configured with the required permissions to discover the AWS resources.
Using Cohesity-provided CloudFormation template	Registers an AWS Organization using a Cohesity-provided CloudFormation template to automatically create required IAM roles and policies to discover the AWS resources.

Register AWS Organization using Preconfigured IAM roles

You can register an AWS Organization using the delegated admin IAM role and IAM member role that are already configured with the required permissions to discover the AWS resources.

Pre-requisites

Before registering AWS at the Organization level, ensure that the following prerequisites are met:

- The Cohesity cluster is deployed in one of the member accounts of your AWS Organization. We recommend you deploy the Cohesity cluster in the delegated administrator account of your AWS Organization.
- AWS Organizations is enabled in your environment.
- **Delegated administrator account ID**:

  A delegated administrator is an AWS member account authorized by the management account to perform AWS Backup administrative tasks across all accounts in your organization.

  To configure a delegated administrator account, run the following commands from your management account:

  1. Enable trusted access for AWS Backup:

     ```
     aws organizations enable-aws-service-access --service-principal backup.amazonaws.com
     ```

  2. Register the member account as delegated administrator:

     ```
     aws organizations register-delegated-administrator --account-id <MEMBER_ACCOUNT_ID> --service-principal backup.amazonaws.com
     ```

     Where, <MEMBER_ACCOUNT_ID> is the AWS account ID of the member account you want to designate as delegated administrator .

- **Delegated administrator IAM role**:

  The IAM role in the delegated admin account that has the required permissions to discover the AWS Organizational Units (OUs) and member accounts within the organization. For more information on permissions, see Comprehensive Permissions List for AWS Source Protection.

- **Member account IAM role**:

  An IAM role that is present in all member accounts within the AWS Organization and is configured with the permissions required to discover and protect the AWS services you want to protect. Cohesity assumes this role to perform resource-level discovery and protection for AWS workloads within each member account. For more information on permissions, see Comprehensive Permissions List for AWS Source Protection.

> **Note:** Cohesity will not discover member accounts that do not have this role configured.

Register AWS Organization

To register AWS at Org-level using preconfigured IAM roles:

1. Navigate to **Data Protection** > **Sources**.
2. Click Register at the top-right of the page and then select **Cloud Platforms > AWS**.

   The Register AWS Source dialog box appears.
3. In **Source Type**, select **AWS Organization**.
4. In **Delegated Admin Account ID**, enter the AWS account ID of the delegated administrator.
5. Turn on **Specify Pre-configured IAM Roles**.
6. In **Delegated Admin Role Name**, enter the IAM role in the delegated admin account. Ensure that IAM role has the required permissions for discovering the OUs and member accounts.
7. In **Member Account Role Name**, enter the IAM role used in all member accounts. Ensure the IAM role has the permissions required to discover the AWS workloads you want to protect.
8. (Optional) To restrict discovery to AWS accounts within specific OUs, turn on **Specify Organizational Unit ID**, and then enter the OU IDs.

   > **Note:** If you specify OU IDs, Cohesity discovers only the accounts that belong to those OUs. Accounts that are directly attached to the organization root (management account) are not discovered. To discover all accounts in the organization, including those attached directly to the organization root, skip this step.
9. From the list of **AWS Services**, review the services selected for protection. All services are selected by default. Clear any services that you do not want to protect:
   - **EC2**
   - **S3**
   - **DynamoDB**
   - **RDS**
   - **Redshift Serverless**

- **Redshift Cluster**
- **DocumentDB**

10. (*Applies only for Amazon S3 buckets protection*) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Under **Inventory Report Location**, provide the following information to create the inventory report:
    a. **AWS Account ID**: Enter the AWS member account ID in your organization where you want to create the inventory report.
    b. **S3 buckets ARN**: Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same cloud region as the Amazon S3 you want to protect.
    c. **Prefix**: Add a prefix value to the name of the inventory report that will be created.
        - The prefix can contain any character, including white spaces.
        - For example, Report-Source Bucket.
        - The prefix should not begin or end with a forward slash (/) .
        - The prefix should not contain consecutive forward slashes.
        - You must not upload any files in the prefix of the inventory report.
        - The inventory report will be created at **<Prefix>/<Path_to_ inventory_report>** when you initiate the protection. Cohesity creates the path to the inventory report configuration on AWS based on the Cohesity cluster, Amazon S3 bucket, and protection.
11. Click **Register**.

The AWS Organization is registered as a source on the Cohesity cluster, and all member accounts within the organization or the specified OUs are discovered. For more information, see What Happens After Registering AWS Organization.

Attach Trust Relationship

The delegated administrator IAM role and the member account IAM role contain the permissions to allow Cohesity to discover AWS workloads in the entity hierarchy within the product UI. However, to enable full discovery, backup, and recovery operations, you must establish the following trust relationships:

- **Between the delegated administrator IAM role and the IAM role attached to Cohesity cluster node (*CohesityInstanceRole*)**: This trust relationship allows the *CohesityInstanceRole* to assume the delegated administrator role, enabling discovery of resources across the AWS organization.
- **Between the member account IAM role and the IAM role attached to Cohesity cluster node (*CohesityInstanceRole*)**: This trust relationship allows the

CohesityInstanceRole to assume the member account IAM role in each member account, enabling backup, restore, and discovery of account-level resources.

These trust relationships use AWS Security Token Service (STS) to allow the *CohesityInstanceRole* to assume the appropriate roles as needed for each operation.

Configure the Trust Relationship

Update the trust policy of both the delegated administrator IAM role and member account IAM role to allow the *CohesityInstanceRole* to assume them:

1. Sign in to the AWS Management Console.
2. Open AWS IAM Console.
3. Navigate to **Roles**.
4. Select the IAM role.
5. Open the **Trust relationships** tab.
6. Choose **Edit trust policy**.
7. Add the *CohesityInstanceRole* ARN to the trust policy.
8. Save the changes.

Example for adding CohesityInstanceRole ARN to the delegated admin role trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<cluster_account_id>:root"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:PrincipalArn": "arn:aws:iam::<cluster_account_id>:role/CohesityInstanceRole"
        }
      }
    }
  ]
}
```

Example for adding CohesityInstanceRole ARN to the member account role trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<cluster_account_id>:role/CohesityInstanceRole",
        "Service": [
          "glue.amazonaws.com",
          "redshift.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Replace <cluster_account_id> with the AWS account ID where Cohesity cluster is deployed.

Register AWS Organization Using CloudFormation Template

You can register an AWS Organization using a Cohesity-provided CloudFormation template to automatically create required IAM roles and policies to discover the AWS resources.

Pre-requisites

Before registering AWS at the Organization level using CFT, ensure that the following prerequisites are met:

- AWS Organizations is enabled in your environment.
- The Cohesity cluster is deployed in one of the member accounts of your AWS Organization. We recommend you deploy the Cohesity cluster in the delegated administrator account of your AWS Organization.
- Choose one of the following AWS CloudFormation StackSet deployment methods:
  - **Deploy from the delegated administrator account**: The user or the IAM role performing the StackSet deployment must have the required permissions to deploy AWS CloudFormation StackSet.
  - **Deploy from the organization management account**: Use the organization management account to deploy the StackSet. No additional delegated-administrator permissions are required. The CloudFormation template automatically identifies the delegated administrator account.

  For information on the permissions required, see the Permissions Required for Deploying the CloudFormation StackSet from the Delegated Administrator Account topic below.
- If you are deploying the AWS CloudFormation StackSet from the delegated administrator account, you must:

- Register that account as a delegated administrator for CloudFormation StackSets in AWS Organizations. Perform the following steps from the management account:
  a. Verify that the delegated administrator account is registered:

     ```
     aws organizations list-delegated-administrators --service-principal member.org.stacksets.cloudformation.amazonaws.com
     ```

  b. If the delegated administrator account is not listed in the output, register it:

     ```
     aws organizations register-delegated-administrator --account-id <delegated_admin_account_id> --service-principal member.org.stacksets.cloudformation.amazonaws.com
     ```

- Enable trusted access between AWS CloudFormation StackSets and AWS Organizations from the delegated administrator account:
  a. Log in to the AWS Management Console of the delegated administrator account.
  b. Navigate to A**WS Organizations** > **Services** > **CloudFormation StackSets** > **Enable trusted access**.

Permissions Required for Deploying the CloudFormation StackSet from the Delegated Administrator Account

If you deploy the CloudFormation StackSet from the delegated administrator account, ensure that the user or role performing the deployment has the permissions required to create and manage AWS CloudFormation StackSets.

The required CloudFormation permissions include:

- `cloudformation:CreateStackSet`
- `cloudformation:UpdateStackSet`
- `cloudformation:DeleteStackSet`
- `cloudformation:DescribeStackSet`
- `cloudformation:DescribeStackSetOperation`
- `cloudformation:ListStackSets`
- `cloudformation:ListStackInstance`

Register AWS Organization

1. Navigate to **Data Protection** > **Sources**.
2. Click Register at the top-right of the page and then select **Cloud Platforms** > **AWS**.

   The Register AWS Source dialog box appears.

3. In **Source Type**, select **AWS Organization**.
4. In **Delegated Admin Account ID**, enter the AWS account ID of the delegated administrator.
5. Ensure **Specify Pre-configured IAM Roles** is turned off.
6. (Optional) To limit discovery to specific OUs, turn on **Specify Organizational Unit ID**, and then enter the OU IDs.
7. Under AWS Services, select the services that you want to protect.
    - **EC2**
    - **S3**
    - **DynamoDB**
    - **RDS**
    - **Redshift Serverless**
    - **Redshift Cluster**
    - **DocumentDB**
8. (*Applies only for Amazon S3 buckets protection*) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Under Inventory Report Location, provide the following information to create the inventory report:
    a. **AWS Account ID**: Enter the AWS member account ID in your organization where you want to create the inventory report.
    b. **S3 buckets ARN**: Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same cloud region as the Amazon S3 you want to protect.
    c. **Prefix**: Add a prefix value to the name of the inventory report that will be created.
        - The prefix can contain any character, including white spaces.
        - For example, Report-Source Bucket.
        - The prefix should not begin or end with a forward slash (/) .
        - The prefix should not contain consecutive forward slashes.
        - You must not upload any files in the prefix of the inventory report.

    The inventory report will be created at <**Prefix>/<Path_to_inventory_report>** when you initiate the protection. Cohesity creates the path to the inventory report on AWS based on the cluster, Amazon S3 bucket, and protection.
9. Click **Next** to generate a CloudFormation Template, which you will use to complete the AWS source registration.

10. Click **Download** to download the CloudFormation Template.
11. Execute the downloaded CFT Stacksets on the delegated Admin account. On running the CFT, IAM roles, and policies are created. For more information, see Troubleshooting.
12. Once the roles and policies are created successfully, click **Register**.

The AWS Organization is registered as a source on the Cohesity cluster, and all member accounts within the organization or the specified OUs are discovered. You can view the registered source by navigating to **Data Protection > Sources** in the cluster UI.

The trust relationship that enables full discovery, backup, and recovery operations is configured by default when the AWS Organization is registered using the CFT template. No additional steps are required to attach the trust relationship.

Run CFT Template

To complete the registration of AWS Organization with the Cohesity cluster, you must run the Cohesity's CloudFormation Template (CFT) either from the delegated administrator account you use for registering or from the management account. Running the CFT will create the IAM roles required for discovering the AWS OUs and member accounts. For more information on registering AWS Organization, see Register AWS Organization Using CloudFormation Template.

Before you begin, ensure you have met the Pre-requisites.

To run the CFT:

1. Sign in to the AWS Management Console of the Delegated administrator AWS account or management account.
2. Open the AWS CloudFormation console.
3. In the CloudFormation console navigation pane, choose **StackSets**.
4. Choose **Create StackSet**.
5. On the Choose a template page, configure the following options, and then click **Next**:
    1. Under **Permissions**, select **Service-managed permissions**. By selecting this option, AWS CloudFormation automatically manages deployment roles across the AWS Organization. The delegated administrator account must have the permissions required to manage StackSets and access AWS Organizations.
    2. Under **Prerequisite – Prepare template**, select **Template is ready**.
    3. Under **Specify template**, choose **Upload a template file** as the **Template source**.
    4. Choose **Choose file** and select the CloudFormation Template you downloaded from the Cohesity cluster UI.

6. On the **Specify stack set details** page, configure the following options, and then click **Next**:
    1. **Stack set name**: enter a name (for example, test-stackset). The name must contain only letters, numbers, and hyphens, and must start with a letter.
    2. **Stack set description – optional**: optionally describe the StackSet's purpose.
    3. Under **Parameters**, fill in the values:
        1. **ClusterAccountId**: the AWS account ID, where the Cohesity cluster is deployed.
        2. **DelegatedAdminAccountId**: The AWS account ID where the delegated administrator role exist for organization discovery and StackSet operations.
        3. **Workflow**: Select one of the following based in your requirement:
            - **source_registration**: Select this option are if you are creating IAM role for registration AWS source.
            - **external_target**: Select this option if you are creating IAM role for cloud tiering or cloud archiving purposes.
            - **both**: Select this option if you are rcreating IAM role for both of the above purposes.
7. On the **Configure stack set options** page, configure the following options, and then click **Next**:
    1. **Tags – optional**: Choose Add new tag to apply key-value metadata to the resources (up to 50 tags per stack). Leave empty if no tags are needed.
    2. **Execution configuration → Managed execution**: Choose how StackSets handles operations.
        1. **Inactive**: StackSets performs one operation at a time.
        2. **Active**: StackSets runs non-conflicting operations concurrently and queues conflicting ones, starting queued operations in the requested order once conflicts clear.
    3. **Capabilities**: Select the checkbox **I acknowledge that AWS CloudFormation might create IAM resources with customized names**. This acknowledgment is required to continue.
8. On the **Set deployment options** page, configure the following options, and then click **Next**:
    1. **Add stacks to stack set**: Choose **Deploy new stacks**.
    2. Under **Accounts > Deployment locations**, choose the targeting method:

1. **Deploy stacks in accounts**: Enter the target accounts directly.

   In **Account numbers**, type the 12-digit account numbers separated by commas, or choose Upload .csv file to populate them from a file.

   Deploy stacks in organisational units: target all accounts in one or more OUs.

3. Under **Specify Regions**, select each Region to deploy to from the dropdown. Stacks are deployed in the order listed; use the up/down arrows to reorder, **Remove** to drop one, or **Add all regions / Remove all regions** to bulk-edit.
4. Under **Deployment options**, set the operation preferences:
   1. **Maximum concurrent accounts – optional**: Number of accounts per Region to deploy to simultaneously. A higher number speeds up the operation. (Default 1.)
   2. **Failure tolerance – optional**: Number of accounts per Region where a deployment can fail before the operation stops. If stopped in one Region, the operation will not continue in other Regions; a lower number is safer. (Default 0.)
   3. **Region concurrency**:
      1. **Sequential**: Deploy into one Region at a time, following the specified Region order.
      2. **Parallel**: Deploy simultaneously across all specified Regions.

      > **Note:** Maximum concurrent accounts must not exceed failure tolerance + 1.

   4. **Concurrency mode**: Specifies how the concurrency level behaves during operation execution.
9. On the Review page:
   1. Review all selections from the previous steps.
   2. Re-confirm the IAM capabilities acknowledgment if prompted.
   3. Choose **Submit** to create the StackSet.

CloudFormation begins creating the StackSet and the stack instances in the targeted accounts and Regions. The StackSet details page opens, where you can monitor the progress and status of each stack instance.

The CFT provisions the following IAM roles and profiles:

- *CohesityInstanceRole*: Attached to Cohesity cluster EC2 nodes; used as the base identity for assuming other roles.

- *CohesityInstanceProfile*: EC2 instance profile that wraps *CohesityInstanceRole*; must be attached to all cluster nodes.
- *CohesitySourceRegistrationRole*: Grants permissions for source registration operations.
- *CohesityDelegatedAdminRole*: Created in the delegated administrator account; allows Cohesity to discover OUs and member accounts across the organization.
- *CohesityBackupS3StagingRole*: Grants permissions for S3 staging operations during backup and recovery.

> **Note:** The trust relationship to enable full discovery, backup, and recovery operations are configured by default when the organization is registered using the CFT template. No additional steps are required.

Troubleshooting

If the StackSet creation fails with an `AWSCloudFormationStackSetAdministrationRole` error when deploying from the delegated administrator account, this indicates one of the following:

- The delegated administrator account is not registered for CloudFormation StackSets.
- **Service-managed permissions** was not selected as the permission model while running te CFT.

What Happens After Registering AWS Organization

Once the AWS Organization is registered as a source on the Cohesity cluster, Cohesity initiates discovery of the AWS Organization hierarchy. Member account and resource discovery is performed asynchronously in the background and may take time depending on the size and complexity of the organization.

As member accounts are discovered, resources in those accounts become available for protection based on the workloads selected during source registration. Any new AWS accounts added to the registered organization or selected OUs are automatically discovered.

You can view the registered source by navigating to **Data Protection > Sources** in the cluster UI.

> **Note:** If discovery fails for a member account due to permission issues, connectivity problems, or other errors, Cohesity skips that account and continues discovery for the remaining accounts. You can review the failure details in the UI and take corrective action as needed.

Manage AWS Organization Registration

After registering your AWS Organization as a source, you can:

- Update the AWS source configuration.
- Refresh the source details.
- Unregister the AWS source.

Update AWS Organization Source Configuration

If you have registered AWS Organization, then you can edit the registered AWS Organization to add or remove the Amazon services protected by Cohesity from your AWS Organization.

> **Note:** Before removing an Amazon service from the source configuration, ensure that you delete all Protection Groups associated with that service.

To edit the AWS Organization source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AWS Organization source and select **Edit**.
3. In the **Edit AWS Source** page, select or unselect the Amazon services.
4. Click **Update**.

Refresh AWS Organization Source

Cohesity automatically refreshes the AWS Organization hierarchy every 24 hours to discover changes such as new or updated accounts. You can also manually refresh the AWS Organization source to fetch the latest changes.

To refresh the AWS source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AWS Organization source and select **Refresh**.

Unregister AWS Organization Source

If you plan to stop taking the backup of your AWS source, you can unregister the AWS source from the Cohesity Cluster.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the AWS source.

To unregister the AWS source:

1. Navigate to **Data Protection** > **Sources.**
2. Click the actions menu (⋮) next to the AWS Organization source and select **Unregister**.

AWS Organization Level Registration Considerations

Consider the following limitations and behaviors when registering AWS at the Organization level:

- Cohesity performs automatic refresh of AWS Organization every 24 hours.
- You must perform a manual refresh to discover the latest account and resource changes.
- The Tags option on the entity hierarchy of the cluster UI displays resources and instances across all those accounts which are discovered as a part of the source registration.
- When the **Native Backup** option is enabled for EC2 instance protection, backups use AWS EBS Direct APIs.
- Fleet instances are not supported for AWS Organization-level registration.
- The CloudSpin feature is not supported for AWS Organization-level registration.
- Backup and restore operations are supported within the same account and across different accounts belonging to the same AWS Organization.
- Cohesity does not support protection of DocumentDB Elastic Clusters when AWS is registered at the organization level.
- Cohesity uses AWS APIs to perform discovery and protection. Ensure that the AWS endpoints for the workloads you want to protect are reachable from all Cohesity cluster nodes.

AWS Account-Level Registration

Use account-level registration to onboard and manage an individual AWS account for protection. This method allows you to register a single AWS account by providing the required credentials and permissions.

Account-level registration is useful when:

- You want granular control over a specific AWS account
- AWS accounts are managed independently
- You do not use AWS Organizations

After registration, Cohesity discovers resources within the account and enables you to configure protection workflows.

> **Note:** Cohesity uses AWS APIs to perform discovery and protection. Ensure that the AWS endpoints for the workloads you want to protect are reachable from all Cohesity cluster nodes.

AWS Account Registration Methods

To protect Amazon sources with Cohesity, you must register your AWS account as a source in the Cohesity cluster and perform backup and recovery operations.

You can register AWS account using one of the following methods:

- Register AWS Account using IAM Role
- Register AWS Account Using IAM User

Cohesity recommends registering the AWS account using the **IAM Role** method for streamlined setup and secure access.

Comparison At a Glance

Registration Method	Description	UI Navigation Path	Cluster Support
Using IAM User	Registers an AWS account using IAM user credentials. You create a policy, create an IAM user with required permissions, and provide access keys during registration.	**Data Protection > Sources > Register > Virtual Machines**.	• On-premises cluster • Virtual Edition cluster • Cohesity Cloud Edition (legacy) on AWS • Cohesity Cloud Edition on AWS

Registration Method	Description	UI Navigation Path	Cluster Support
Using IAM Role	Registers an AWS account using a Cohesity-provided CloudFormation template to automatically create required IAM roles and policies.	• **Data Protection > Sources > Register > Virtual Machines.** • **Data Protection > Sources Register > Cloud Platforms > AWS.**	• Cohesity Cloud Edition (legacy) on AWS • Cohesity Cloud Edition on AWS

> **Note:** Cohesity does not support the IAM user registration method for protecting AWS S3 buckets. To protect AWS S3 buckets, you must register the AWS account using the IAM role method.

Register AWS Account using IAM Role

If you are choosing to register the AWS account using the IAM role method, then you will execute a CloudFormation template (CFT) provided by Cohesity to automate the creation of necessary AWS resources, such as IAM roles and policies, which are required for Cohesity to access and protect your AWS services. Executing the CFT ensures all required permissions and roles are set up correctly, enabling secure backup and recovery operations from Cohesity.

> **Note:** You can register an AWS account using the IAM role only on Cohesity Cloud Edition on AWS.

Prerequisites

IAM User Permissions to Execute CFT

The IAM user you use to run the CloudFormation Template and to create and view the stack must have the required permissions. The following table lists the permissions the IAM user must have to run the CloudFormation template.

> **Note:** Ensure to add these permissions manually.

Permission	Purpose
`cloudformation:CreateStack`	Create a new stack.
`cloudformation:UpdateStack`	Modify the stack during updates or re-registration
`cloudformation:DeleteStack`	Delete a stack.
`cloudformation:ValidateTemplate`	Validate the CFT before deployment.
`cloudformation:DescribeStacks`	View the status of stacks.
`cloudformation:ListStacks`	View the status of stacks.
`cloudformation:GetTemplate`	Fetch template for an existing stack.
`cloudformation:GetTemplateSummary`	View a summary of template resources and parameters.
`cloudformation:DescribeStackEvents`	View stack creation progress/errors.
`cloudformation:DescribeStackResources`	View resources created in the stack.
`cloudformation:CreateUploadBucket`	Create a temporary bucket used for uploading artifacts needed by the stack.
`iam:CreateRole`	Create required IAM roles.
`iam:DeleteRole`	Delete IAM roles when cleaning up registrations.
`iam:GetRole`	Retrieve details of existing roles.
`iam:UpdateRole`	Modify existing IAM role configuration.
`iam:TagRole`	Add tags to IAM roles created by the template.
`iam:UntagRole`	Remove tags from IAM roles.
`iam:PutRolePolicy`	Attach inline policies to IAM roles.
`iam:DeleteRolePolicy`	Delete inline policies from IAM roles.

Permission	Purpose
`iam:GetRolePolicy`	Retrieve inline policies attached to IAM roles.
`iam:CreateInstanceProfile`	Create instance profiles needed for EC2-based operations.
`iam:GetInstanceProfile`	Retrieve instance profile details.
`iam:AddRoleToInstanceProfile`	Attach IAM roles to the instance profile.
`iam:RemoveRoleFromInstanceProfile`	Detach IAM roles from the instance profile.
`iam:DeleteInstanceProfile`	Delete instance profiles during cleanup.
`iam:AttachRolePolicy`	Attach managed policies to IAM roles.
`iam:DetachRolePolicy`	Remove managed policies from roles.
`iam:ListPolicies`	List available policies for attachment/validation.
`iam:ListRoles`	Retrieve all roles to validate role existence.
`iam:PassRole`	Allow CloudFormation to pass the IAM role to AWS services.
`s3:CreateBucket`	Create an S3 bucket for storing backup or exported data.
`s3:PutObject`	Upload objects to an S3 bucket.
`s3:GetObject`	Retrieve objects from an S3 bucket.
`s3:ListBucket`	View bucket contents to validate template uploads.

Register AWS Account

Before registering an AWS account as a source, make sure you have completed all prerequisite tasks.

Option1 (Recommended)

1. Navigate to **Data Protection > Sources**.
2. Click **Register** at the top-right of the page and then select **Cloud Platforms** > **AWS**.
   The Register AWS Source dialog box appears.
3. **In Source Type, select AWS Account**.

4. Enter the **AWS Account ID**.
5. From the list of AWS Services, select the AWS services you want to protect:
   - EC2
   - EBS
   - S3
   - DynamoDB
   - RDS
   - Redshift Serverless
   - Redshift Cluster
   - DocumentDB
6. (*Applies only for Amazon EC2 instance protection*) **Specify Fleet Settings**: Enable this option to:
   - Specify the VPC and subnet in which the EC2 fleet instance should be launched.
   - Add tags to the EC2 fleet instance.

   By default, this option is disabled and the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

   Specify the VPC and Subnet

   From the **Fleet Instance location** drop-down list select one of the following options to specify the VPC and subnet in which the fleet instance should be launched.

   - **Same Subnet as protected VMs**: The fleet instance will be launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.
   - **Select New**: The fleet instance will be launched in the specified VPC and subnet of the specified region. If the AWS source is not registered with the Cohesity cluster, the cluster cannot detect the AWS region, and the associated VPC and subnets. To configure a new VPC and subnet for the fleet instance, you must first register the source and then update the registered source to specify the VPC and subnet. For more information, see Register AWS Account using IAM Role.

   If you have not selected the **Fleet Instance location** after enabling the **Specify Fleet Settings** option, then the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

   Add a Tag

   Tags enable you to categorize and identify your fleet instances. A tag consists of a key and an optional value, both of which you define.

   To add a tag to the fleet instances, click **Add Tags** and enter a tag key and value. You can add up to 47 tags to a fleet instance.

Delete a Tag

To delete an existing tag, click the **cross** icon next to the tag.

7. (*Applies only for Amazon S3 buckets protection*) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Enable **Specify S3 Inventory Report Location**, provide the following information to create the inventory report:
   - **S3 buckets ARN**: Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same AWS account and cloud region as the Amazon S3 you want to protect.
   - **Prefix**: Add a prefix value to the name of the inventory report that will be created.
     - The prefix can contain any character, including white spaces.

       For example, *Report-Source Bucket*.
     - The prefix should not begin or end with a forward slash (/) .
     - The prefix should not contain consecutive forward slashes.
     - You must not upload any files in the prefix of the inventory report.

     The inventory report will be created at *<Prefix>/<Path_to_inventory_report>* when you initiate the protection. Cohesity creates the path to the inventory report on AWS based on the cluster, Amazon S3 bucket, and protection.
8. Click **Next** to generate a CloudFormation Template, which you will use to complete the AWS source registration.
9. Click **Download** to download the CloudFormation Template.
10. Run the CloudFormation Template in your AWS account to create the IAM roles and policies that Cohesity requires for authentication. For more information, see Execute the CloudFormation Template.
11. Once the roles and policies are created successfully, click **Register**.

Option 2

To register AWS account as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Select **Register** > **Virtual Machines**.
3. From the **Select Hypervisor Source Type** drop-down, select **AWS : IAM /Role**.
4. Choose **Category** as **Standard or Gov**.
5. Choose the **Authentication Method** as an **IAM Role**.
6. Click the **CloudFormation Template** link to download the CFT file.

7. Run the CloudFormation Template in your AWS account to create the IAM roles and policies that Cohesity requires for authentication. For more information, see Execute the CloudFormation Template.
8. Specify your **AWS Account ID**.
9. Specify the **IAM Role** name.
10. (*Applies only for Amazon EC2 instance protection*) **Specify Fleet Settings**: Enable this option to:
    - Specify the VPC and subnet in which the EC2 fleet instance should be launched.
    - Add tags to the EC2 fleet instance.

    By default, this option is disabled and the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

    Specify the VPC and Subnet

    From the **Fleet Instance location** drop-down list select one of the following options to specify the VPC and subnet in which the fleet instance should be launched.

    - **Same Subnet as protected VMs**: The fleet instance will be launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.
    - **Select New**: The fleet instance will be launched in the specified VPC and subnet of the specified region. If the AWS source is not registered with the Cohesity cluster, the cluster cannot detect the AWS region, and the associated VPC and subnets. To configure a new VPC and subnet for the fleet instance, you must first register the source and then update the registered source to specify the VPC and subnet. For more information, see Register AWS Account using IAM Role.

    If you have not selected the **Fleet Instance location** after enabling the **Specify Fleet Settings** option, then the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

    Add a Tag

    Tags enable you to categorize and identify your fleet instances. A tag consists of a key and an optional value, both of which you define.

    To add a tag to the fleet instances, click **Add Tags** and enter a tag key and value. You can add up to 47 tags to a fleet instance.

    Delete a Tag

    To delete an existing tag, click the **cross** icon next to the tag.
11. (*Applies only for Amazon S3 buckets protection*) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Enable **Specify S3 Inventory Report Location**, provide the following information to create the inventory report:

- **S3 buckets ARN**: Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same AWS account and cloud region as the Amazon S3 you want to protect.
- **Prefix**: Add a prefix value to the name of the inventory report that will be created.
  - The prefix can contain any character, including white spaces.

    For example, *Report-Source Bucket*.
  - The prefix should not begin or end with a forward slash (/) .
  - The prefix should not contain consecutive forward slashes.
  - You must not upload any files in the prefix of the inventory report.

  The inventory report will be created at *<Prefix>/<Path_to_inventory_report>* when you initiate the protection. Cohesity creates the path to the inventory report on AWS based on the cluster, Amazon S3 bucket, and protection.

12. Click **Register**.

Register an AWS C2S Source Using IAM Role

Cohesity supports Commercial Cloud Services (C2S) region in AWS backup and recovery workflows. C2S is AWS's secret region offered by AWS to the US Intelligence community. You can register an AWS C2S source only using an IAM role. Contact Cohesity Support to enable the feature.

1. Navigate to **Data Protection > Sources**.
2. Select **Register** > **Virtual Machines**.
3. From the **Select Hypervisor Source Type** drop-down, select **AWS : IAM User/Role**.
4. Choose **Category** as **C2S**.
5. Configure the following parameters:

Field Name	Description
Base URL	URL of the Commercial Cloud Services (C2S) server that contains the Amazon storage provided by the C2S agency.
Agency	The name of the agency that is providing the Commercial Cloud Services (C2S) server.
Mission	The account name for accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.

Field Name	Description
Role	The IAM role to use when accessing the Amazon storage which is provided by the C2S agency.
Password	The password to read the Client Private Key when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Server CA Trusted Certificate	Upload the Certificate Authority Trusted Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Certificate	Upload the Client Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Private Key	Upload the Client Private Key file to use when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.

6. (*Applies only for Amazon EC2 instance protection*) **Specify Fleet Settings**: Enable this option to:
   - Specify the VPC and subnet in which the EC2 fleet instance should be launched.
   - Add tags to the EC2 fleet instance.

   By default, this option is disabled and the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

   Specify the VPC and Subnet

   From the **Fleet Instance location** drop-down list select one of the following options to specify the VPC and subnet in which the fleet instance should be launched.
   - **Same Subnet as protected VMs**: The fleet instance will be launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.
   - **Select New**: The fleet instance will be launched in the specified VPC and subnet of the specified region. If the AWS source is not registered with the Cohesity cluster, the cluster cannot detect the AWS region, and the associated VPC and subnets. To configure a new VPC and subnet for the fleet instance, you must first register the source and then update the registered source to specify the VPC and subnet. For more information, see Register AWS Account using IAM Role.

   If you have not selected the **Fleet Instance location** after enabling the **Specify Fleet Settings** option, then the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

Add a Tag

Tags enable you to categorize and identify your fleet instances. A tag consists of a key and an optional value, both of which you define.

To add a tag to the fleet instances, click **Add Tags** and enter a tag key and value. You can add up to 47 tags to a fleet instance.

Delete a Tag

To delete an existing tag, click the **cross** icon next to the tag.

7. (*Applies only for Amazon S3 buckets protection*) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Enable **Specify S3 Inventory Report Location**, provide the following information to create the inventory report:
   - **S3 buckets ARN**: Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same AWS account and cloud region as the Amazon S3 you want to protect.
   - **Prefix**: Add a prefix value to the name of the inventory report that will be created.
     - The prefix can contain any character, including white spaces.

       For example, *Report-Source Bucket*.
     - The prefix should not begin or end with a forward slash (/) .
     - The prefix should not contain consecutive forward slashes.
     - You must not upload any files in the prefix of the inventory report.

     The inventory report will be created at *<Prefix>/<Path_to_inventory_report>* when you initiate the protection. Cohesity creates the path to the inventory report on AWS based on the cluster, Amazon S3 bucket, and protection.
8. Click **Register**.

Configure a New VPC and Subnet for Fleet Instance

The EC2 fleet instance will be launched in the specified subnet when you back up the Amazon EC2 instance to the Cohesity cluster. To configure a new VPC and subnet for the fleet instance:

1. In the Cohesity Dashboard select **Data Protection** > **Sources**.
2. Navigate the hierarchy to the registered AWS source for which you want to configure a new VPC and subnet for the fleet instance.
3. Click **Edit** next to the **Specify Subnet for Fleet Instance** text , or click the actions menu next to the AWS source and select **Edit**.

   The Edit Source page appears.

4. Enable the **Specify Fleet Settings** option if you have not enabled when registering the AWS source.
5. From the **Fleet Instance Location** drop-down list, select **Select New**.
6. Select a region, and its VPC and subnet in which the fleet instance should be launched.

   > **Note:** As the Amazon EC2 sources are mostly located in one region, you can select the VPC and subnet of only one region.

7. Click **Save**

Run CloudFormation Template

To complete the registration of AWS account with the Cohesity cluster, you must run the Cohesity's CloudFormation Template (CFT) on the AWS account you are registering. You can download CFT from the Cohesity cluster UI as you register the AWS source.

When you execute the CFT, it creates an AWS IAM role and the policies for accessing all IAM roles and KMS keys in your account.

Execute the CloudFormation Template

1. Log in to AWS Management Console.
2. Navigate to **Services** and search for **CloudFormation**.
3. Click **Create stack** > **With new resources** to create a stack under AWS CloudFormation using the downloaded CloudFormation Template.
4. In the **Prerequisite - Prepare template** page, select **Upload a template file**.
5. Click **Choose File** to browse and upload the JSON CloudFormation template. Click **Next**.
6. Enter a **Stack name** and specify the AWS Account ID where the Cloud Edition cluster is running.
7. Select one of the following based in your requirement:
   - **source_registration**: Select this option are if you are creating IAM role for registration AWS source.
   - **external_target**: Select this option if you are creating IAM role for cloud tiering or cloud archiving purposes.
   - **both**: Select this option if you are rcreating IAM role for both of the above purposes.
8. Click **Next**.

9. Acknowledge the creation of IAM resources by selecting the checkbox, then click **Create Stack**.
10. Once the stack creation is successful, view the **Resources** tab to check the created IAM roles.
11. The following IAM roles and profiles are created:
    - *CohesityInstanceRole*: Attached to Cohesity cluster EC2 nodes; used as the base identity for assuming other roles.
    - *CohesityInstanceProfile*: EC2 instance profile that wraps *CohesityInstanceRole*; must be attached to all cluster nodes.
    - *CohesitySourceRegistrationRole*: Grants permissions for source registration operations.
    - *CohesityBackupS3StagingRole*: Grants permissions for S3 staging operations during backup and recovery.

Attach the IAM Role to Cohesity Nodes

To complete registration, attach the *CohesityInstanceProfile* role to all nodes in the cluster:

1. Log in to AWS Management Console.
2. Navigate to **Services** > **EC2** > **Instances**.
3. Select the cluster created. For a multi node, the IAM role must be attached on all nodes.
4. Select **Actions** > **Security** > **Modify IAM Role**.
5. In the **IAM Role** field, enter the name/search for the Instance Profile (*CohesityInstanceProfile*) that was created.
6. Click **Update IAM role**.

   The IAM role is successfully attached to that EC2 instance.

> **Note:** If any new node is added to the cluster, then this role has to be attached to every newly added node in the cluster as well.

Complete AWS Source Registration

Once the IAM role is attached, complete the AWS source registration on the cluster. For more information, see Register AWS Account using IAM Role.

Register AWS Account Using IAM User

The IAM User method allows you to register an AWS account using user credentials instead of IAM roles. You manually create a policy, create an IAM user with the required

permissions, and provide the user's access keys during registration.

Prerequisites

For registering using IAM User method, you must first create an AWS policy with required permissions, and then create an IAM user with this policy attached. The permissions required are provided in the following procedure.

1. Create an AWS Policy
2. Create an IAM User

Create an AWS Policy

To restrict Cohesity's access to your AWS resources, Cohesity recommends creating an IAM policy with only the permissions required for backup and recovery.

To create a new policy with permissions to create the cluster:

1. Log in to the Amazon AWS console using your AWS account with admin privileges.
2. In the top left of the menu bar, select **Services**.
3. Under **Security, Identity & Compliance**, select **IAM**.
4. In the left frame, click **Policies**.
5. Click **Create policy**.
6. In the **Create policy** page, select the **JSON** tab.
7. Update the default JSON code with the permissions required for the workloads your want to protect. For information on the permissions required, see Minimum Permissions for Individual Workloads and Workflows.
8. Click **Review policy**.
9. In the **Name** field, enter a name of this policy, such as `CohesityPolicy` and click **Create policy**.

Create an IAM User

Once you create a policy with the required permissions, you can create a new IAM user and attach the policy that you created previously.

To create an IAM user:

1. Log in to the Amazon AWS console using your AWS account with admin privileges.
2. In the top left of the menu bar, select **Services**.
3. Under Security, Identity & Compliance, select **IAM**.
4. In the left frame, click **Users**.

5. Click **Add users**.
6. In the User name field, enter a user name such as `CohesityCE`.
7. For **Access type**, select Programmatic access and click **Next: Permissions**.
8. Click **Attach existing policies directly**, then search for the name of the policy you created previously and select that policy.
9. Click **Next: Tags**.
10. Optionally, enter tags for the policy and then click **Next: Review**.
11. Review your choices and then click **Create User**.
12. In the **Secret access key** column, click **Show**.
13. Copy and paste the **Access key ID** and the Secret access key into a text file or click **Download .csv** to save the information onto your desktop or a corporate share.

    These keys are required when registering AWS Source with the Cohesity cluster.
14. Click **Close**.

Register AWS Account Using IAM User Method

To register the AWS Account as the source:

1. Select **Data Protection** > **Sources**.
2. Select **Register** > **Virtual Machines**.
3. From the **Select Hypervisor Source Type** drop-down, select **AWS : IAM User/Role**.
4. Select the AWS category **Standard** or **Gov**.
5. Choose the **Authentication Method** as an **IAM User**.
6. Enter the information described below and click **Register**.

Authentication Method	Steps
Access Key ID	Specify the **Access Key ID** for the new IAM user that was created to back up a AWS Cloud source or convert and clone Amazon EC2 to an AWS Cloud Source. For general information about AWS keys, see Access Keys (Access Key ID and Secret Access Key).
Secret Access Key	Specify the **Secret Access Key** for the new IAM user that was created to back up a AWS Cloud Source or convert and clone Amazon EC2 to an AWS Cloud Source. For general information about AWS keys, see Access Keys (Access Key ID and Secret Access Key).

Authentication Method	Steps
ARN	Specify the Amazon Resource Name (**ARN**) of the IAM user in the following format: `arn:aws:iam::<account-id>:user/<username>` For more information, see Amazon Resource Names (ARNs).

7. (*Applies only for Amazon EC2 instance protection*) **Specify Fleet Settings**: Enable this option to:
   - Specify the VPC and subnet in which the EC2 fleet instance should be launched.
   - Add tags to the EC2 fleet instance.

   By default, this option is disabled and the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

   Specify the VPC and Subnet

   From the **Fleet Instance location** drop-down list select one of the following options to specify the VPC and subnet in which the fleet instance should be launched.
   - **Same Subnet as protected VMs**: The fleet instance will be launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.
   - **Select New**: The fleet instance will be launched in the specified VPC and subnet of the specified region. If the AWS source is not registered with the Cohesity cluster, the cluster cannot detect the AWS region, and the associated VPC and subnets. To configure a new VPC and subnet for the fleet instance, you must first register the source and then update the registered source to specify the VPC and subnet. For more information, see Register AWS Account Using IAM User.

   If you have not selected the **Fleet Instance location** after enabling the **Specify Fleet Settings** option, then the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are to be protected.

   Add a Tag

   Tags enable you to categorize and identify your fleet instances. A tag consists of a key and an optional value, both of which you define.

   To add a tag to the fleet instances, click **Add Tags** and enter a tag key and value. You can add up to 47 tags to a fleet instance.

   Delete a Tag

   To delete an existing tag, click the **cross** icon next to the tag.

8. (*Applies only for Amazon S3 buckets protection*) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Enable **Specify S3 Inventory Report Location**, provide the following information to create the inventory report:
    - **S3 buckets ARN**: Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same AWS account and cloud region as the Amazon S3 you want to protect.
    - **Prefix**: Add a prefix value to the name of the inventory report that will be created.
        - The prefix can contain any character, including white spaces.

          For example, *Report-Source Bucket*.
        - The prefix should not begin or end with a forward slash (/) .
        - The prefix should not contain consecutive forward slashes.
        - You must not upload any files in the prefix of the inventory report.

      The inventory report will be created at *<Prefix>/<Path_to_inventory_report>* when you initiate the protection. Cohesity creates the path to the inventory report on AWS based on the cluster, Amazon S3 bucket, and protection
9. Click **Register**.

Configure a New VPC and Subnet for Fleet Instance

The EC2 fleet instance will be launched in the specified subnet when you back up the Amazon EC2 instance to the Cohesity cluster. To configure a new VPC and subnet for the fleet instance:

1. In the Cohesity Dashboard select **Data Protection** > **Sources**.
2. Navigate the hierarchy to the registered AWS source for which you want to configure a new VPC and subnet for the fleet instance.
3. Click **Edit** next to the **Specify Subnet for Fleet Instance** text , or click the actions menu next to the AWS source and select **Edit**.

   The Edit Source page appears.
4. Enable the **Specify Fleet Settings** option if you have not enabled when registering the AWS source.
5. From the **Fleet Instance Location** drop-down list, select **Select New**.
6. Select a region, and its VPC and subnet in which the fleet instance should be launched.

> **Note:** As the Amazon EC2 sources are mostly located in one region, you can

> select the VPC and subnet of only one region.

7. Click **Save**

Manage AWS Acccount Registration

After registering your AWS account as a source, you can:

- Update the AWS source configuration.
- Refresh the source details.
- Unregister the AWS source.

Update AWS Account Source Configuration

You can edit the registered AWS account to add or remove the Amazon services protected by the Cohesity from your AWS account.

To edit the AWS account source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AWS account source and select Edit.
3. In the **Edit AWS Source** page, select or unselect the Amazon services.
4. Click **Update**.

Refresh AWS Account Source

You can refresh the AWS account source configuration and fetch the latest changes on your AWS account.

To refresh the AWS source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the AWS account source and select **Refresh**.

Unregister AWS Account Source

If you plan to stop taking the backup of your AWS source, you can unregister the AWS source from the Cohesity cluster.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the AWS source.

To unregister the AWS source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AWS account source and select **Unregister**.

Amazon Elastic Compute Cloud Instances

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Amazon Elastic Compute Cloud (EC2) instances in your AWS account.

Plan and Prepare for AWS EC2 Protection

Prerequisites

Required AWS Service Endpoints

AWS Resource Type	Service Endpoints
IAM	iam.amazonaws.com
CloudFormation	cloudformation.<regionId>.amazonaws.com
EC2	ec2.<regionId>.amazonaws.com
KMS	kms.<region>.amazonaws.com
S3	s3.<region>.amazonaws.com s3-fips.<region>.amazonaws.com
SSM	ssm.<region>.amazonaws.com
EBS	ebs.<region>.amazonaws.com
RDS	rds.<region>.amazonaws.com

Recovery Prerequisites

Same Account, Same Region:

No prerequisites.

Different Account, Same Region:

- Target AWS account should be registered as a data source in the same destination cloud region.
- If you have enabled **Encryption Settings** for an EC2 instance that is already

encrypted, then the EC2 instance will be recovered with the encryption provided in the **Encryption Settings**.

Different Account, Different Region:

- Target AWS account should be registered as a data source in the same destination cloud region.
- If you have enabled **Encryption Settings** for an EC2 instance that is already encrypted, then the EC2 instance will be recovered with the encryption provided in the **Encryption Settings**.

Considerations

General Considerations

- When recovering an AWS EC2 instance running Microsoft Windows OS with UEFI Preferred boot mode, the process will utilize a base AMI based on the 'Current Instance Boot Mode' of the source EC2. For example, if the source EC2 was configured as UEFI-preferred, the recovery will automatically use an AMI that supports UEFI boot mode.
- Cohesity does not support recovery of EC2 instances to a different availability zone.
- Backups for ARM-based Amazon EC2 instances from the marketplace are only supported using the EBS diff API-based backup method.
- You cannot exclude EBS volumes at the object-level (manual volume exclusion) for auto-protected EC2 instances. To exclude EBS volume for auto-protect VMs, you must use the **Volume Exclusion Settings** option in the Additional Settings.

File-Level Recovery (FLR) Considerations

- File-level recovery is not supported if EC2 instances are backed up using Snapshot Manager as the protection type.
- File-level recovery is not supported for EC2 instances using the XFSv5 file system.
- You can perform file and folder recovery to a Linux EC2 instance only if it uses x86_64 architecture.

  When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.

AWS C2S Considerations

- Amazon EC2 Simple Systems Manager (SSM) is not available in C2S. You must install the Cohesity Agent on the target machine for FLR.
- AWS diff APIs are not available in C2S.
- All VM restores are fleet-based.

Networking and Configuration Issues

Issue: The backup of the AWS VMs will fail with the following error if you have configured to launch the fleet instance in the VPC and subnet of a region different from the protected VM's region: `VPC and Subnet are not specified for the fleet instance. Please specify the same by editing the source`.

Workaround: Update the registered AWS source to launch the fleet instance in the same subnet as protected VMs.

Support for HTTP Proxy Server

HTTP Proxy is used if the cluster using the proxy server is behind a private network and cannot reach the internet to go through the proxy. Cohesity cluster supports HTTP proxy for AWS sources for performing the following workflows:

- On-demand CloudSpin
- Policy-based CloudSpin
- Cloud-native backup and recovery
- Disaster recovery failback workflows

To enable HTTP proxy, contact Cohesity Support.

Protect AWS EC2 Instances

In Cohesity, AWS EC2 instances are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Create or Edit a Standard Policy.

> **Note:**
> - Backup and recovery of AWS EC2 requires a proxy VM to be launched on the cloud source. For enhanced security, after creating the proxy VM on the cloud source, Cohesity now automatically deploys an X.509 certificate .
> - Cohesity preserves the Instance Metadata Service v2 (IMDSv2) configuration of Amazon EC2 instances during protection.

To protect AWS EC2 instances:

1. Select **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then click **Protect > Virtual Machines**. The **New Protection** page appears.

3. Click **Add Objects**, and from the **Registered Source** drop-down list, select the AWS source you want to protect. You can also register a new source with the Cohesity cluster by clicking the **Register Source** option.
4. From the objects list, select the EC2 instances you want to protect.

   Optionally, you can:

   - Click the icon to enable Auto Protect at various hierarchy levels (account, AWS region, or Availability zone levels). When this option is enabled at a particular level, all the EC2 instances that are added to that level in the future are automatically added to the Protection Group and are protected from the next protection run.
   - Click the **Tags** icon on the right and select tags to:
     - Auto Protect ( ) the VMs having these tags.
     - Exclude (⊘) the auto protection of VMs having these tags.

     If a VM has multiple tags with Auto Protect and Exclude applied, exclusion takes precedence. These excluded VMs are also excluded even if a parent object is auto protected.

     The **Tags** button will not be displayed if the Cohesity cluster detects that no sources have tags assigned to them.
   - Exclude any volumes of the EC2 instance from protection. To exclude a volume:
     - Click the edit icon displayed next to the EC2 instance of which you want to exclude the volume.

       Unselect the volume which you want to exclude from protection.

       > **Note:**
       > - You cannot exclude the root volume from protection.
       > - You cannot exclude EBS volumes at the object level (manual volume exclusion) for auto-protected EC2 instances. To exclude EBS volume for auto-protect VMs, you must use the **Volume Exclusion Settings** option in the Additional Settings.

       For more information, see Exclude EBS Volumes from EC2 Instance Backup.
5. Click **Continue** to go back to the **New Protection** page.
6. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

7. From the Policy drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
8. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, click **Protect**.
9. Select a **Protection Group Type**:
    - **Native Snapshot**—Back up EC2 instances in an AWS cloud service using the native snapshot API of the cloud service.
    - **Snapshot Manager**—Backup EC2 instances in an AWS cloud service using the Cloud Snapshot Manager APIs of the cloud service.
10. Optionally, from the **Policy** drop-down list, you can change the policy you selected earlier policy or create a new Protection Policy.
11. Select **Storage Domain** and select an existing storage domain or create a new one by selecting Add Storage Domain.
12. **Start Time**: Available only if the selected policy is set to Backup Daily. Indicates when the Protection Group should run. The current time is displayed by default which can be changed. You can change the time zone by selecting a different time zone.
13. **SLA**: The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full**: Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
14. If you need to change any of the advanced settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
15. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.

Additional Settings	Description
End Time	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Indexing	This option is toggled on by default to index the data of the Protection Group. Indexing improves files and folders search. Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.
Volume Exclusion Settings	EC2 disks can be excluded based on disk tags using simple query rules. For example, the query, `(type = log AND environment IN (qa, dev)) OR exclude = yes`, excludes all volumes from backup for which the volume tags meet the above condition. For more details, see Exclude EBS Volumes from EC2 Instance Backup. **Note:** Exclusion of disks enabled at the object-level (manual exclusion) will take precedence over the exclusion configured here.
Cancel Run at Quiet Start Time	*Available only if the selected policy has at least one quiet time period.* Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. If the option is toggled off, then after a protection run starts, it continues to execute even when a quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period.

Additional Settings	Description
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
Description	Enter a brief description about the Protection Group.

Exclude EBS Volumes from EC2 Instance Backup

Cohesity supports excluding EBS volumes from EC2 instance backup. The following are important points to note:

- By default, all volumes are included as part of the EC2 instance backup.
- You cannot exclude the root volume.
- You cannot exclude EBS volumes at the object level (manual volume exclusion) for auto-protected EC2 instances. To exclude EBS volume for auto-protect VMs, you must use the **Volume Exclusion Settings** option in the Additional Settings.
- If you exclude a volume from the EC2 instance backup and if any system information is stored on the excluded volume, recovering the EC2 instance fails.

  For example, in the `/etc/fstab` file, if you do not specify the `nofail` option for the excluded volume, recovering the Linux instance fails.

You can exclude EBS volumes from EC2 instance backup using the following two methods:

- Manual volume exclusion (object-level)
- Volume exclusion using rules (Protection Group-level)

Manual Volume Exclusion

To exclude EBS volumes manually while creating/editing the Protection Group, perform the following steps:

1. Follow steps 1-4 in Protect AWS EC2 Instances.
2. After you select the EC2 instance, the **All Volumes** option is displayed:

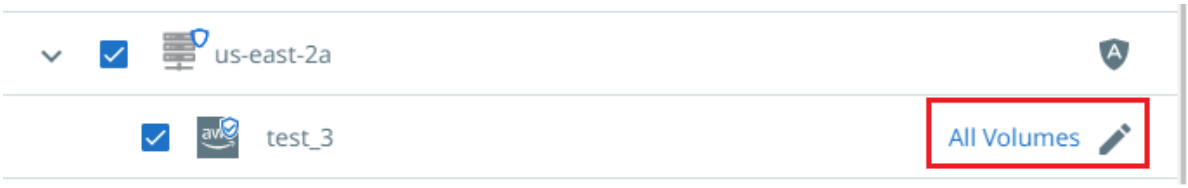

   To edit the volume selection after creating the Protection Group, hover over the EC2 instance, click the action icon (⋮) next to the Protection Group, and select **Edit**.
3. Click **All Volumes**.

   The volumes that are included as a part of the backup are displayed:

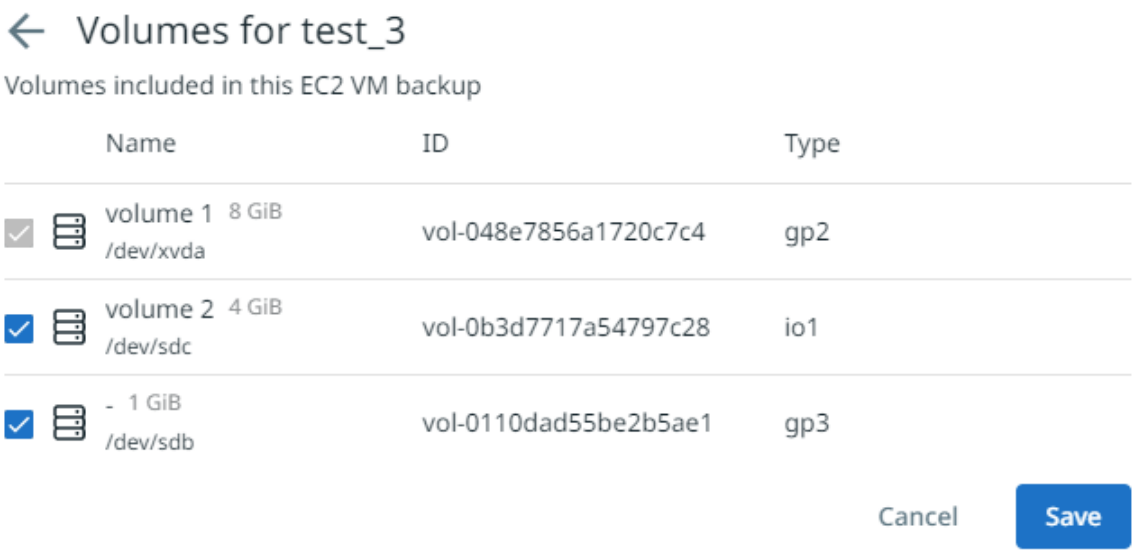

4. Select the volume to be excluded by unchecking the corresponding check box.
5. Click **Save**.

The volumes excluded are displayed next to the corresponding EC2 instance:

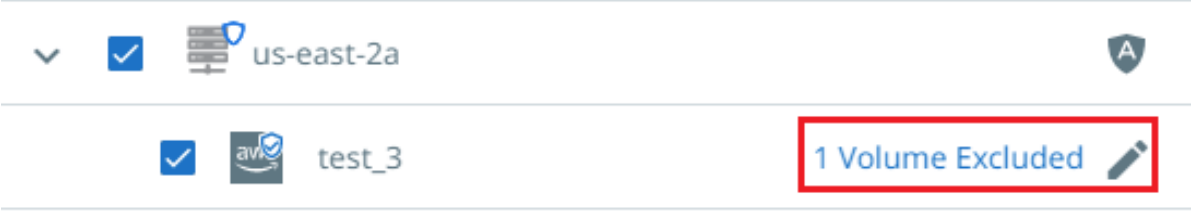

Volume Exclusion Using Rules

To exclude EBS volumes after the Protection Group is created, complete the following steps:

1. Follow steps 1-10 in Protect AWS EC2 Instances.
2. In the Additional Settings section, set the rules using queries in the Volume Exclusion Settings field. EC2 disks can be excluded based on disk tags using simple query rules. For example, the query, (type = log AND environment IN (qa, dev)) OR exclude = yes, excludes all volumes from the backup for which the volume tags meet the above condition.

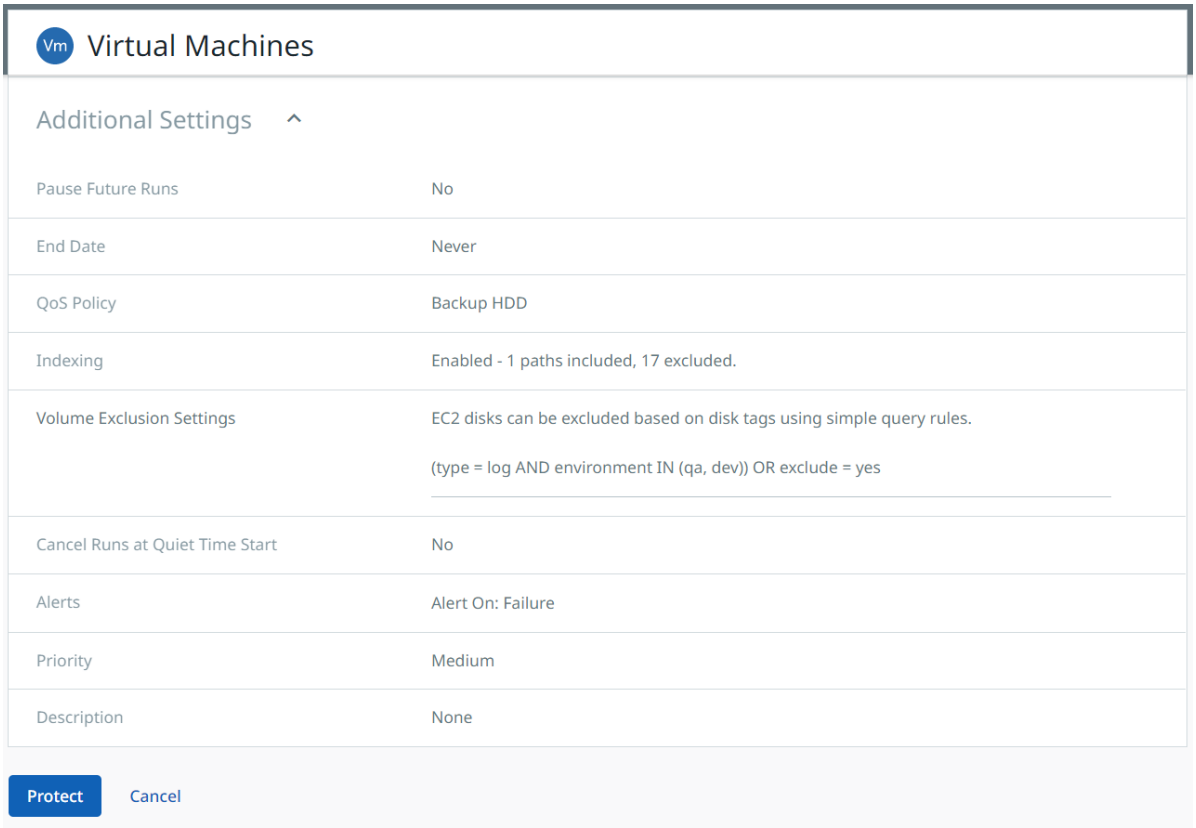

3. Click **Protect**.

To determine whether EBS volumes have been excluded from the backup, complete the following steps:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, click on the protection job.
3. On the **Runs** tab, click on the backup job that is complete.
4. On the **Run Details** page, click on the VM name. The **Backup Task Activity** pop-up window lists the details of all the tasks.

   The details of the excluded volume are displayed:

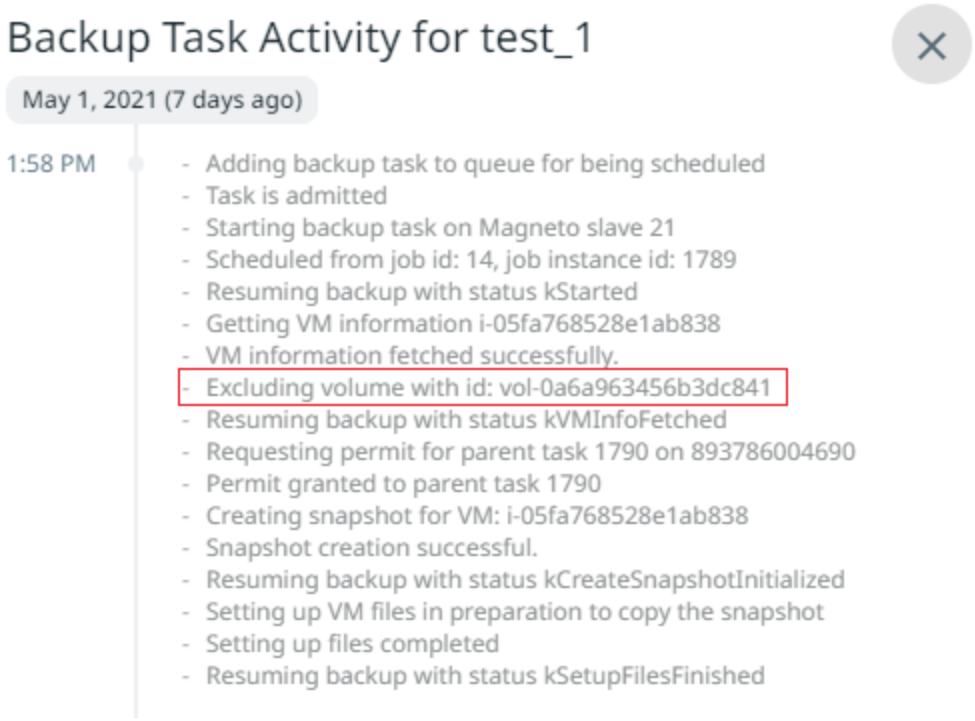

AWS Backup Use Cases

The following are the use cases for backing up Amazon EC2 instances:

Backup On-prem Data Using Cloud Edition

Cloud Edition supports the backup and recovery of Virtual Machines (such as a VM in a vCenter server) to a Cohesity cluster Cloud Edition running in AWS Cloud Service.VMs can be recovered to a new location or back to the original location.

Backup Amazon EC2 to On-Premises or Cloud Edition Cluster using Native Snapshot

Cohesity cluster supports backing up virtual servers using the native snapshot of the cloud service. For example, a virtual server Protection Group backs up Amazon EC2 instances from AWS Cloud service to an on-premises Cohesity cluster. The result of this backup is a snapshot on the on-premises Cohesity cluster.

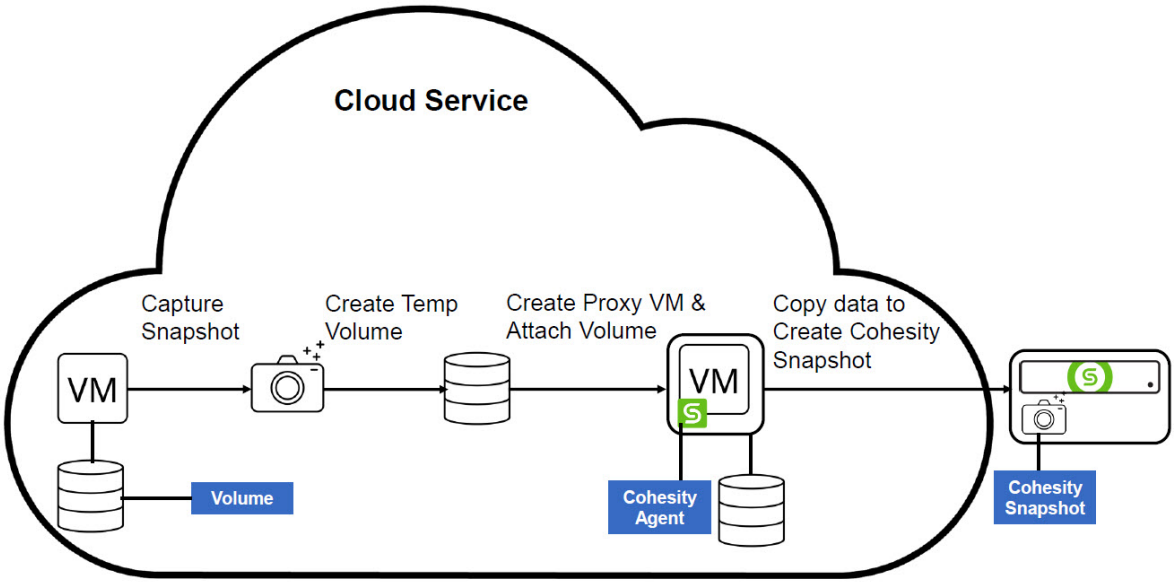

Cohesity Cloud Edition also supports backing up Amazon EC2 instances using the native snapshot of the Cloud Service. For example, a Virtual Server Protection Group backs up Amazon EC2 instances to a Cloud Edition Cohesity cluster running in the same Cloud service. The result of this backup is Snapshots on the Cloud Edition cluster.

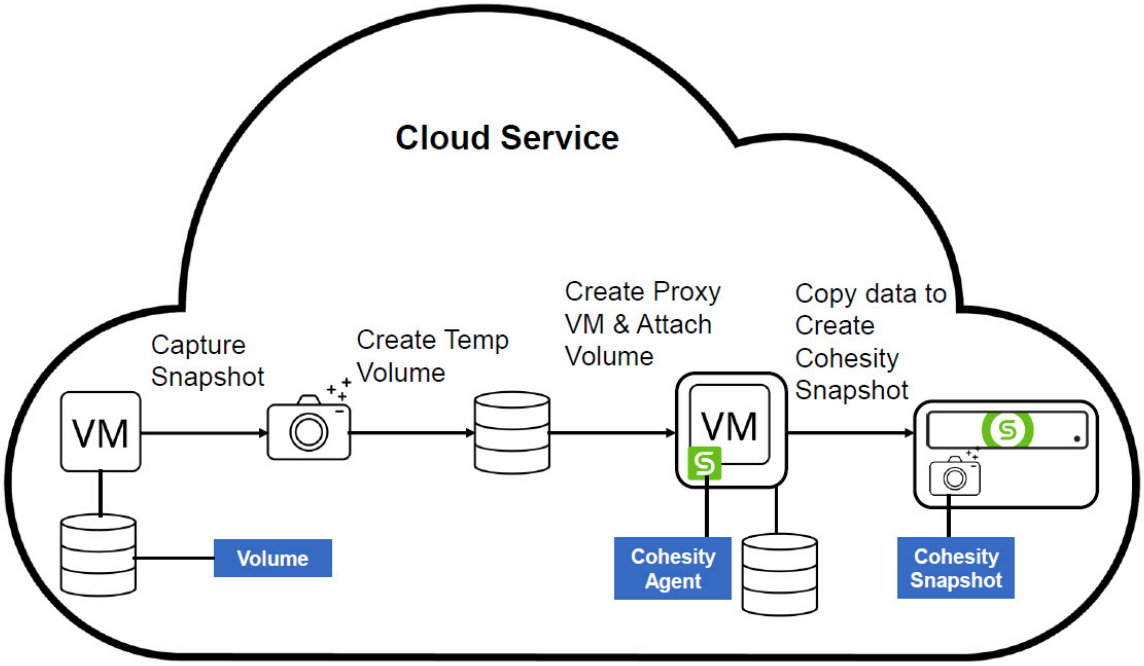

Backup Amazon EC2 Using Snapshot Management

Cohesity uses Snapshot Management APIs to create AWS snapshots for the services such as Amazon EC2 and Amazon RDS. When a Protection Group with Snapshot Manager enabled is

executed, only the metadata is stored on the Cohesity cluster. All snapshot data gets stored in AWS itself.

To enable this feature, you must select Snapshot Manager as the Protection Group Type while creating the Protection Group for protecting AWS service. By default, the snapshots are stored on the same region and account as the AWS source. You can recover these snapshots to a different region and account. Additionally, the snapshots stored in the same AWS region and account as the source can be replicated to a different AWS account or region. For more information, see Use AWS Snapshot Management for Replication.

Use AWS Snapshot Management for Replication

If the policy attached to the Snapshot Manager-enabled Protection Group has replication to AWS enabled, you can replicate the data to different AWS regions or accounts.

To replicate the AWS snapshots to another AWS region:

1. Register AWS Source.
2. Create a Policy with replication to AWS enabled.
3. Create a Protection Group:
    - Choose Snapshot Manager as the Protection Group Type.
    - Attach the policy with replication enabled.

This type of replication does not require a Cohesity cluster to serve as the destination for the replication of data. The AWS region serves as the destination. Cohesity replicates the snapshots to the EBS bucket in the selected region.

Manage Recovery

This chapter includes the following topics:

Recover Files or Folders for AWS Native Snapshots to the New Server

Prerequisites

- Firewall port 50051 must be open on the target EC2 instances. For more information, see Manage Firewall Ports
- The target EC2 instances must be reachable with private IP.
- The target EC2 instances must have the IAM role “AmazonEC2RoleforSSM”.
- The AWS SSM agent must be installed on target EC2 instances.

> **Note:** For enhanced security, when installing the agent on the target EC2, Cohesity automatically deploys an X.509 certificate.

> **Note:** You can perform files and folder recovery to an EC2 Linux instance only if the EC2 Linux instance is of x86_64 architecture.

Recover Files and Folders to New Server

Perform the following steps to recover files and folders to the new server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on Cohesity.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on Cohesity.

If you select..	Then..
Files and Folders	a. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. For more information, Use Wildcard * in Search for Recovery. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. b. From the list of search results, select the required files or folders to recover. c. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: i. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. ii. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. iii. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. d. Go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial server or Protection Group name. You can also specify the wildcard character *. For more information, Use Wildcard * in Search for Recovery. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up the browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears.

If you select..	Then..
	2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. 8. Go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Server Type	Volume Type	

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next: Recover Options**: If you select this option, then continue with the step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':', it is replaced with '_COLON_' in the archive.
     Example:

- The original name is "a:b" then the archived name is "a_COLON_b"
- The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
Example:

- The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **New Server** to recover the files or folders to a new server.
6. **Source**: From the **Source** drop-down menu, select a registered server to which you want to recover the files and folders. You can also register a new source by clicking **Register Source**.
7. **Target**: From the Target (for VM and physical server) or Volume (for NAS) drop-down menu, select a VM or volume to which the files and files are to be recovered.

   > **Note:** For recovering files and folders of a VMware VM, ensure that VMware Tools are installed and running on the target VM. The target list displayed includes only VMs with VMware Tools installed and running.

8. **Recover to**: In the **Recover to** field, change the default location to which the files or folders are recovered. By default, the files and folders are recovered to the /**tmp**/ directory for VMs.

   For NAS, an example path on an alternate location could be /tmp/recover. You can recover encrypted files/folders to their original location only.

   Example for Windows alternate location:

   - D:\cohesity_restore\1.txt
   - D:\cohesity_restore\bar

   User Account Requirements for VM Protection may need to be met to recover files to a Windows VM.

   If recovering to a standalone Hyper-V host or cluster, you must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.
9. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

- C:\restore\1.txt
- C:\restore\cohesity<taskid>_0\1.txt
- C:\restore\bar
- C:\restore\cohesity<taskid>_1\bar

If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

- **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.
- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

10. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Files or Folders for AWS Native Snapshots to the Original Server

Prerequisites

- Firewall port 50051 must be open on the target EC2 instances. For more information, see Manage Firewall Ports
- The target EC2 instances must be reachable with private IP.
- The target EC2 instances must have the IAM role “AmazonEC2RoleforSSM”.
- The AWS SSM agent must be installed on target EC2 instances.

> **Note:** For enhanced security, when installing the agent on the target EC2, Cohesity automatically deploys an X.509 certificate.

> **Note:** You can perform files and folder recovery to an EC2 Linux instance only if the EC2 Linux instance is of x86_64 architecture.

Recover Files and Folders to Original Server

Perform the following steps to recover files and folders to the original server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on Cohesity.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on Cohesity.

If you select..	Then..
Files and Folders	a. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. For more information, Use Wildcard * in Search for Recovery. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. b. From the list of search results, select the required files or folders to recover. c. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: i. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. ii. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. iii. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. d. Go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial server or Protection Group name. You can also specify the wildcard character *. For more information, Use Wildcard * in Search for Recovery. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up the browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option.

If you select..	Then..
	3. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. 8. Go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Server Type	Volume Type	

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next: Recover Options**: If you select this option, then continue with the step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':', it is replaced with '_COLON_' in the archive.
     Example:

- The original name is "a:b" then the archived name is "a_COLON_b"
- The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
Example:

- The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **Original Server** to recover the files or folders to the original server.
6. **Recover to Original Path**: To recover to a different location in the original server, disable the **Recover to Original Path** option, and then provide the location to which the files or folders are to be recovered in the **Recover To** field. By default, the files and folders will be recovered to the original location.
7. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

     - C:\restore\1.txt
     - C:\restore\cohesity<taskid>_0\1.txt
     - C:\restore\bar
     - C:\restore\cohesity<taskid>_1\bar

     If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
   - **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.

- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

8. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover AWS EC2 Instance to New Location

To recover an AWS EC2 instance to a new location:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **VMs**.

   The Virtual Machines page appears.
3. Search for an EC2 instance or Protection Group that contains the snapshots to recover from by entering characters of the EC2 instance name, Protection Group name, or by specifying the wildcard character *.

   A list of EC2 instance displays. For more information, Use Wildcard * in Search for Recovery.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more EC2 instances or Protection Groups that contain the snapshots to recover.
   - If you have selected a Protection Group, then all the EC2 instances backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
   - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select an EC2 instance or Protection Group, the Storage Domain, and Source filters will automatically get locked.
5. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a...	Then...
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the **Protection** Run drop-down menu, select the snapshot from which you want to recover the EC2 instance. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a...	Then...
EC2 instance	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. b. Select the snapshots to be recovered. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived in the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** and provide the following details:
    a. **Source**: Select a registered AWS account as the new recovery destination.
    b. **Region**: Select a destination AWS region.
    c. **Key Pair**: Select an AWS key pair to be associated with the recovered EC2 Instance(s).
    d. **Subnet**: Select the subnet where the recovered EC2 Instance(s) will be launched.
    e. **Network Security Group**: Select the network security group to be associated with the recovered EC2 Instance(s).
8. To recover AWS EC2 instances after protection, Cohesity creates fleet instances to transfer data from the snapshots stored on the Cohesity cluster. Enable **Specify Fleet Settings** to:
    - Specify the VPC and subnet in which the EC2 fleet instance should be launched.

      From the Fleet Instance location drop-down list select one of the following options to specify the VPC and subnet in which the fleet instance should be launched.

        a. **Same Subnet as protected VMs**: The fleet instance will be launched in the same VPC and subnet of the Amazon EC2 instances that are being protected.
        b. **Select New**: The fleet instance will be launched in the specified VPC and subnet of the specified region. If the AWS source is not registered with the Cohesity cluster, the cluster cannot detect the AWS region, and the associated VPC and subnets. To configure a new VPC and subnet for the fleet instance, you must first register the source and then update the registered source to specify the VPC and subnet.

      If you have not selected the **Fleet Instance location** after enabling the **Specify Fleet Settings** option, then the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are being protected.
    - **Add tags to the EC2 fleet instance**

      Tags enable you to categorize and identify your fleet instances. A tag consists of a key and an optional value, both of which you define.

      To add a tag to the fleet instances, click **Add Tags** and enter a tag key and value. You can add up to 47 tags to a fleet instance.

      To delete an existing tag, click the **cross** icon next to the tag.

By default, this option is disabled and the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are protected.

9. Under **Recovery Options**, edit the default settings. Click the **edit** icon to edit the settings:
    a. **Rename**: Add a Prefix and/or Suffix to the recovered Amazon EC2 instances.
    b. **Power State**: To ensure that the restored EC2 instances are functioning properly, the EC2 instances are powered on by default after the restoration process. However, you can enable the **Power Off** option if you prefer to keep the VMs powered off once Cohesity verifies their functionality. If you are recovering VMs from a Storage Domain with Cloud Tier enabled, Cohesity recommends keeping the VMs powered off.
    c. **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
    d. **Task Name**: Change the default name of the recovery task.
10. Click **Recover**.

Cohesity begins to restore the selected Amazon EC2 instances.

Recover AWS EC2 Instance to Original Location

To recover an AWS EC2 instance to the original location:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **VMs**.

   The Virtual Machines page appears.
3. Search for an EC2 instance or Protection Group that contains the snapshots to recover from by entering characters of the EC2 instance name, Protection Group name, or by specifying the wildcard character *.

   A list of EC2 instance displays. For more information, Use Wildcard * in Search for Recovery.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more EC2 instances or Protection Groups that contain the snapshots to recover.
    - If you have selected a Protection Group, then all the EC2 instances backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.

- You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select an EC2 instance or Protection Group, the Storage Domain, and Source filters will automatically get locked.

5. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a...	Then...
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the **Protection** Run drop-down menu, select the snapshot from which you want to recover the EC2 instance. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a...	Then...
EC2 instance	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. b. Select the snapshots to be recovered. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the EC2 instance to the original location.
8. To recover AWS EC2 instances after protection, Cohesity creates fleet instances to transfer data from the snapshots stored on the Cohesity cluster. Enable **Specify Fleet Settings** to:
    - Specify the VPC and subnet in which the EC2 fleet instance should be launched.

      From the Fleet Instance location drop-down list select one of the following options to specify the VPC and subnet in which the fleet instance should be launched.

        a. **Same Subnet as protected VMs**: The fleet instance will be launched in the same VPC and subnet of the Amazon EC2 instances that are being protected.

        b. **Select New**: The fleet instance will be launched in the specified VPC and subnet of the specified region. If the AWS source is not registered with the Cohesity cluster, the cluster cannot detect the AWS region, and the associated VPC and subnets. To configure a new VPC and subnet for the fleet instance, you must first register the source and then update the registered source to specify the VPC and subnet.

      If you have not selected the **Fleet Instance location** after enabling the **Specify Fleet Settings** option, then the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are being protected.
    - **Add tags to the EC2 fleet instance**

      Tags enable you to categorize and identify your fleet instances. A tag consists of a key and an optional value, both of which you define.

      To add a tag to the fleet instances, click **Add Tags** and enter a tag key and value. You can add up to 47 tags to a fleet instance.

      To delete an existing tag, click the **cross** icon next to the tag.

      By default, this option is disabled and the EC2 fleet instance is launched in the same VPC and subnet of the Amazon EC2 instances that are protected.
9. Under **Recovery Options**, edit the default settings. Click the **edit** icon to edit the settings:
    a. **Rename**: Add a Prefix and/or Suffix to the recovered Amazon EC2 instances.

    b. **Power State**: To ensure that the restored EC2 instances are functioning properly, the EC2 instances are powered on by default after the restoration process. However, you can enable the **Power Off** option if you prefer to keep the VMs powered off once Cohesity verifies their functionality. If you are

recovering VMs from a Storage Domain with Cloud Tier enabled, Cohesity recommends keeping the VMs powered off.

c. **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.

d. **Task Name**: Change the default name of the recovery task.

10. Click **Recover**.

Recover VMs with Disks > 1 TB from AWS Native API Snapshots

When recovering from backups that were captured using AWS native snapshotting and the disks being restored are greater than 1 TB, the Cohesity cluster uses Proxy VMs to upload the data to the AWS cloud similar to how the Cohesity cluster creates Proxy VMs to download the data from the AWS Cloud to an on-premises Cohesity cluster or Cohesity Cloud Edition Cluster.

Note the following requirements for supporting the recovery of VMs with disks over 1 TB:

- The VPC where the VM will be recovered to, must be reachable from the Cohesity cluster through a VPN. For more information, see VPN Connections.
- Port 50051 must be open in the network security group where the VMs will be recovered to.

Disaster Recovery of VMs Using AWS Cohesity Cloud Edition

To provide disaster recovery failover and failback, Cohesity supports the Cloud Migration of VMs between hypervisor servers (VMware vCenter), cloud services (AWS and Azure) and Cohesity Cloud Edition Cohesity clusters. For Hyper-V, Cohesity supports only failover to Azure.

Disaster recovery with Cohesity Cloud Edition consists of the following two use cases:

- Back Up VMs to On-Prem, Replicate and Failover to AWS Cloud
- Back Up AWS Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server

Cohesity Import and Export for AWS

Cohesity supports an internal import and export feature that is fast and cost-effective. The number of temporary fleet instances created in the user account is limited by limiting the number of import tasks running at a time. Hence, even if other workloads are not present in the cluster, the number of import tasks will remain the same. For details on setting the number of import tasks, contact Cohesity Support.

Supported Softwares

Cohesity supports import and export from the following software versions:

- CentOS/RHEL 6, 7, 8
- Ubuntu 14, 16, 18
- Windows Server 2012, 2016, 2019

Limitations

- Import of VMs with Logical Volume Manager (LVM) on root volume spanning across multiple disks is not supported.
- This feature is not supported for:
  - Agent-based backups
  - Full CloudSpin

Back Up VMs to On-Prem, Replicate and Failover to AWS Cloud

To provide disaster recovery failover and failback, Cohesity supports the Cloud Migration of VMs between hypervisor servers (VMware vCenter), cloud services (AWS and Azure) and Cohesity Cloud Edition Cohesity clusters. For Hyper-V, Cohesity supports only failover to AWS. Cohesity Cohesity Cloud Edition supports the replication of VMs captured on an on-premises Cohesity cluster to Cohesity Cloud Edition as shown by the example figure below. Starting from the left in the following figure, the use case shows how Protection Groups running on the on-premises Cohesity cluster captures Snapshots of VMs and stores these Snapshots on the on-premises Cluster. The Protection Policies associated with these Jobs, are configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity Cloud EditionCohesity cluster.

From the Snapshots stored on the Cohesity Cloud Edition, you can failover VMs. The failover process creates new VMs running in the cloud service (such as AWS or Azure). As part of the Failover process, the VMs are converted and spun up in the cloud service.

> **Note:** Multi-Disk Windows and Linux EC2 instances are supported to failover to AWS and failback from AWS.

Cohesity supports the failover of VMware VMs with UEFI boot mode enabled to AWS.

Failover to AWS supports the following UEFI boot mode enabled operating systems:

- Windows Server 2012, 2016, 2019
- Ubuntu 16, 18

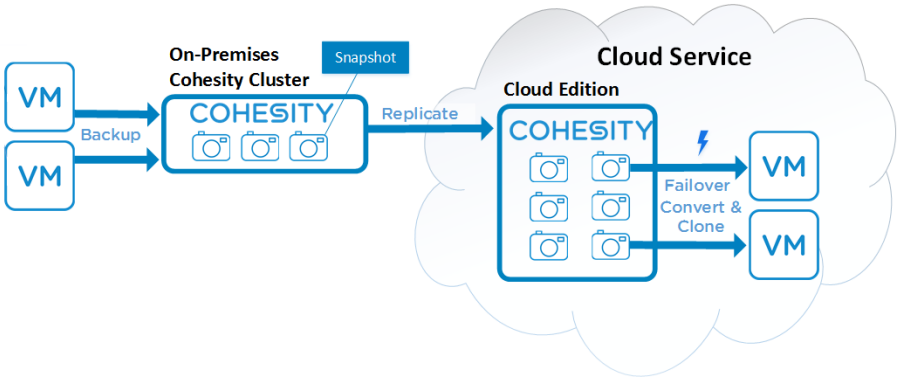

> **Important:** To support the failback (recovery) of these VMs back to the original Hypervisor server (such as VMware vCenter), Cohesity agents must be installed on the Windows VMs prior to backing up the snapshots on the on-premises Cohesity cluster. Failback is described in Back Up AWS Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server

Verify the following requirements prior to backing up the VMs:

- If you are planning to perform failback (recovery) of the backed-up VMs back to the Hypervisor server from the Cloud Service, then:
  - Before registering the VMs of the Hypervisor server with the on-premises Cohesity cluster, you must install Cohesity agent on the VMs if you are planning to back up the cloud VM back to the Cohesity cluster using the **Cohesity Protection Service** Protection Group type.
  - You do not need to install the Cohesity agent on the VMs if you are planning to back up the cloud VM back to the Cohesity cluster using native snapshots.

    > **Note:** You must install the Cohesity agent on the Windows VMs before registering the VMs of the Hypervisor server with the on-premises Cohesity cluster. The Cohesity agent is required to bring all the disks on the VMs online.

- Failover and Failback of VMware VM to AWS supports the following operating systems:
  - CentOS 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit
  - RHEL 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit
  - SUSE Linux Enterprise Server 11 64 bit, 12 64 bit
  - Ubuntu 14.x 64-bit, 16 64-bit, 18 64-bit, 20 64-bit

  - OEL 5.11 - 5.8 64-bit, 6.x 64-bit, 7.0 - 7.9 64-bit
  - Debian 6.0.0-6.0.8 64 bit, 7.0.0-7.8.0 64 bit, 8.0.0 64 bit
  - Fedora Server 19-21 64 bit
  - Windows 10 Desktop Edition, Microsoft Windows Server 2012, Microsoft Windows Server 2016, and Microsoft Windows Server 2019, Windows Server 2022.
- Cloning VMware VMs to an AWS Cloud has the following requirements:
  - Cloning snapshots to an AWS Cloud is supported for VMware VMs with specific Linux operating systems. For information on supported operating systems, see Supported Software
  - Cloning Snapshots to an AWS Cloud is supported for VMware VMs with the Windows operating system.
  - The disk size of the VMware VMs must be a multiple of 1 MB.
  - VMs that have one to ten disks can be cloned.
  - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

    ```
    UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
    defaults,nofail 0 0
    ```

    For detailed instructions, see Edit /etc/fstab to use UUID.

    For AWS, be aware of the following requirements:

    - In addition, the fstab entry for **CentOS 7.4**, **RHEL 7.4**, **RHEL 6.7**, **SUSE 11** and **Ubuntu 16** operating systems must specify the `nofail` option as shown in the following example:

      ```
      UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
      defaults,nofail 0 0
      ```

    - In addition, the fstab entry for the **Ubuntu 14** operating system must specify the `nobootwait` option as shown in the following example:

      ```
      UUID=2b262696-41a6-47c9-b173-d02717ee0e10 /mnt/expvol ext4
      defaults,nobootwait 0 0
      ```
- For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.

To back up VMs to an on-premises Cluster, replicate and failover to Cohesity Cloud Edition:

1. Register the Hypervisor server with the on-premises Cohesity cluster
2. Set up replication between the on-premises Cohesity cluster and the Cohesity Cloud Edition Cohesity cluster. For instructions, see Replication and Remote Access Setup.
3. On the on-premises Cohesity cluster, create a Policy that replicates to a Cohesity Cloud Edition Cohesity cluster as shown in the top figure. For instructions, see Create or Edit a Standard Policy.
4. On the on-premises Cohesity cluster, create a Virtual Server Protection Group (such as `Win2k12`) that has **Cloud Migration** enabled. For detailed instructions, see Add or Edit a Protection Group for Virtual Servers. Associate the Policy created in the previous step with this Protection Group.

   When selecting Objects to be protected by a Protection Group, you can only select VMs with VMware tools installed, Remote Desktop enabled, and for Windows VM Cohesity agents also must be installed.

   When you create a Protection Group on the on-premises Cohesity cluster that replicates to a Cohesity Cloud Edition Cohesity cluster, after the first Snapshot is replicated to the Cohesity Cloud Edition Cluster, an Inactive copy of the Protection Group is created on the Cohesity Cloud Edition Cluster that received the Snapshot. The new Inactive Protection Group on the Cohesity Cloud Edition Cohesity cluster has the same name as the Protection Group on the on-premises Cohesity cluster, for example: `Win2k12`. The VM must be powered on and connected to the network when the Protection Group is running.
5. You can failover from the Snapshots that were replicated to the Cohesity Cloud Edition Cluster. The failover process creates new VMs running in the cloud service (such as AWS or Azure). As part of the Failover process, the VMs are converted and spun up in the cloud service.

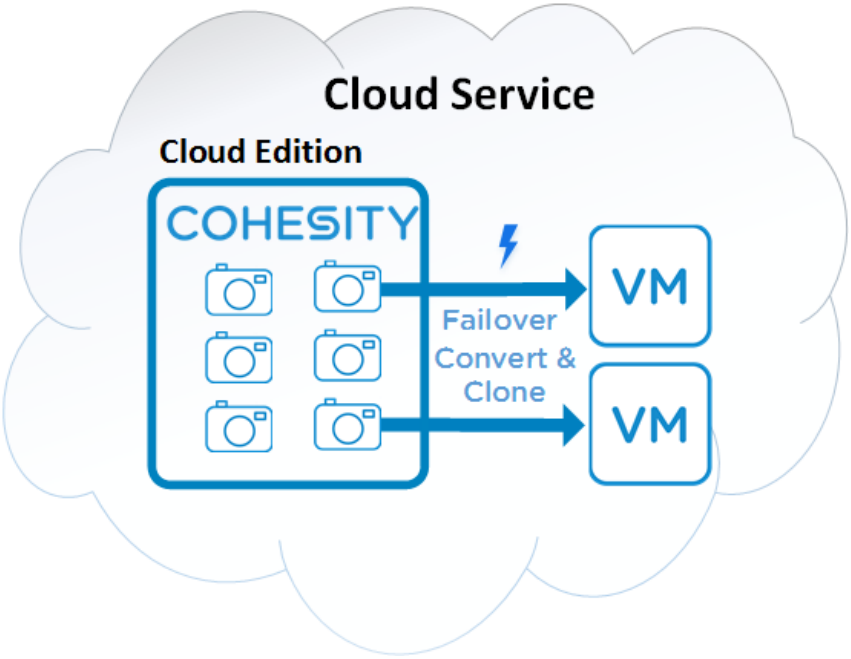

To failover from the Snapshots:

a. On the Cohesity Cloud Edition Cohesity cluster that contains the replicated Snapshots, select **Data Protection > Protection** and browse for the Inactive Protection Group that was automatically created.

b. Click on the inactive Protection Group name.

c. Click the action icon (⋮) next to the Protection Group and select **Failover**.

d. From the **Failover to Source** drop-down, select the name of the registered Cloud Source where the VMs will be cloned to.

e. Click **Continue to Clone**.

f. Follow the instructions provided in Set Clone Options. You can also add custom tags to all the resources created during failover.

A clone task is created and the VM is created in the AWS cloud service and is powered up if that option is selected.

6. Follow the steps in the next part of the Disaster Recovery workflow which is documented in Back Up AWS Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server.

You can also use CloudSpin (test/dev clone) workflow to failover individual VMs within the Protection Group instead of failing over all VMs in the Protection Group. For more information, see Test & Dev Using CloudSpin.

Back Up AWS Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server

To provide disaster recovery failover and failback, Cohesity supports the Cloud Migration of VMs between hypervisor servers (VMware vCenter), cloud services (AWS and Azure) and Cohesity Cloud Edition Cohesity clusters.

> **Note:**
> - For Hyper-V, Cohesity supports only failover to AWS. Failback is not supported.
> - Failover of VMs born in the cloud to an on-premises VMware vCenter is not supported. This topic describes the failback of VMs that were originally backed up from an on-premises VMware vCenter server and subsequently failed over to AWS using the process described in Back Up VMs to On-Prem, Replicate and Failover to Cloud.

This topic describes the second use case of disaster recovery using Cloud Edition—how to backup cloud VMs to the Cloud Edition cluster, replicate to the on-premise cluster, and then failback to the Hypervisor server.

For information on how the VMs are backed up in the AWS cloud service, see the first process of disaster recovery using Cloud Edition documented in Back Up VMs to On-Prem, Replicate and Failover to AWS Cloud.

Cohesity supports the failback of VMware VMs with UEFI boot mode enabled from AWS. Failback from AWS supports the following UEFI boot mode enabled operating systems:

- Windows Server 2012, 2016, 2019
- Ubuntu 16, 18

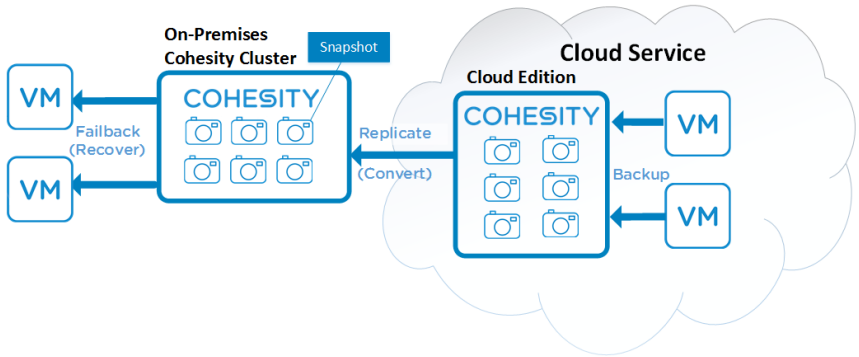

As illustrated in the figure above, on the Cohesity Cloud Edition running in the Cloud Service, a Protection Group using native snapshots or Cohesity Protection Service backs up the cloud VMs (as indicated by the backup arrow in the figure). The policy associated with the Protection Group specifies that these Snapshots are replicated to the on-premises Cohesity cluster (as indicated by the Replicate arrow). From those Snapshots stored on the on-premises Cohesity cluster, you can failback or recover the VMs back to their original

location or a new location in your Hypervisor server such as a vCenter Server (as indicated by the Failback (Recover) arrow).

To back Up Cloud VMs to Cohesity Cloud Edition, replicate to an on-premises Cohesity cluster and failback to a hypervisor server:

Verify the following requirements prior to failback:

- The VMs that are backed up to the Cloud Service must have VMware Tools and Remote Desktop installed before the original Protection Group captures the Snapshots from the VMware vCenter Server as described in the Back Up VMs to On-Prem, Replicate and Failover to AWS Cloud.
- If you are backing up the Cloud VMs to the Cohesity Cohesity Cloud Edition Cluster using the Cohesity **Protection Service** Protection Group type, the VMs must have the Cohesity agent installed inside the guest operating system. The agent backs up the VM from inside the guest OS (guest-level backup). This is unlike the Native Snapshots type, which operates at the hypervisor level and does not require an agent on the VMs. For installation instructions, see Deploy Cohesity Agent on Virtual Servers.
- For AWS cloud source:
  - Configure the VPC and tunnels so that the source VMs are reachable using private IP addresses.
  - Cohesity fleet instances must have ports 50051 and 22 open, so that the Cohesity cluster can communicate with fleet instances. Cohesity creates one network security group in each VPC, and attaches to the fleet instance.

    > **Note:** By default, the inbound rules at the security group will be configured with the CIDR range of the network where the Cohesity cluster is configured. Contact Cohesity Support if you want to configure the security group with a different CIDR range.

  - If you are backing up from an AWS Cloud Source, you must use a AWS Cloud Source that has been granted adequate permissions to back up.

  For more on the native snapshot process for creating the Cohesity snapshot, see AWS Backup Use Cases.

Steps:

1. If not already configured, set up replication between the Cohesity Cloud Edition Cluster and the on-premises Cohesity cluster. For instructions, see Replication and Remote Access Setup.
2. On the Cohesity Cloud Edition Cluster, create a Policy that replicates from the Cohesity Cloud Edition Cluster to the on-premises Cohesity cluster as shown in the

figure. For instructions, see Create or Edit a Standard Policy.

3. On the Cohesity Cloud Edition Cluster, refresh the Object hierarchy for the Source of the VMs that will be backed up:
   a. Select **Data Protection > Sources**.
   b. Click the action icon (⋮) next to the Source and select **Refresh**.
4. On the Cohesity Cloud Edition Cluster, create a Virtual Server Protection Group (such as `BackToPrem`) with Protection Group Type option as **Native Snapshots** or **Cohesity Protection Service**. Associate the Policy created in the previous step with this Protection Group. For detailed instructions, see Add or Edit a Protection Group for Virtual Servers. If **Cohesity Protection Service** is selected as the Protection Group type, then the VMs that are backed up to the Cloud Service (failovered VMs) must have a Cohesity agent installed.

   When selecting Windows VM (Objects) to failover, you can only select Snapshots of VMs that have the Cohesity agent installed when they were created with the Back Up VMs to On-Prem, Replicate and Failover to AWS Cloud process. The VM should be powered on when the Protection Group is running.

   When you create a Protection Group on the Cohesity Cloud Edition Cluster that replicates to an on-premises Cohesity cluster, after the first Snapshot is replicated to the on-premises Cluster, an Inactive copy of the Job is created on the on-premises Cluster that received the Snapshot. The new Inactive Job on the on-premises Cohesity cluster has the same name as the Protection Group on the Cohesity Cloud Edition Cluster, for example: `BackToPrem`.
5. You can failback (recover) the VMs to the original Hypervisor server or an alternate Hypervisor server:
   a. On the on-premises Cohesity cluster that contains the replicated Snapshots, select **Data Protection > Protection** and browse for the Inactive Job that was automatically created.

   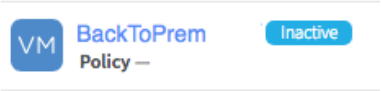

   b. Click on the inactive Job name.
   c. Browse for a protection run that captured a Snapshot that you want to recover.
   d. Click the action icon (⋮) next to the Job and select **Recover VM**.
   e. Follow the instructions provided in Recover VMs.

   The VM is created in the Hypervisor server and is powered up if that option is selected. As part of the convert and clone process, the VMware Tools will be uninstalled. You must reinstall the VMware Tools.

> **Important:** Do not change the size of the partition or the disk of the recovered VM.

Disaster Recovery of Views in AWS Cohesity Cloud Edition

Cohesity Cohesity Cloud Edition supports the disaster recovery of Views as shown by the following example figures.

Views can be used to store data such as files and directories. For example, other backup products can store data on a view that is hosted on the Cohesity cluster. To access a view from your system, you must mount the view.

On the left of the first figure, shows systems accessing a View hosted on an on-premises Cohesity cluster using the NFS or SMB protocols. A Protection Group running on the on-premises Cohesity cluster captures Snapshots of the View and stores these Snapshots on the Cluster. The Protection Policy associated with this Job, is configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity cluster Cohesity Cloud Edition. From these Snapshots stored in the Cohesity Cloud Edition, you can clone a new View that can be accessed by systems using the NFS or SMB protocols as shown on the right of the first figure.

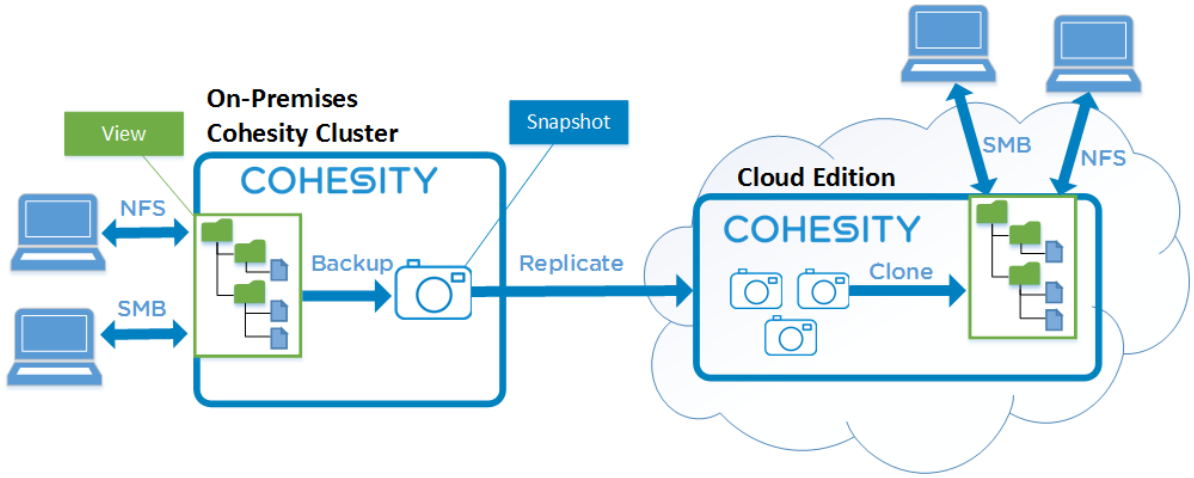

You can backup Views hosted on the Cohesity Cloud Edition as shown on the right of the following figure. A Protection Group captures a Snapshot of a View and stores the Snapshot on the Cohesity cluster Cohesity Cloud Edition. The Protection Policy associated with this Job, is configured to replicate (or copy) Snapshots from the Cohesity cluster Cohesity Cloud Edition to the on-premises Cohesity cluster. From these View Snapshots stored in the on-premises Cohesity cluster, you clone a View and make it available for mounting as shown on the left of the following figure.

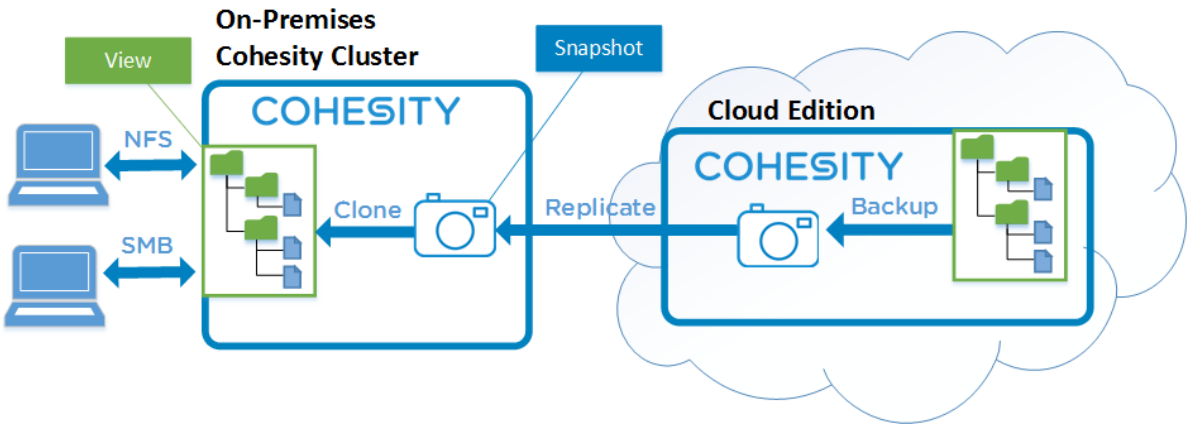
On-Premises
Cohesity Cluster
View
Snapshot
COHESITY
NFS
SMB
Clone
Replicate
Cloud Edition
COHESITY
Backup

AWS Clone

This chapter includes the following topics:

Clone VMs from On-Premises to AWS Cloud Service (On Demand CloudSpin)

Cohesity supports the conversion of snapshots of VMs from an on-premises Cohesity cluster to an AWS Cloud service. After conversion, the VM runs in the Cloud service.

> **Note:** This use case is not currently supported for the Google Cloud Platform.

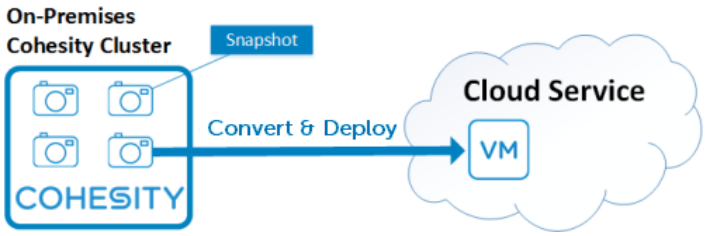

Prerequisites and Considerations

- For the list of VMs that supports CloudSpin to AWS, see Supported Software.
- Cohesity supports the CloudSpin of VMware VMs with UEFI boot mode enabled.
- Cohesity supports both on-demand and policy-based CloudSpin for snapshots stored on the Cohesity cluster (primary copy).
- Cohesity does not support on-demand and policy-based CloudSpin for archived or CloudRetrieved snapshots.
- Cohesity supports on-demand CloudSpin for snapshots that are replicated to another Cohesity cluster.
- Cloning VMware VMs to AWS
  - Cloning snapshots to an AWS Cloud is supported for VMware VMs with specific Linux operating systems. For information on supported operating systems, see Supported Software
  - Cloning Snapshots to an AWS Cloud is supported for VMware VMs with the Windows operating system.
  - The disk size of the VMware VMs must be a multiple of 1 MB.
  - VMs that have one to ten disks can be cloned.
  - When cloning a VMware VM that has 1 TB or greater disks, Cohesity launches a temporary EC2 instance and copies data from the Cohesity cluster to the instance. The instance can use the same VPC and subnet that's used for the clone, or you can provide a different VPC and subnet, but they must be

reachable from the cluster.

- When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

  ```
  UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
  defaults,nofail 0 0
  ```

  For detailed instructions, see Edit /etc/fstab to use UUID.

  For AWS, be aware of the following requirements:

  - In addition, the fstab entry for **CentOS 7.4**, **RHEL 7.4**, **RHEL 6.7**, **SUSE 11** and **Ubuntu 16** operating systems must specify the `nofail` option as shown in the following example:

    ```
    UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
    defaults,nofail 0 0
    ```

  - In addition, the fstab entry for the **Ubuntu 14** operating system must specify the `nobootwait` option as shown in the following example:

    ```
    UUID=2b262696-41a6-47c9-b173-d02717ee0e10 /mnt/expvol ext4
    defaults,nobootwait 0 0
    ```

- When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.

  > **Note:**
  > **For Convert and Deploy**: It is mandatory to install the Cohesity agent on the VMs. Windows does not show the data disks if the Cohesity agent is not installed, irrespective of the SAN/Disk Policy.
  >
  > **For Disaster Recovery**: It is mandatory to install the Cohesity agent on multi-disk Windows VMs. Windows does not show the data disks if the Cohesity agent is not installed, irrespective of the SAN/Disk Policy. In addition, the Protection Group capturing the Snapshot must have **Cloud Migration** enabled. For detailed instructions, see Add or Edit a Protection Group for Virtual Servers.

- For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
- Cohesity does not support CloudSpin of VMs from VMware Cloud (VMC) to AWS.

Convert and Deploy a VM from On-Premises to AWS

Verify that the above prerequisites are met prior to backing up the VMs.

1. Create a Protection Group to capture Snapshots of VMs. For more information, see Add or Edit a Protection Group for Virtual Servers.
2. Register the Cloud service that the VMware VMs are cloned **to**. For detailed instructions, see either Register an Azure Cloud Source at Subscription Level or Register AWS Source.
3. Log in to the Cohesity Dashboard of the Cohesity cluster that contains the Snapshots to clone from. For example, log in to the on-premises Cohesity cluster.
4. To clone and convert the Virtual Server (VM), follow the instructions provided in Clone VMs.
5. After the Clone task is completed, log in to the new VM in the Cloud service.

   For AWS:

   a. Log in to the Amazon AWS console using your AWS account.
   b. In the top left of the menu bar, select **Services**.
   c. Select **Compute** > **EC2**.
   d. Select **INSTANCES** > **Instances**.
   e. Browse for the new Virtual Machine name.
   f. Copy the new IP address from the **Private IPs** field.
   g. Log in to the VM using the IP address. Paste the IP address in the browser.
6. Multiple disks are automatically mounted on the Windows VM, if the following conditions are true before the Protection Group captures the Snapshot of a Windows VM:

   - A Cohesity agent is installed on the VM.
   - VMware tools are installed on the VM.
   - VM is running

   If any of the conditions were not true when the Job captured the Snapshot, the `C:` drive is mounted but you must mount the other drives such as `D:` as described here:

   a. If Remote Desktop is installed on the new cloned VM, log in to the new VM, start Remote Desktop Connection, enter the IP of the new cloned VM.
   b. In the new VM, launch the Windows Disk Management.

c. Right click on a disk that is not mounted and select **Online**. Repeat for all other offline disks.

Tear Down a Launched VM from Cohesity Dashboard

1. In the Cohesity Dashboard, select **Test & Dev**.
2. Click the actions menu (⋮) next to the clone task you want to tear down and select **Tear Down**.

AWS Replication

This chapter includes the following topics:

Replicating VMs in AWS Cloud Edition

Cohesity Cohesity Cloud Edition supports the replication of VMs captured on an on-premises Cohesity cluster to Cohesity Cloud Edition as shown by the example figure below.

Starting from the left in the following figure, the use case shows how Protection Groups running on the on-premises Cohesity cluster captures Snapshots of VMs and stores these Snapshots on the on-premises Cluster. The Protection Policies associated with these Jobs, are configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity Cloud EditionCohesity cluster.

Starting from the left in the following figure, the bottom workflow shows how the Cohesity Cloud Edition can directly back up a VM that is not in the Cloud. For example a VM in a vCenter server can be backed up directly to Cohesity Cloud Edition running in a Cloud service such as AWS or Azure.

You can recover VMs from the Snapshots stored in the Cohesity Cloud Edition running in the Cloud service (such as AWS).

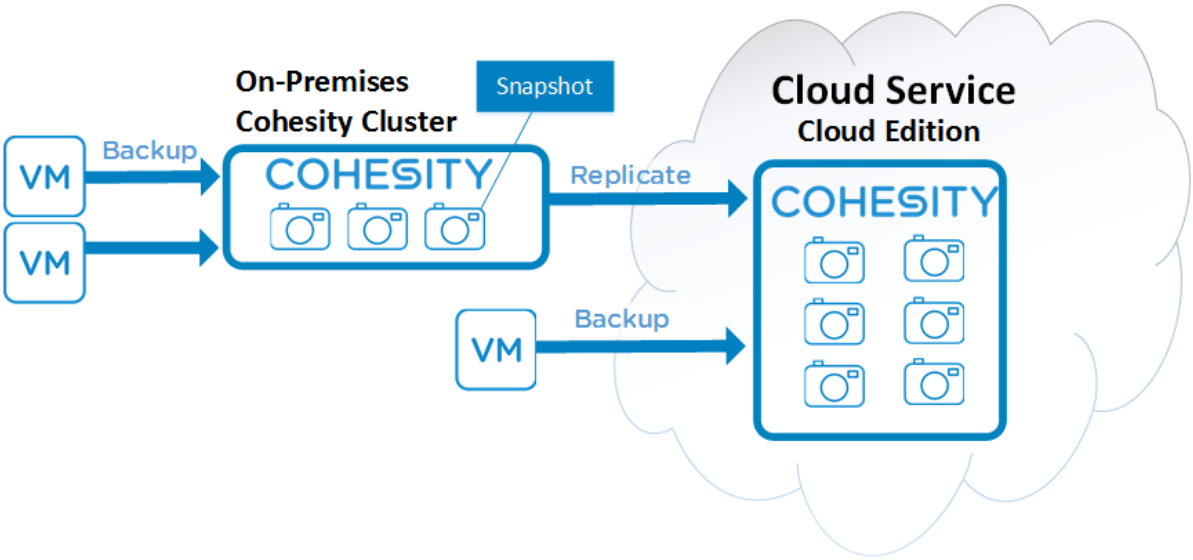

Replicate and Clone VMs in AWS Cloud Edition

Cohesity Cohesity Cloud Edition supports Test/Dev (cloning) as shown by the following example figure.

Starting from the left in the following figure, the top workflow shows how Protection Groups running on the on-premises Cohesity cluster captures Snapshots of the VMs and stores these Snapshots on the on-premises Cluster. The Protection Policies associated with these Jobs, are configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity clusterCohesity Cloud Edition.

Starting from the left in the following figure, the bottom workflow shows how the Cohesity Cloud Edition can directly back up a VM that is not in the Cloud. For example a VM in a vCenter server can be backed up directly to Cohesity Cloud Edition running in a Azure Cloud service.

From the Snapshots stored in the Cohesity Cloud Edition on the Azure Cloud service, you can spin up test/dev environments using cloning.

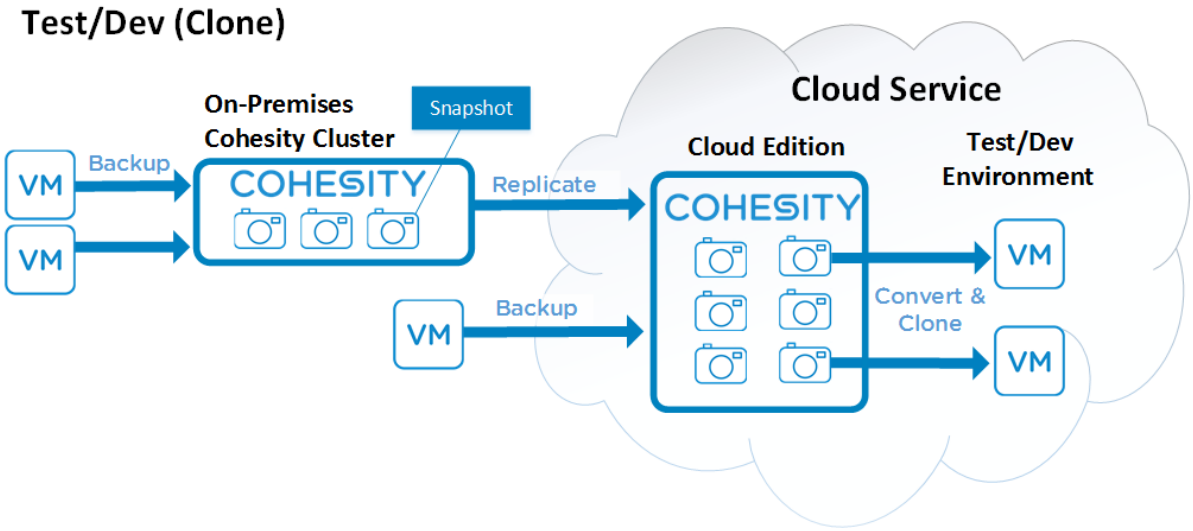

To convert and clone VMs to a Cloud service:

1. Verify the following requirements prior to backing up the VMs:
    - For the list of VMs that supports CloudSpin to AWS, see Supported Software.
    - All the VMware vSphere versions supported by Cohesity for backup is supported for cloning to AWS. For the list of supported VMware vSphere versions, see Supported Software.

    Cloning Snapshots to an Azure Cloud is supported for VMware and Hyper-V VMs with the following operating systems:

    - Linux CentOS 7 on LVM
        - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
        - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

          ```
          UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
          defaults,nofail 0 0
          ```

          For detailed instructions, see Edit /etc/fstab to use UUID.
    - Windows

- You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
- For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
- When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. The Cohesity agents installed must be of the same version as the Cohesity Cluster version. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.
- Due to a Windows issue in builds 14393 and build 15063, a deployed Windows 10 or Windows 2016 VM may not be connected to the network. Resolve this by installing fix KB4073562 for build 14393 or KB4016240 for build 15063 or by updating Windows. For more information, see Cannot connect remotely to a Windows 10 or Windows Server 2016 VM in Azure because of netvsc.sys.
- Make the disks online by default using the following SAN policy commands:
  - `diskpart`
  - `san policy=OnlineAll`

  For more information, see the SAN Policy documentation.

- Red Hat Enterprise Linux
  - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
  - For the list of supported RHEL versions, see Supported Software.
  - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

    ```
    UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
    defaults,nofail 0 0
    ```

    For detailed instructions, see Edit /etc/fstab to use UUID.
- The disk size of the VMware VMs must be a multiple of 1 GB.
- Cohesity does not support the CloudSpin of VMs whose OS Disk is formatted with XFS version 5 to Azure. This limitation applies to both Hyper-V and VMware VMs, regardless of their operating system type.

- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to Azure.

Cloning Hyper-V VMs to Azure

- VMs that have one to five disks can be cloned. Each disk must be smaller than 2 TB.
- SUSE Linux VMs are *not* supported.
- All the Hyper-V versions supported by Cohesity for backup is supported for cloning to Azure. For the list of supported Hyper-V versions, see Supported Software.
- Generation 2 VMs are not supported by all VM sizes in Azure. For information on Generation 2 supported VM series, see Microsoft document.
- Cloning of Hyper-V VMs to Azure Stack Hub is not supported.

Cloning VMware VMs to an AWS Cloud has the following requirements:

- Cloning snapshots to an AWS Cloud is supported for VMware VMs with specific Linux operating systems. For information on supported operating systems, see Supported Software
- Cloning Snapshots to an AWS Cloud is supported for VMware VMs with the Windows operating system.
- The disk size of the VMware VMs must be a multiple of 1 MB.
- VMs that have one to ten disks can be cloned.
- When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

  ```
  UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
  defaults,nofail 0 0
  ```

  For detailed instructions, see Edit /etc/fstab to use UUID.

  For AWS, be aware of the following requirements:

  - In addition, the fstab entry for **CentOS 7.4**, **RHEL 7.4**, **RHEL 6.7**, **SUSE 11** and **Ubuntu 16** operating systems must specify the `nofail` option as shown in the following example:

    ```
    UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
    defaults,nofail 0 0
    ```

  - In addition, the fstab entry for the **Ubuntu 14** operating system must specify the `nobootwait` option as shown in the following example:

    ```
    UUID=2b262696-41a6-47c9-b173-d02717ee0e10 /mnt/expvol ext4
    defaults,nobootwait 0 0
    ```

- When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.

  > **Note:**
  > **For Convert and Deploy**: It is mandatory to install the Cohesity agent on the VMs. Windows does not show the data disks if the Cohesity agent is not installed, irrespective of the SAN/Disk Policy.
  >
  > **For Disaster Recovery**: It is mandatory to install the Cohesity agent on multi-disk Windows VMs. Windows does not show the data disks if the Cohesity agent is not installed, irrespective of the SAN/Disk Policy. In addition, the Protection Group capturing the Snapshot must have **Cloud Migration** enabled. For detailed instructions, see Add or Edit a Protection Group for Virtual Servers.

- For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.

2. Create a Protection Group to capture Snapshots of VMs. For more information, see Add or Edit a Protection Group for Virtual Servers. For example, you could create a Protection Group on a Cohesity cluster running in Azure as shown in the bottom workflow in the figure above.
3. Register the Cloud service that the VMware VMs are cloned **to**. For detailed instructions, see either Register an Azure Cloud Source at Subscription Level or Register AWS Source.
4. Log in to the Cohesity Dashboard of the Cohesity cluster that contains the Snapshots to clone from. For example, log in to the Cohesity cluster running in Azure or AWS.
5. Clone the VM in the Cloud service such as Azure. For detailed instructions, see Clone VMs.
6. After the Clone task is completed, log in to the new VM in the Cloud service.

   For Azure:

   a. Log in to Microsoft Azure Portal.
   b. Search for the new Virtual Machine name.
   c. Click on the VM name.
   d. Under the **Subscription name**, copy the IP address of the new VM.

e. Log in to the VM using the IP address. If Remote Desktop was installed in the VM that was cloned, you can log in using the Remote Desktop.

For AWS:

a. Log in to the Amazon AWS console using your AWS account.
b. In the top left of the menu bar, select **Services**.
c. Select **Compute** > **EC2**.
d. Select **INSTANCES** > **Instances**.
e. Browse for the new Virtual Machine name.
f. Copy the new IP address from the **Private IPs** field.
g. Log in to the VM using the IP address. Paste the IP address in the browser.

7. Multiple disks are automatically mounted on the Windows VM, if the following conditions are true before the Protection Group captures the Snapshot of a Windows VM:

   - A Cohesity agent is installed on the VM.
   - VMware tools are installed on the VM.
   - VM is running

   If any of the conditions were not true when the Job captured the Snapshot, the `C:` drive is mounted but you must mount the other drives such as `D:` as described here:

   a. If Remote Desktop is installed on the new cloned VM, log in to the new VM, start Remote Desktop Connection, enter the IP of the new cloned VM.
   b. In the new VM, launch the Windows Disk Management.
   c. Right click on a disk that is not mounted and select **Online**. Repeat for all other offline disks.

Clone VMs to AWS Cloud Using CloudSpin

On a Cohesity cluster with a Protection Group that captures Snapshots of VMs with a CloudSpin schedule, the cluster clones VMs to an AWS cloud service. You can then launch the VMs to a running state.

What happens during the CloudSpin process is dependent on the cloud service type:

Full	Incremental
When you clone to an AWS Cloud using CloudSpin, the Cohesity cluster copies the Snapshot to an S3 bucket in the specified AWS account and then converts the snapshot from VMDK format to an AMI format using the AWS import service.	On the first CloudSpin run, the Cohesity cluster copies the full VM snapshot data to EBS volumes. The number of EBS volumes corresponds to the number of backed-up VM disks. Once copying is done, Cohesity uses AWS import service to generate an AMI from the EBS volume corresponding to the root volume and generates the EBS snapshots for all other volumes, if any. For subsequent runs, only incremental changes are copied between the snapshots. Incremental CloudSpin to AWS requires a Fleet Instance. The cluster communicates with the Fleet Instance using private IP. Therefore, the on-prem cluster must be on a VPN to AWS VPC.

Cohesity supports CloudSpin on both on-premises Cohesity cluster and Cohesity Cloud Edition Cohesity clusters:

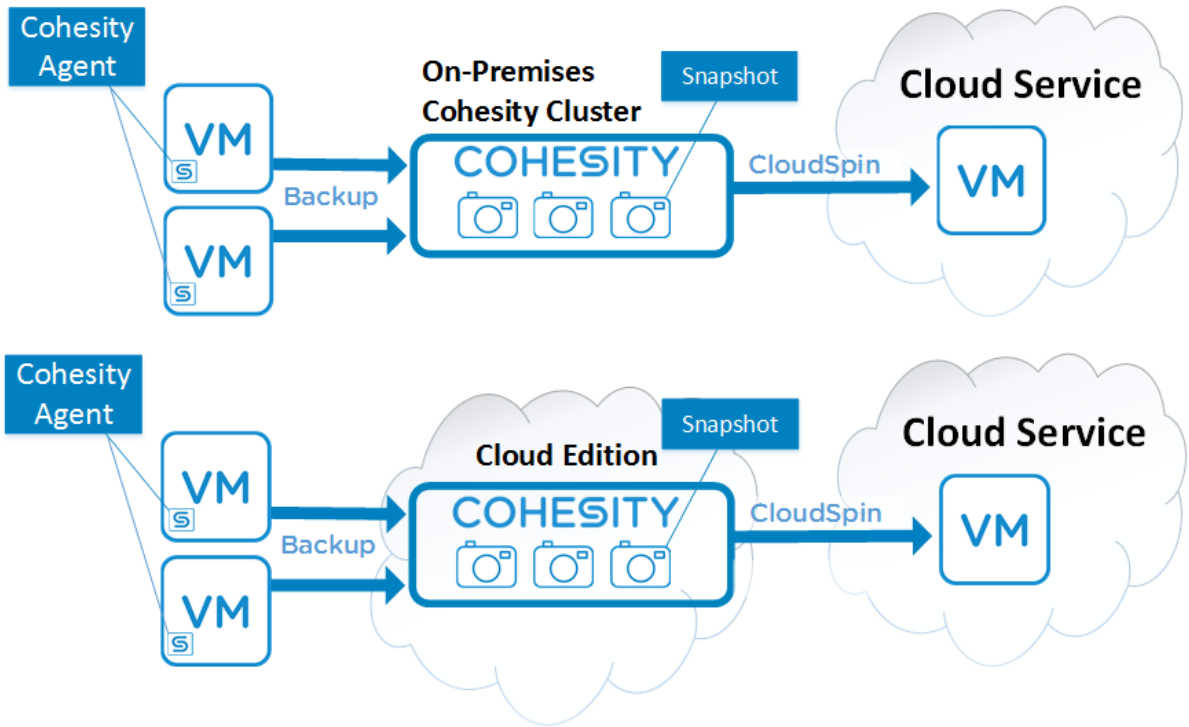

To use CloudSpin to clone Snapshots to AWS:

1. Verify the following requirements prior to backing up the VMs:
    - For the list of VMs that supports CloudSpin to AWS, see Supported Software.
    - All the VMware vSphere versions supported by Cohesity for backup is supported for cloning to AWS. For the list of supported VMware vSphere versions, see Supported Software.
    - Cloning VMware VMs to an AWS Cloud has the following requirements:
        - Cloning snapshots to an AWS Cloud is supported for VMware VMs with specific Linux operating systems. For information on supported operating systems, see Supported Software
        - Cloning Snapshots to an AWS Cloud is supported for VMware VMs with the Windows operating system.
        - The disk size of the VMware VMs must be a multiple of 1 MB.
        - VMs that have one to ten disks can be cloned.
        - Due to an AWS limitation, performing a CloudSpin to an instance type that has an ENA is not supported. For more information, see Enhanced Networking on Linux.
        - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

          ```
          UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4 defaults,nofail 0 0
          ```

          For detailed instructions, see Edit /etc/fstab to use UUID.

          For AWS, be aware of the following requirements:

            - In addition, the fstab entry for **CentOS 7.4**, **RHEL 7.4**, **RHEL 6.7**, **SUSE 11** and **Ubuntu 16** operating systems must specify the `nofail` option as shown in the following example:

              ```
              UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4 defaults,nofail 0 0
              ```

            - In addition, the fstab entry for the **Ubuntu 14** operating system must specify the `nobootwait` option as shown in the following example:

              ```
              UUID=2b262696-41a6-47c9-b173-d02717ee0e10 /mnt/expvol ext4 defaults,nobootwait 0 0
              ```

        - When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must

be installed on the VM and the VM must be running. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.

- For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.

2. Register a Cloud Source. Select **Data Protection > Sources** and from the **Register Source** drop-down, select **AWS**. A cloud source is a cloud service instance where the snapshots are converted and copied **to**. For details, see Register AWS Source.
3. For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
4. Create a Policy with a CloudSpin schedule that converts and copies Snapshots to the Cloud Source registered in the previous step. For detailed instructions, see Create or Edit a Standard Policy. Protection Groups that use this policy will automatically convert and copy Snapshots backed up by the Protection Group to a Cloud Source using the specified CloudSpin schedule.

   For AWS:

Full	Incremental
In Full CloudSpin, the AMI and snapshots are created in the specified Region.	Specify the Region, VPC, and Subnet for an incremental CloudSpin. The AMI and snapshots are created in the specified Region. VPC and Subnet are required for the following reason: Incremental CloudSpin requires connectivity between the Cohesity cluster and AWS through VPC and Subnet - During incremental CloudSpin, Cohesity launches a temporary EC2 instance in the provided VPC and Subnet. The temporary instance is used to copy data from Cohesity cluster to AWS. An instance launched in this subnet must be reachable from the Cohesity cluster.

5. Create a Protection Group that uses the Policy created in the previous step.

   After the Protection Group run and the CloudSpin Copy Task are successfully completed, the VM can be launched.
6. Launch the VM created in the Cloud Service:

   Launch VM from the Cohesity Dashboard:

   a. Select **Data Protection > Protection**.
   b. Click the Job that backed up the VM to launch in the Cloud.

c. Browse for a protection run that contains a Snapshot to convert to a VM. The protection run must have the Status of **Success** and the associated CloudSpin Copy task must be successful .

d. Click the date link of the protection run in the Start Time column to view the protection run details.

e. Click the **CloudSpin Task** tab.

f. To launch, complete one of the following actions:

- Click the actions menu (⋮) and select **Launch Instance**.
- Click **Launch All Instances**.

To tear down a launched VM, see Tear Down a Launched VM from Cohesity Dashboard.

Launch VM from AWS Console:

a. Log in to the Amazon AWS console and select **EC2**. The **EC2 Dashboard** displays.

b. In the left frame, select **IMAGES > AMIs**.

c. In the search field, enter part of the VM name. The generated VM name consists of the Job name, the original name and the Snapshot time.

d. Select a VM and click **Launch**.

e. Enter the required settings and click **Review and Launch**.

To prevent a tear down of a cloned CloudSpin VM, see the How to prevent a tear down of a cloned CloudSpin VM KB article.

Tear Down a Launched VM from Cohesity Dashboard

- In the Cohesity Dashboard, select **Test & Dev**.
- For the clone task to tear down, click the actions menu (⋮) and select **Tear Down**.

Consideration

- AWS Security Token Service (AWS STS) must be enabled in the region where a VM is being Cloud-spun.
- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to AWS.
- Cohesity supports both on-demand and policy-based CloudSpin for snapshots stored on the Cohesity cluster (primary copy).
- Cohesity does not support on-demand and policy-based CloudSpin for archived or CloudRetrieved snapshots.
- Cohesity supports on-demand CloudSpin for snapshots that are replicated to another Cohesity cluster.

Edit /etc/fstab to use UUID

You must edit the `/etc/fstab` file and change the (non-LVM) entries to use UUIDs to identify the disk partitions. These changes must be made to the `/etc/fstab` file before the Protection Group (which captures the Snapshots to be cloned) is run.

For example, the following entry in the `/etc/fstab` table must be changed to use the UUID to identify the partition instead of the device file name of `/dev/sbd1`.

```
/dev/sdb1 /mnt/disk1    ext4    defaults        0 0
```

> **Note:**
> - Use labels or UUIDs instead of device names, as device names may change when other disks are added or removed.
> - Kernel name descriptors must not be used in configuration files as they are not persistent and may change during each boot. UUIDs or PARTUUIDs are preferred over labels.
> - Using only the device name in the absence of UUID or label might cause issues as these names are not unique. This can cause overwrites and eventual loss of an entire volume indexing itself.
> - Each volume must be uniquely identified by an identifier like UUID or label. **A missing unique identifier (UUID/label) can result in dummy volume names being used instead of actual volumes.**

You can determine the UUID by running the `blkid` command, for example:

```
# blkid /dev/sdb1
/dev/sdb1 UUID="12345cac-77bc-46b1-9f79-cb180e438c12" TYPE="ext4"
```

Edit the `/etc/fstab` as root and for this entry replace the device file name with the returned UUID, for example:

```
12345cac-77bc-46b1-9f79-cb180e438c12 /mnt/disk1 ext4    defaults        0 0
```

> **Note:** LVM volume entries should not be updated. LVM volumes start with the `/dev/mapper` prefix, for example: `/dev/mapper/centos-root` and `/dev/mapper/centos-swap`.

Repeat these steps for all entries in the `/etc/fstab` that are not LVM volume entries.

Amazon Elastic Block Store

Amazon Elastic Block Store (EBS) is a scalable, high-performance block storage service designed for use with Amazon EC2 instances. EBS volumes function as virtual hard disks and provide persistent storage independent of the lifecycle of the EC2 instance.

Cohesity uses AWS EBS Direct APIs to provide agentless, block-level backup and recovery for both standalone and attached EBS volumes.

> **Important:** The backup and recovery of Amazon EBS volume is supported only on Cohesity Cloud Edition on AWS.

Why Use EBS Volume Protection?

You can use EBS volume protection when you need to protect and recover storage volumes independently of the EC2 instances to which they are attached. Key benefits include:

- **Independent volume protection** – Protect EBS volumes without relying solely on EC2 instance backups.
- **Faster and more flexible recovery** – Recover individual volumes directly when only the volume data is required.
- **Avoid full instance recovery** – Restore a specific EBS volume without recovering the entire EC2 instance, reducing recovery time and operational overhead.
- **Support for standalone volumes** – Protect EBS volumes that are not attached to any EC2 instance.

Plan and Prepare for Amazon EBS Volume Protection

To protect Amazon EBS volume with Cohesity, you must register your AWS account as a source in the Cohesity cluster and perform backup and recovery operations. Ensure you meet the pre-requisites and understand the considerations before registering the AWS account. For information on registering AWS account, see Register AWS Source.

Supported Amazon EBS Volume Types

Cohesity supports protection of the following Amazon EBS volume types:

- General Purpose SSD (gp2, gp3)
- Provisioned IOPS SSD (io1, io2)
- Throughput Optimized HDD (st1)
- Cold HDD (sc1)

Required Permissions for EBS Volume Data Protection

For information on the permissions required to protect EBS volume, see Minimum Permissions for Individual Workloads and Workflows.

Ports Used for Communication

For information on port numbers required, seeAmazon Elastic Block Store.

Considerations for EBS Volume Data Protection

Review and understand the following limitations and known issues before you protect EBS volumes:

- If an EBS volume attached to an EC2 instance is protected through both EC2 instance protection and EBS volume protection, each protection is accounted for independently and contributes separately to protected-size calculations.
- Cohesity does not support:
  - Restoring an EBS volume to an alternate AWS account or AWS Region.
  - Restoring an EBS volume directly as an attached volume to an EC2 instance.
  - File-level recovery (FLR) from EBS volume backups.
  - EBS volume protection when AWS is registered at the organization level.
- Restores create a new EBS volume in a detached state.
- If required, you must manually attach the restored volume to an EC2 instance.
- AWS supports up to five pending snapshots per EBS volume. AWS regional snapshot limits may also affect backup and recovery operations.
- Cohesity automatically retries AWS API requests when throttling occurs. Backup and recovery operations may take longer during periods of high AWS API activity.

Backup Amazon EBS Volume

Depending on the services you selected while registering the AWS source, the Cohesity cluster automatically discovers all supported workloads in the specified AWS account ID. The discovered workloads appear under **Data Protection > Sources**. When you open the registered AWS account source, each workload type is displayed on its own tab. Select the **EBS** tab to view all discovered EBS volumes. For more information on registering the AWS, see Register AWS Source.

Cohesity discovers both standalone EBS volumes and EBS volumes attached to EC2 instances. You can filter discovered volumes by attachment state:

- Detached (standalone volumes)
- Attached
- All

The EBS tab displays the entity hierarchy for Amazon EBS, allowing you to navigate from the account level down to individual volumes.

Amazon EBS Volume Entity Hierarchy

The Amazon EBS entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Account	The AWS account registered with Cohesity for Amazon EBS volume protection.
AWS Region	The AWS region where the Redshift Serverless resources reside.
Availability Zone	The availability zone within the region.
EBS Volume	The individual volume that can be selected for backup and recovery.

Backup Scope by Amazon EBS Volume Hierarchy

The Cohesity cluster allows you to back Amazon EBS volumes at different levels of the entity hierarchy. You can choose to back up:

- **AWS Account** – Backup all the Amazon EBS volumes within the registered AWS account.
- **AWS Region** – Backup all the Amazon EBS volumes within a specific region of the registered account.
- **Availability Zone** – Backup all the Amazon EBS volumes within a specific availability zone.
- **EBS Volume** – Backup individual Amazon EBS volume.

Supported Backup Types

Cohesity supports both full and incremental backup for Amazon EBS volumes:

- **Full backup**: The first backup of an EBS volume is always a full backup.
- **Incremental backup**: Subsequent backups capture only the blocks that have changed since the previous backup, reducing backup time and storage consumption.

The backup behavior depends on the AWS snapshot retention setting configured in the Protection Group for the EBS volume:

Retain Last AWS Snapshot	Backup Behavior
Enabled	After the initial full backup, subsequent backups are incremental. The most recent AWS snapshot is retained in your AWS account and used to identify changed blocks for subsequent backups.
Disabled	After the initial full backup, subsequent backups are also full backups. Because no AWS snapshot is retained after the backup completes, Cohesity cannot identify changed blocks for subsequent backups.

Protect Amazon EBS Volume

In Cohesity, Amazon Redshift Serverless namespace databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

To protect Amazon EBS volume:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **EBS**.
3. In the **New Protection** page, click **Add Objects** and select your registered source, or click **Register Source** to add a new one.
4. Select the volumes for protection using one of the following views:
   - **List View**: Use this option to select the volumes for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select volumes for protection at various hierarchy levels.

     Optionally,

     You can click the icon to enable Auto Protect at various hierarchy levels (account, AWS region, or Availability zone levels). When this option is enabled

at a particular level, all the volumes that are added to that level in the future are automatically added to the Protection Group and are protected from the next protection run.

You can exclude the volumes from auto-protection within a hierarchy level by clicking the auto-protect icon ( ) next to the volume. The volumes excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the volumes for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all volumes that match the selected tag for protection and backup operations.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added volumes with this tag.
  - You can exclude the volumes from auto-protection within a tag by clicking the auto-protect icon ( ) next to the volume. The volumes excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude volumes with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue.**
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. Create or select an existing policy from the **Policy** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	Retain AWS Snapshot	Use the **Retain last AWS snapshot** option to retain the most recent AWS EBS snapshot in your AWS account after a backup completes. When enabled, the most recent AWS snapshot is retained and used to identify changed blocks for incremental backups. Only the changed data is transferred to the Cohesity cluster, which can reduce backup duration and data transfer compared to full backups. When disabled, the AWS snapshot created during the backup is deleted after the backup completes. Because no previous AWS snapshot is available for comparison, each backup is performed as a full backup. > **Note:** The first backup of an EBS volume is always a full backup, regardless of the Retain AWS Snapshot setting.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon EBS Volume Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new EBS volumes to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for the EBS volume protection.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for the EBS volume protection.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for the EBS volume protection.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for the EBS volume protection.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
    - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon EBS Volume

You can recover the Amazon EBS volume only to its original location

> **Note:** We recommend that you review the Plan and Prepare for Amazon EBS Volume Protection.

To recover Amazon EBS volume:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **EBS**.
3. Search for the EBS volume by entering characters of the EBS volume name, or by specifying the wildcard character *.
4. From the list of search results, select the EBS volume to recover.
    - By default, the latest snapshot of the EBS volume is selected for recovery.
    - To recover from a different snapshot, hover over the snapshot and click the **Edit** icon. The **Edit Recovery Point for <*volume_name*>** page appears.
    - By default, snapshots are displayed in the List mode. You can switch to the Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

    To select a snapshot from the **List** mode:

    a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*volume_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date

range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

b. Select the snapshots to be recovered.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

To select a snapshot from the **Timeline** mode:

a. Select the **Timeline** tab from the top-right corner of the **Edit recovery point for <*volume_name*>** page.

b. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

c. Select the snapshot to be recovered by clicking the corresponding blue dot.

d. By default, the snapshot is recovered from the local Cohesity cluster.

e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, type Yes to confirm.

> **Note:**
> - Currently, recovering to original location is only supported.
> - Restoring to the original location overwrites the existing EBS volume.

7. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Click the Edit icon and provide a new name by adding Prefix and/or Suffix to the recovered Amazon EBS volume.
Task Name	Update the default name of the recovery task.

8. Click **Recover**.

Amazon Simple Storage Service Bucket

Cohesity's software-defined platform provides a simple, fast, and cost-effective data protection solution for the Amazon Simple Storage Service (S3) buckets.

To protect the Amazon S3 bucket, Cohesity leverages the Amazon S3 inventory report. The inventory report contains the list of all the objects available on the Amazon S3 bucket you selected for protection. Cohesity uses this report to perform the first full backup. This report

is also used to periodically reconcile the list of objects. The subsequent incremental backups are performed by using the AWS EventBridge capability.

With this protection approach, Cohesity can back up multi-billion objects at a faster rate.

> **Note:**
> - This is an Early Access feature on Physical and Virtual Edition clusters. Contact your Cohesity account team to enable the feature.
> - This feature is not supported on Cohesity Cloud Edition (legacy) on AWS (formerly known as AWS Cloud Edition cluster)
> - This feature is Generally Available (GA) on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster).

Plan and Prepare for Amazon S3 Bucket Protection

Before protecting your Amazon S3 bucket, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon S3 with Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Supported Amazon S3 Storage Class

Cohesity supports the data protection of the following Amazon Simple Storage Service Bucket (S3) storage class:

- Amazon S3 Standard
- Amazon S3 Intelligent-Tiering
- Amazon S3 Standard-IA
- Amazon S3 One Zone-IA

Permission for AWS S3 Inventory Report

- To write objects to the Amazon S3 bucket, you must add the `S3:PutObject` permission to the S3 bucket policy attached to the Amazon S3 bucket where you want to create the inventory report.

  The following is an example of an Amazon S3 bucket policy that allows **s3.amazonaws.com** to write (Put) objects in the Amazon S3 bucket:

```
{
    "Version": "2012-10-17",
    "Id": "S3-Console-Auto-Gen-Policy-1698064515475",
    "Statement": [
        {
            "Sid": "InventoryAndAnalyticsExamplePolicy",
            "Effect": "Allow",
            "Principal": {
                "Service": "s3.amazonaws.com"
            },
            "Action": [
                "s3:PutObject"
            ],
            "Resource": "arn:aws:s3:::<s3bucketname_you_chose_for_inventory>/*",
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "<account_Id>",
                    "s3:x-amz-acl": "bucket-owner-full-control"
                }
            }
        }
    ]
}
```

- If you're using Server-Side Encryption with AWS Key Management Service (SSE-KMS) for the bucket where you want to create the inventory report, ensure that the KMS key associated with that bucket has the necessary permissions to allow Amazon S3 to access it. To do this, add the following permissions to the KMS key policy:

```
{
            "Sid": "Allow Amazon S3 use of the KMS key",
            "Effect": "Allow",
            "Principal": {
                "Service": "s3.amazonaws.com"
            },
            "Action": "kms:GenerateDataKey",
            "Resource": "*"
        }
```

Create a Lifecycle Rule on Amazon S3

To delete the older inventory reports from the Amazon S3 bucket, you must create a lifecycle rule on the Amazon S3 bucket. You can delete all the inventory reports older than 30 days. For information on creating a lifecycle rule, see Amazon documentation.

Create AWS IAM Role for Glue Job

Cohesity uses AWS Glue jobs to process and sort Amazon S3 inventory reports during full and baseline incremental backups of AWS S3 buckets.

For these Glue jobs to run successfully, an AWS IAM role named `AWSGlueServiceRole` must exist in your AWS account with the required permissions and trust relationship.

If the AWS source is registered at the organization level, ensure that the `AWSGlueServiceRole` IAM role and permissions are configured in each member account containing Amazon S3 buckets to be protected.

If the AWS source is registered at the account level using IAM role, then the role required for Glue job is created by default.

If the AWS source is registered at the account level using IAM user, then you must manually create the `AWSGlueServiceRole` using the steps described below:

Before You Begin

Before creating a new IAM role:

- Verify whether the `AWSGlueServiceRole` already exists in your AWS account or member account.
- If the role exists, ensure it has:
  - The required managed policies
  - The correct trust relationship
- If the role does not exist, create it using one of the supported methods described below.

Verify Whether the IAM Role Exists

Perform the following steps to verify whether the `AWSGlueServiceRole` already exists:

1. Log in to the **AWS Management Console**.
2. Navigate to **IAM** > **Roles**.
3. Search for **`arn:aws:iam::<account-id>:role/service-role/AWSGlueServiceRole`**.

If the role exists, verify the following:

- Required Permissions
- Required Trust Relationship

Required Permissions

The AWSGlueServiceRole must have the following policies attached:

- **AWSGlueServiceRole**

  (AWS-managed policy for Glue service execution)
- **AmazonS3FullAccess**

  Required to:
  - Read S3 inventory reports
  - Write processed objects to the destination S3 path

  ##### Required Trust Relationship

  The role must allow AWS Glue to assume it. Verify that the trust relationship includes the following:

```
{

    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",

            "Principal": {

                "Service": "glue.amazonaws.com"
            },

            "Action": "sts:AssumeRole"
        }
    ]
}
```

Create the IAM Role

If `AWSGlueServiceRole` does not exist, create it using one of the following methods, depending on whether the required managed policy already exists in your AWS account. If the AWS source is registered at the organization level, ensure that the `AWSGlueServiceRole` is created in each member account containing Amazon S3 buckets you want to protect.

- Create the IAM Role Using AWS Glue Getting Started Wizard
- Create the IAM Role Using CloudFormation Template

Create the IAM Role Using AWS Glue Getting Started Wizard

This method is recommended when the customer managed policy AWSGlueServiceRole-EZCRC-s3Policy does not already exist in your AWS account or member account.

> **Important:** Creating the role from the AWS Glue Getting Started Wizard may fail if a role or policy with the same name already exists. In such cases, we recommend you create the IAM Role using the CloudFormation Template.

Perform the following steps in the applicable AWS account or member account to create the IAM role and attach the required policy:

1. Log in to the AWS Management Console.
2. Navigate to **AWS Glue**.
3. Click **Set up roles and users**.
4. On the **Choose IAM users and roles for AWS Glue** page, retain the default selections and click **Next**.
5. On the **Grant Amazon S3 access** page, retain the default settings and click **Next**.
6. On the **Choose a default service role** page:

1. Ensure the role name is **AWSGlueServiceRole** and the ARN is **`arn:aws:iam::<account-id>:role/service-role/AWSGlueServiceRole`**.
2. Click **Next**.

7. On the **Review and confirm** page, click **Apply changes**.

   The IAM role is created.
8. Navigate to the IAM service and confirm that the role, `AWSGlueServiceRole` exists.
9. Confirm the **AWSGlueServiceRole** managed policy is attached to the role.
10. Confirm that the trust relationship attached to the role allows the AWS Glue service to assume the role:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": "glue.amazonaws.com"
            },
            "Action": "sts:AssumeRole"
        }
    ]
}
```

11. Add **AmazonS3FullAccess** policy to this role to allow accessing the inventory report and to put the objects on the destination path, using the steps below.
12. From the **Permissions** tab of the role, click **Add permissions** > **Attach policies**.
13. Select **AmazonS3FullAccess**.

Create the IAM Role Using CloudFormation Template

Creating `AWSGlueServiceRole` using the CloudFormation template is recommended when the **AWSGlueServiceRole** managed policy already exists in the applicable AWS account or member account.

You can use the following CloudFormation template for creating the IAM Role:

```
AWSTemplateFormatVersion: "2010-09-09"
Description: "Creates an IAM role for AWS Glue jobs to assume for Cohesity Backups"
Resources:
  AWSGlueServiceRole:
    Type: AWS::IAM::Role
```

```
    Properties:
      RoleName: AWSGlueServiceRole
      Path: /service-role/
      Description: Role for Glue jobs to assume for Cohesity Backups
      ManagedPolicyArns:
        - arn:aws:iam::aws:policy/service-role/AWSGlueServiceRole
        - arn:aws:iam::aws:policy/AmazonS3FullAccess
      AssumeRolePolicyDocument:
        Version: "2012-10-17"
        Statement:
          - Effect: Allow
            Principal:
              Service: glue.amazonaws.com
            Action: sts:AssumeRole
      Tags:
        - Key: UniqueTag
          Value: cohesity
Outputs:
  RoleName:
    Description: Name of the created IAM Role
    Value: !Ref AWSGlueServiceRole
  RoleArn:
    Description: ARN of the created IAM Role
    Value: !GetAtt AWSGlueServiceRole.Arn
```

To run the CloudFormation Template:

> **Note:** If the AWS source is registered at the organization level, deploy the CloudFormation template in each member account containing Amazon S3 buckets you want to protect.

1. Log in to the AWS Management Console.
2. Navigate to **All services** and select **CloudFormation**.
3. Click **Create stack** > **With new resources (standard)** to create a stack using the downloaded CloudFormation Template.
4. In the **Prerequisite - Prepare template** page, select **Upload a template file**.
5. Click **Choose File** to browse the JSON CloudFormation Template.

6. Click **Next**.
7. Provide a Stack name. The stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).
8. Click **Next**.
9. Retain all the default stack configuration options and click **Next**.
10. On the Review page, select the checkboxes that need to be acknowledged and click **Create stack**.

Permission for Amazon Key Management Service (KMS)

If the Amazon S3 bucket you want to protect is encrypted with Server-side encryption with AWS Key Management Service keys (SSE-KMS) or Dual-layer server-side encryption with AWS Key Management Service keys (DSSE-KMS), then for Cohesity to access the Amazon S3 bucket, you must perform one of the following actions:

- Add the IAM role created by the Cloud Formation template to the AWS KMS user.
- Add the following permission to the Key policy attached to the AWS KMS:
  - kms:Encrypt
  - kms:Decrypt
  - kms:ReEncrypt*
  - kms:GenerateDataKey*
  - kms:DescribeKey

  For example:

```
{
    "Version": "2012-10-17",
    "Id": "AccessKeyId",
    "Statement": [
        {
            "Sid": "Allow use of the key to cohesity role",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::<AWS-ACCOUNT>:role/<ROLE-NAME>"
            },
            "Action": [
                "kms:Encrypt",
                "kms:Decrypt",
                "kms:ReEncrypt*",
                "kms:GenerateDataKey*",
                "kms:DescribeKey"
            ],
            "Resource": "*"
        }
    ]
}
```

Considerations

- Cohesity does not support:
  - Browse and recover an object in an Amazon S3 bucket. However, you can recover multiple objects by specifying the object prefix in the recovery task under the S3 Prefixes to Recover option.
  - The backup of older versions of the Amazon S3 versioned bucket. Only the latest version of the versioned Amazon S3 bucket is backed up.
  - Cross-region backup and recovery, i.e., the backup and recovery of Amazon S3 buckets located in a different cloud region than where your data is backed up.
- The Amazon S3 buckets where you want to create the inventory report and the Amazon S3 bucket you want to protect must be in the same region.
- If the SQS is deleted between backups, all the changes between these backups will be skipped in the next incremental backups.
- Cohesity will skip the backup of Amazon S3 objects that are present in the following access tiers of the Amazon S3 Intelligent Tiering during the protection:
  - Archive Access Tier
  - Deep Archive Access Tier
- Cohesity does not support restoring only metadata. The metadata of Amazon S3 objects will be restored only if the object itself is also restored.
- Cohesity does not remove older inventory reports from the Amazon S3 bucket. Therefore, you must create a lifecycle rule to remove these objects from the Amazon S3 bucket to avoid storage issues.

Backup Amazon S3 Bucket

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads including Amazon S3 associated with the specified AWS Source. The discovered workloads appear under **Data Protection > Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **S3** tab to view all discovered S3 buckets. Depending on the registration scope, you can navigate from the organization level or account level down to individual S3 buckets. For information on registering AWS, see Register AWS Source.

Amazon S3 Entity Hierarchy

Based on the registration level of your AWS source, the S3 tab displays the S3 entity hierarchy, allowing navigation from the organization or account level down to individual S3 buckets. The **Amazon S3** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with the Cohesity cluster for S3 protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The organizational units within the AWS organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain S3 buckets.
AWS Account	The standalone AWS account registered with the Cohesity cluster for S3 protection. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the S3 buckets reside.
Amazon S3	The S3 bucket within the AWS source.

Backup Scope by Amazon S3 Hierarchy

The Cohesity cluster allows you to back up Amazon S3 at different levels of the entity hierarchy. Depending on the registration level of your AWS source, you can choose to back up:

- **AWS Organization** – Protect all S3 buckets across all member accounts in the registered AWS organization.
- **Organizational Unit (OU)** – Protect all S3 buckets within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all S3 buckets within a specific member account in the registered AWS organization.
- **AWS Account** – Protect all S3 buckets within the registered standalone AWS account.
- **AWS Region** – Protect all S3 buckets within a specific AWS Region.
- **S3 Bucket** – Protect an individual S3 bucket.

Backup Process

Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. When you register the AWS source, you can specify the Amazon S3 bucket where you want to create the inventory report. The requirements for this bucket depend on how the source is registered:

- **Organization-level registered source**: The Amazon S3 bucket you specify must be within the same cloud region as the Amazon S3 buckets selected for protection.

- **Account-level registered source**: The Amazon S3 bucket you specify must be within the same account and the same cloud region as the Amazon S3 buckets selected for protection.

When you initiate the Amazon S3 buck et protection, AWS will create an inventory report. It may take up to 48 to 72 hours for AWS to create an inventory report. This inventory report will contain the list of all the objects available on the Amazon S3 bucket you selected for protection. Cohesity uses this inventory report to perform the first full backup of the Amazon S3 bucket. Once Cohesity performs the first full backup, Cohesity utilizes Amazon EventBridge and SQS queues to perform incremental backups of the Amazon S3 bucket.

With this protection approach, Cohesity can back up multi-billion Amazon S3 objects at a faster rate.

Supported Backup Types

Cohesity supports both full and Incremental backups for Amazon S3.

Protect Amazon S3 Bucket

In Cohesity, Amazon S3 buckets are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

Before you Begin

Before protecting your Amazon S3 buckets, ensure you have met the prerequisites and understood the considerations. For more information, see Plan and Prepare for Amazon S3 Bucket Protection.

Add Protection to Your Registered Amazon S3 Buckets

To protect your Amazon S3 buckets:

1. Select **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then click **Protect** > **Amazon S3**.

   The New Protection page appears.
3. Click **Add Objects**, and from the **Registered Source** drop-down list, select the AWS source you want to protect. You can also register a new source with the Cohesity cluster by clicking the **Register Source** option.
4. From the objects list, select the S3 buckets you want to protect using one of the following views or options:

- **List View**: Use this option to select the S3 buckets for protection from a flat, searchable list.
- **Hierarchy View**: Use this option to select S3 buckets for protection at various hierarchy levels.

  Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new S3 buckets added.

  If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:

  - AWS Organization
  - Organizational Unit (OU)
  - Member Account
  - AWS Region

  If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:

  - AWS Account
  - AWS Region

  You can exclude the S3 buckets from auto-protection within a hierarchy level by clicking the auto-protect icon ( ) next to the S3 bucket. The S3 buckets excluded from auto-protection are displayed with the exclude ( ) icon.
- **Tags**: Use this option to select the S3 buckets for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all S3 buckets that match the selected tag for protection.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added S3 buckets with this tag.
  - You can exclude the S3 buckets from auto-protection within a tag by clicking the auto-protect icon ( ) next to the S3 bucket. The S3 buckets excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude S3 buckets with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.

6. Provide a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. Choose a policy to specify backup frequency and retention. If you don't have a policy, you can create one. Ensure that the backup frequency set in the policy for Amazon S3 bucket is at most weekly. Backups may fail if the frequency is set for more than a week.
8. If you wish to change or configure any of the additional settings, select **More Options** and perform the below steps, or else, click **Protect**.
9. Under **Settings**, edit the following options if necessary:
    - **Storage Domain**: To store the backed up data, select an existing storage domain or create a new one by selecting Add Storage Domain.
    - **Start time**: Indicates what time the protection run should start. Enter the **Start Time** and select **AM** or **PM**. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone here.
    - **SLA**: Defines how long the administrator expects a protection run to take. Enter:
        - **Full**. The number of minutes you expect a full protection run, which captures all the blocks in an object, to take.
        - **Incremental**. The number of minutes you expect an incremental protection run, which captures only the changed blocks in an object, to take.
10. Under **Additional Settings**, configure the following option:
    - **Skip File on Errors**: Enable this option to continue the protection run even if any error is encountered when backing an Amazon S3 object. By default, this option is enabled. By disabling this option, the protection run will fail if one of the objects in Amazon S3 encounters an error.
    - **Enable ACL Backups**: Enable this option to backup ACL. You can backup ACLs only if ACLs are configured in the Amazon S3 bucket you chose to protect. By default, this option is disabled.
11. Click **Protect**.

Cohesity starts backing up the Amazon S3 buckets you selected.

Manage the Amazon S3 Bucket Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

You cannot update the registered AWS source and the assigned Storage Domain in the Protection Group settings.

Add Amazon S3 Buckets to the Protection Group

You can add new Amazon S3 buckets to an existing Protection Group for backup. The newly added buckets are backed up from the next protection run, and it will be a full backup.

To add Amazon S3 buckets to the Protection Group:

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you created for protecting the Amazon S3 bucket.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Select the Amazon S3 bucket you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added Amazon S3 buckets will be backed up in the next scheduled protection run.

Remove Amazon S3 Buckets from the Protection Group

From the Edit Protection page, you can remove Amazon S3 buckets from the Protection Group that you do not want to back up in the next protection run.

To remove the Amazon S3 buckets from the Protection Group:

1. Select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for protecting the Amazon S3 buckets.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Unselect the Amazon S3 buckets that you want to remove from the Protection Group.
6. Click **Save**.

The removed Amazon S3 buckets will not be backed up in the next scheduled protection run.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirement. In the Protection Policy, you can edit the settings listed in the Create or Use an Existing Protection Policy.

To edit the policy settings:

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon S3 buckets.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Policy** field, click **Edit**.

   The Edit Policy page appears.
5. Edit the settings listed in the Create or Edit a Standard Policy and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you created for protecting the Amazon S3 buckets.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, from the drop-down list in the policy field, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-Demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to backup Amazon S3 buckets. It provides you options to back up all the objects or only selected objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- If configured with replication targets, you can select those targets.
- If configured with archival targets, you can select those targets.

To start a protection run:

1. Select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for protecting the Amazon S3 buckets.
3. Click the action icon next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:

a. In the **Backup** drop-down list, select any one of the following:
   - Backup all Objects in the Protection Group
   - Backup only selected Objects in the Protection Group
   - All objects that failed last run

b. In the **Backup Type** drop-down list, select the backup type.

c. In the **Replicate** section, select the replication target, the retention duration, and enable the **Replicate only fully successful runs** checkbox if you want to replicate the data only after a successful protection run.

d. In the Archive section, select the archive target, the retention duration, and enable the **Archive only fully successful runs** checkbox if you want to archive the data only after a successful protection run.

5. Click **Run**.

Pause Future Protection Runs

You can stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for protecting the Amazon S3 buckets.
3. Click the action icon next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run **Pause Future Runs**.

To resume a paused protection run:

1. Select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the action icon next to the Protection Group and select **Resume**.

Delete the Protection Group

Cohesity provides the following ways to delete the Protection Group associated with the Amazon S3 buckets:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Protection Group are not deleted

- **Delete Protection Group and Snapshots**: Delete the Protection Group and all Snapshots, any local Snapshots captured by this Protection Group and stored on this Cluster are deleted. In addition, any archived Snapshots stored on external targets that were captured by this Protection Group are also deleted.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for protecting the Amazon S3 buckets.
3. Click the action icon next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group**? dialog box, select any one of the following options:
    - **Delete Protection Group Only**
    - **Delete Protection Group and Snapshots**

If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Run.

If the Protection Group is replicating to a remote cluster, the inactive job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the inactive job.

Recover Amazon S3 Bucket

After you protect your Amazon S3 buckets, you can recover them to their original location or a new location using Cohesity.

Recover Amazon S3 to the Original Location

To recover your protected Amazon S3 buckets to their original location:

1. Navigate to **Data Protection** > **Recoveries**. Click **Recover** on the top-right of the page and then select Amazon S3 as the recovery source.
2. Search for a S3 bucket or Protection Group that contains the snapshots to recover from by entering characters of the S3 bucket name, Protection Group name, or by specifying the wildcard character *.

    You can optionally filter the search results by source, Protection Group, or date.
3. From the list of search results, select an S3 bucket you want to recover.

    By default, the latest snapshot is selected for recovery.

    To recover from a different snapshot, perform the following steps:

a. From the selected lists, click the edit icon.

   The Edit recovery point for <S3bucket_name> page appears. By default, the page displays the latest snapshot available.

b. Optionally, you can also select a snapshot from a different date by performing one of the following:

   - Selecting the date from the **Choose a date** drop-down list. Once you select the date, the Edit recovery point page displays the available snapshots on the specified date.
   - Clicking the **List** tab. The List tab displays all the available snapshots of the selected S3 bucket.

c. Select the snapshots to be recovered.

   By default, the snapshot is recovered from the local cluster.

d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 5.

4. Click **Next: Recover Options**.
5. Under **Recover To**, select **Original Location**.
6. Enable **Overwrite Existing Bucket** if you want to recover the S3 bucket by overwriting the original S3 bucket. By default, this option is disabled.

   If the data in the source S3 object and target S3 are the same, then the target S3 will not be updated even if the **Overwrite Existing Bucket** option is enabled. If the target bucket is versioned and the S3 object name of the target and the source match, then a new version of the S3 object will be created.
7. In the **Object Prefix** field, add prefix value to the names of the S3 objects that are recovered by this task. For information about the S3 object naming convention, see the AWS documentation.
8. Select your **Recovery Options**:

   a. **S3 Prefixes to Recover**: To perform granular recovery of the Amazon S3 bucket, enable the **Turn ON** to recover specific prefixes option and provide the prefix of the objects to be recovered.

      For example, `Finance/statement/document`.

      You can click the **add** icon to provide multiple prefixes.

      The maximum number of prefixes that you can provide per recovery task is 100.

   b. **Continue on Error**: Enable this option if you want to continue the recovery even if one of the S3 objects encounters an error. By default, this option is

disabled and the recovery operation will fail if one of the objects encounters an error.

c. **Task Name**: Change the default name of the recovery task.

9. Click **Recover**.

Cohesity begins to restore the selected Amazon S3 bucket.

Recover Amazon S3 to New Location

To recover your protected Amazon S3 buckets to a new location:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click Recover on the top-right of the page and then select Amazon S3 as the recovery source.
3. Search for a S3 bucket or Protection Group that contains the snapshots to recover from by entering characters of the S3 bucket name, Protection Group name, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, or date.
4. From the list of search results, select an S3 bucket you want to recover.

   By default, the latest snapshot is selected for recovery.

   To recover from a different snapshot, perform the following steps:

   a. From the selected lists, click the **edit** icon.

      The Edit recovery point for <S3bucket_name> page appears. By default, the page displays the latest snapshot available.

   b. Optionally, you can also select a snapshot from a different date by performing one of the following:

      - Selecting the date from the **Choose a date** drop-down list. Once you select the date, the Edit recovery point page displays the available snapshots on the specified date.
      - Clicking the **List** tab. The List tab displays all the available snapshots of the selected S3 bucket.

   c. Select the snapshots to be recovered.

      By default, the snapshot is recovered from the local cluster.

      Optionally, you can also select snapshots archived to the cloud by clicking the cloud icon if available in the **Location** option.

   d. Click Select Recovery Point and then go to step 5.
5. Click **Next: Recover Options**.

6. Under **Recover To**, select any one of the following options based on where you want to restore the S3 bucket and then provide the required information:
   - **Original Location:** Restore the S3 bucket to the original S3 bucket.
   - **New Location**: Restore the S3 to an alternate S3 bucket.

   Based on the registration level of the AWS source, select the recovery destination hierarchy:
   - If the AWS source is registered at the organization level, select the following:
     - **Source**: Select the AWS Organization where you want to recover the S3 bucket. You can recover the S3 bucket within the same organization as the source S3 bucket.
     - **AWS Account**: Select the member account to use as the recovery destination.
     - **Region**: Displays the AWS region of the source S3 bucket. Cross-region recovery is not supported; therefore, the region is selected automatically and cannot be changed.
     - **S3 Bucket**: Select the S3 bucket to which you want to recover.
   - If the AWS source is registered at the account level, select the following:
     - **AWS Account**: Select the registered AWS account to use as the recovery destination.
     - **Region**: Displays the AWS region of the source S3 bucket. Cross-region recovery is not supported; therefore, the region is selected automatically and cannot be changed.
     - **S3 Bucket**: Select the S3 bucket to which you want to recover.
7. Enable **Overwrite Existing Bucket** if you want to recover the S3 bucket by overwriting the original S3 bucket. By default, this option is disabled.

   If the data in the source S3 object and target S3 are the same, then the target S3 will not be updated even if the **Overwrite Existing Bucket** option is enabled. If the target bucket is versioned and the S3 object name of the target and the source match, then a new version of the S3 object will be created.
8. In the **Object Prefix** field, add prefix value to the names of the S3 objects that are recovered by this task. For information about the S3 object naming convention, see AWS documentation.
9. Select your **Recovery Options**:
   a. **S3 Prefixes to Recover**: To perform granular recovery of the Amazon S3 bucket, enable the **Turn ON** to recover specific prefixes option and provide the prefix of the objects to be recovered.

      For example, `Finance/statement/document`.

You can click the **add** icon to provide multiple prefixes.

The maximum number of prefixes that you can provide per recovery task is 100.

b. **Continue on Error**: Enable this option if you want to continue the recovery even if one of the S3 objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.

c. **Task Name**: Change the default name of the recovery task.

10. Click **Recover**.

Cohesity begins to restore the selected Amazon S3 bucket.

Amazon Aurora MySQL

Amazon Aurora MySQL is a fully managed, MySQL-compatible, cloud-based relational database service offered by Amazon Web Services (AWS). Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Amazon Aurora MySQL clusters.

Backup Methods

You can protect your Amazon Aurora MySQL using one of the following backup methods:

- **AWS Snapshot methods**: Cohesity uses this option to protect Amazon Aurora MySQL clusters. For protecting the Amazon Aurora MySQL clusters, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon Aurora MySQL cluster.
- **Ingest method**: Cohesity uses this option to protect the Amazon Aurora MySQL cluster at the database level by ingesting the backup data to the Cohesity cluster.

> **Important:**
> - Cohesity supports the protection of Amazon Aurora MySQL only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster) and Cohesity Cloud Edition (legacy) on AWS.
> - On Cohesity Cloud Edition on AWS, Cohesity supports the protection of Amazon Aurora MySQL at both the cluster level (via AWS snapshot-based protection) and the database level (via ingest-based protection).
> - On Cohesity Cloud Edition (legacy) on AWS, Cohesity supports the backup and recovery of Amazon Aurora MySQL only at the cluster level.

This topic provides information on how you can protect Amazon Aurora MySQL on Cohesity Cloud Edition on AWS.

Plan and Prepare for Amazon Aurora MySQL

Before protecting Amazon Aurora MySQL with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon Aurora MySQL with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Authenticate for Ingest-based Protection

If you are planning to protect the Amazon Aurora MySQL using Ingest-based protection method, you must authenticate the Cohesity cluster to discover the databases of the Amazon Aurora MySQL cluster using one of the following methods:

Authentication Method	Permissions	Example Command

Username and **password** of the Amazon Aurora MySQL	• SELECT
	• INSERT
	• UPDATE
	• DELETE
	• CREATE
	• DROP
	• RELOAD
	• PROCESS
	• REFERENCES
	• INDEX
	• ALTER
	• SHOW DATABASES
	• LOCK TABLES
	• CREATE VIEW
	• SHOW VIEW
	• CREATE ROUTINE
	• ALTER ROUTINE
	• SHOW ROUTINE
	• EVENT
	• TRIGGER
	• SET_USER_ID
		Create the Database user:
		`CREATE USER 'dbuser'@'%' IDENTIFIED BY 'StrongPassword!';`
		Grant below permission to user:
		`GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, PROCESS, REFERENCES, INDEX, ALTER, SHOW DATABASES, LOCK TABLES, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, SHOW ROUTINE, EVENT, TRIGGER, SET_USER_ID ON *. WITH GRANT OP'`
		`FLUSH PRIVILEGES;`

IAM-authenticated MySQL user		Create the Database user, using the master login:
		`CREATE USER <iamuser> IDENTIFIED WITH AWSAuthenticationPlugin as 'RDS';`
		Grant below permission to user:
		`GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, PROCESS, REFERENCES, INDEX, ALTER, SHOW DATABASES, LOCK TABLES, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, SHOW_ROUTINE, EVENT,TRIGGER, SET_USER_ID ON . TO `<iamuser>`@'%' WITH GRANT OPTION; FLUSH PRIVILEGES;`

For information on the steps to authenticate the Cohesity cluster to the Amazon Aurora for MySQL in the Cohesity cluster UI, see Protect Amazon Aurora MySQL Using Ingest Method.

Ports Used for Communication

For information on port numbers required, seeAWS Cloud Database Workloads.

Considerations for AWS Aurora MySQL Data Protection

Review and understand the following considerations before you protect Amazon Aurora MySQL:

- Cohesity performs only full backup of Amazon Aurora MySQL.
- Cohesity protects only the user databases within the Amazon Aurora MySQL. Protection of system databases are not supported.
- The backup and recovery of Amazon Aurora MySQL cluster at database level is supported only on Cohesity Cloud Edition on AWS.
- Point-in-time restore (PITR) is supported when recovering Amazon Aurora MySQL at cluster-level.
- Restoring a backup taken from a higher Amazon Aurora MySQL database version to a lower version is not supported.
- You cannot perform ingest-based and snapshot-based protections within the same Protection Group.
- If the databases of an Aurora cluster are discovered, then to perform snapshot-based protection, you must select all the databases of that Aurora cluster.
- If you select Aurora clusters for which no databases are discovered, then you cannot perform ingest-based protection.
- Cross-account restore of an Amazon Aurora MySQL cluster protected using the AWS Snapshot protection type is not supported.
- When recovering encrypted Aurora clusters to the same AWS account but a different AWS region, you must create a KMS encryption key in the target region with the same alias name as the KMS key used to encrypt the source Aurora cluster.
- When performing cross-region cross-account recovery of an Amazon Aurora cluster encrypted with a customer-managed KMS key, ensure the following:
    - The source account's KMS key alias is available in both the source and destination regions of the destination account.

  Example:

  If A1 is the source account with region R1, and A2 is the destination account with region R2, ensure the KMS key alias from A1 is available in both R1 and R2 of A2.

    - The destination account has permission to access the source account's KMS key alias.
    - The option group from the source account is available in both the source and destination regions of the destination account.

Backup Amazon Aurora MySQL

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads including Amazon Aurora MySQL associated with the AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **Aurora MySQL** tab to view all discovered clusters.

The **Aurora MySQL** tab displays the entity hierarchy for the Amazon Aurora MySQL environment, allowing navigation from the organization or account level down to individual databases. For information on registering AWS, see Register AWS Source.

Authenticate for Ingest-based Protection

Aurora MySQL Entity Hierarchy

The **Aurora MySQL** tab displays the entity hierarchy for the Aurora MySQL cluster, allowing you to navigate from the organization or account level down to individual databases.

The **Aurora MySQL** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with Cohesity for Aurora MySQL protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The organizational units within the AWS organization that contain member accounts
Member Account	The AWS member accounts within the organization that contain Aurora MySQL clusters.
AWS Account	The AWS account registered with Cohesity for Aurora MySQL protection. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the Aurora MySQL resources reside.
Availability Zone	The availability zone within the region where the Aurora MySQL cluster is deployed.
Aurora MySQL Cluster	The specific Aurora MySQL cluster within the selected availability zone.
Database	The individual databases hosted within the Aurora MySQL cluster.

Backup Scope by Aurora MySQL Hierarchy

The Cohesity cluster allows you to back up Amazon Aurora MySQL at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level.

You can choose to back up:

- **AWS Organizatio**n – Protect all Aurora MySQL clusters and databases across all member accounts in the registered AWS organization.

- **Organizational Unit (OU)** – Protect all Aurora MySQL clusters and databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all Aurora MySQL clusters and databases within a specific member account in the registered AWS organization.
- **AWS Account** – Protect all Aurora MySQL clusters and databases within the registered standalone AWS account.
- **AWS Region** – Protect all Aurora MySQL clusters and databases within a specific AWS Region.
- **Availability Zone** – Protect all Aurora MySQL clusters and databases within a specific Availability Zone.
- **Aurora MySQL Cluster** – Protect all databases within a specific Availability Zone.
- **Database** – Protect individual databases within an Aurora MySQL cluster by providing the database credentials.

Aurora MySQL Backup Methods

Cohesity provides the following Protection Types for the protecting Amazon Aurora MySQL:

- **AWS Snapshot method**: Cohesity uses this option to protect Amazon Aurora MySQL clusters. For protecting the Amazon Aurora MySQL clusters, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon Aurora MySQL cluster.
- **Ingest method**: Cohesity uses this option to protect the Amazon Aurora MySQL cluster at the database level by ingesting the backup data to the Cohesity cluster.

Supported Backup Types

Cohesity performs only full backup of Amazon Aurora MySQL.

Protect Amazon Aurora MySQL Using Ingest Method

With the Ingest protection method, Cohesity cluster protects the Amazon Aurora MySQL at the database level.

Authenticate for Aurora MySQL Protection

To protect the Amazon Aurora MySQL at the database level using the Ingest method, you must provide the credentials of the Amazon Aurora MySQL cluster while creating the protection. By providing the credentials, Cohesity discovers the databases as objects for protection.

Perform the following steps for authentication:

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AWS source and click into it.
3. Click the **Aurora MySQL** tab.

4. Select the checkbox for the hierarchy level (AWS Organization, OUs, account, region, availability zone, or cluster level), and then click **Database Credentials** above the checkboxes.

   The Database Credentials for <cluster> appears.

5. Select one of the following authentication methods for Cohesity to discover the databases of the Amazon Aurora MySQL clusters:

   - **Credentials**: Enter the **Username**, and **Password** for the Amazon Aurora MySQL clusters.
   - **IAM Authentication**: Enter the **Username** created within the database for IAM authentication.

   For information on the authentication requirements, see Authenticate for Ingest-based Protection

6. Click **Save**.

The databases of the Amazon Aurora MySQL clusters are discovered.

Protect Amazon Aurora MySQL Database

To protect your Amazon Aurora MySQL using Ingest method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS MySQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon Aurora MySQL databases for protection using one of the following views:

   - **List View**: Use this option to select the Amazon Aurora MySQL databases for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select databases for protection at various hierarchy levels.

     Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon Aurora MySQL clusters.

     - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
       - **AWS Organization**
       - **Organizational Unit (OU)**
       - **Member Account level**

- **AWS Region**
- **Availability Zone**
- **Aurora Cluster**

- If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
  - **AWS Account**
  - **AWS Region**
  - **Availability Zone**

    Aurora Cluster

  You can also exclude the Amazon Aurora MySQL databases from auto-protection within an Organization, OU, member account, account, region, or availability zone by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the databases for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all databases with that tag for protection.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon Aurora MySQL databases with this tag.
  - You can exclude the Amazon Aurora MySQL databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude databases with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The Ingest-based policy is enabled by default. Create or select an existing policy from the **Policy (Cohesity Ingest)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Protect Amazon Aurora MySQL Using AWS Snapshot Method

With the AWS Snapshot protection method, Cohesity cluster protects the Amazon Aurora MySQL of your AWS account at the cluster level.

To protect your Amazon Aurora MySQL cluster using AWS snapshot method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **Aurora MySQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon Aurora MySQL clusters for protection using one of the following views:
    - **List View**: Use this option to select the Aurora MySQL clusters for protection from a flat, searchable list.

- **Hierarchy View**: Use this option to select clusters for protection at various hierarchy levels.

  Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new Aurora MySQL clusters added to your AWS source.

  - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
    - **AWS Organization**
    - **Organizational Unit (OU)**
    - **Member Account**
    - **AWS Region**
    - **Availability Zone**
  - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
    - **AWS Account**
    - **AWS Region**
    - **Availability Zone**

  You can also exclude the Amazon Aurora MySQL clusters from auto-protection within an account, region, or availability zone by clicking the auto-protect icon ( ) next to the cluster. The clusters excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the Amazon Aurora MySQL clusters for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all Amazon Aurora MySQL clusters with that tag for protection.

  Optionally,

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon Aurora MySQL clusters with this tag.
  - You can exclude the Amazon Aurora MySQL clusters from auto-protection within a tag by clicking the auto-protect icon ( ) next to the cluster. The clusters excluded from auto-protection are displayed with the exclude ( ) icon.

- You can exclude clusters with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The AWS Snapshot-based policy is enabled by default. Create or select an existing policy from the **Policy (AWS Snapshot)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon Aurora MySQL Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon Aurora for MySQL cluster or databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon Aurora MySQL.

3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Aurora MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Aurora MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Aurora MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.

- Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.

5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon Aurora MySQL

Depending on the backup method you used for protecting Amazon Aurora MySQL, you can:

Recover Amazon Aurora MySQL at Cluster Level

Recover Amazon Aurora MySQL at Database Level

Recover Amazon Aurora MySQL at Cluster Level

You can recover the Amazon Aurora MySQL at cluster level if you have protected the Amazon Aurora MySQL using the AWS Snapshot Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon Aurora MySQL.

To recover Amazon Aurora MySQL at cluster level:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **Aurora MySQL**.
3. Select **Cluster** as the recovery task.
4. Search for the cluster by entering characters of the cluster name, or by specifying the wildcard character *.
5. From the list of search results, select the cluster to recover. By default, the latest snapshot of the cluster is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for <*cluster_name*>** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*cluster_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. You can also select any point on the green solid line on the slider to

restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.

b. Select the snapshots to be recovered.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

- To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the cluster and then provide the required information:
   - **Original Location**: Restore the cluster to the same region.

     > **Note:** If an Aurora cluster with the same name already exists, the restore will fail. Ensure the original cluster is not present or restore with a different name.

   - **New Location**: Restore the Aurora cluster to a different region. You can restore the Aurora cluster to the same or a different region within the same AWS account as the original Aurora cluster.

     Based on the registration level of the AWS source, select the recovery destination hierarchy:

     - If the AWS source is registered at the organization level:
       - **Source**: Select the AWS organization that contains the destination Aurora cluster. You can recover the Aurora cluster within the same organization as the source Aurora cluster.
       - **AWS Account**: Select the member account to which you want to restore the Aurora cluster.
       - **Region**: Select a region to restore the Aurora cluster.

- **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered Aurora cluster.
- **Subnet**: Select a subnet in the Amazon VPC to store the recovered Aurora cluster.
- **Availability Zone**: Select an availability zone in AWS to recover the Aurora cluster.

- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account to which you want to restore the Aurora cluster.
  - **Region**: Select a region to restore the Aurora cluster.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered Aurora cluster.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered Aurora cluster.
  - **Availability Zone**: Select an availability zone in AWS to recover the Aurora cluster.

To overwrite any existing Aurora cluster with the same name, enable **Overwrite existing Aurora cluster(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of Amazon Aurora MySQL cluster is not supported.
> - If an Aurora cluster with the same name already exists, the restore will fail.

8. **Database Instance Identifier**: Specify the unique key that identifies the database cluster that will be recovered.
9. **DB Option Group**: Select an option group that contains the option you want to attach to the Aurora cluster that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
10. **DB Parameter Group**: Select the database parameter group to associate with the Aurora cluster.
11. **DB Port**: Specify the TCP/IP port that the Aurora cluster will use for application connections. The connection string of any application connecting to the Aurora cluster must specify the port number of the Aurora cluster. Both the security group applied to the Aurora cluster and your company's firewalls must allow connections to this port. The default port is 3306.

12. Configure the following Recovery Options:
    a. **Multi A-Z Deployment**: Enable this option if you want the Aurora cluster to be recovered to have a standby cluster deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the Aurora cluster to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the Aurora cluster to automatically upgrade when a new minor database engine
13. **Task name**: Optionally, change the default name of the recovery task
14. Click **Recover**.

Recover Amazon Aurora MySQL at Database Level

You can recover the Amazon Aurora MySQL at database level if you have protected the Amazon Aurora MySQL using the Amazon Ingest Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon Aurora MySQL.

To recover Amazon Aurora MySQL:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **Aurora MySQL**.
3. Select **Database** as the recovery task.
4. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
5. From the list of search results, select the database to recover. By default, the latest snapshot of the database is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for *<database_name>*** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*database_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
  b. Select the snapshots to be recovered.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.
- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.

6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:
   - **Original Location**: Restore the database to the original Amazon Aurora MySQL cluster.

     > **Note:** Restoring to the original location overwrites the existing database.

     Type **Yes** to confirm.
   - **New Location**: Restore the database to a different Amazon Aurora MySQL cluster. Cohesity supports cross-region and cross-account restores of Amazon Aurora MySQL at the database level.

     Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source**: Select the AWS organization that contains the destination Amazon Aurora MySQL cluster. You can recover data only within the same AWS Organization as the source.
  - **AWS Account**: Select the member account that contains the destination Amazon Aurora MySQL cluster.
  - **Cluster**: Select the Amazon Aurora MySQL cluster to which you want to restore the database.
- If the AWS source is registered at the account level:
  - **AWS Account**: Select the registered AWS account that contains the destination Amazon Aurora MySQL cluster.
  - **Cluster**: Select the Amazon Aurora MySQL cluster to which you want to restore the database.

To overwrite any existing database with the same name, enable **Overwrite existing database(s)** and type **Yes** to confirm.

> **Note:** Restoring to a new location also replaces any existing database with the same name in the selected cluster.

8. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

9. Click **Recover**.

Amazon Aurora PostgreSQL

Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Amazon Aurora PostgreSQL clusters.

Backup Methods

You can protect your Amazon Aurora PostgreSQL using one of the following backup methods:

- **AWS Snapshot methods**: Cohesity uses this option to protect Amazon Aurora PostgreSQL clusters. For protecting the Amazon Aurora PostgreSQL clusters, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon Aurora PostgreSQL cluster.
- **Ingest method**: Cohesity uses this option to protect the Amazon Aurora PostgreSQL cluster at the database level by ingesting the backup data to the Cohesity cluster.

> **Important:**
> - Cohesity supports the protection of Amazon Aurora PostgreSQL only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster) and Cohesity Cloud Edition (legacy) on AWS.
> - On Cohesity Cloud Edition on AWS, Cohesity supports the protection of Amazon Aurora PostgreSQL at both the cluster level (via AWS snapshot-based protection) and the database level (via ingest-based protection).
> - On Cohesity Cloud Edition (legacy) on AWS, Cohesity supports the backup and recovery of Amazon Aurora PostgreSQL only at the cluster level.

This topic provides information on how you can protect Amazon Aurora PostgreSQL on Cohesity Cloud Edition on AWS.

Plan and Prepare for Amazon Aurora PostgreSQL Protection

Before protecting Amazon Aurora PostgreSQL with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon Aurora PostgreSQL with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Before You Begin

To protect Amazon Aurora PostgreSQL with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Authenticate for Ingest-based Protection

When creating a Protection Group to protect Amazon Aurora PostgreSQL using the Ingest-based protection method, you must authenticate the PostgreSQL client running on the Cohesity cluster to discover the databases of the Amazon Aurora PostgreSQL cluster. This can be done using one of the following authentication methods:

- **Username** and **password** of the Amazon Aurora PostgreSQL clusters.

  Ensure the user has the following permissions:

  - `psql -h cluster_fqdn -U postgres`
  - `CREATE USER dbuser WITH PASSWORD '<password>' CREATEDB;`
  - `GRANT rds_superuser TO dbuser;`

- **IAM-authenticated PostgreSQL user**.

  If you choose the IAM-authenticated PostgreSQL user method, ensure the following:

  - Enable IAM DB authentication on the DB cluster:
    i. From the AWS Management Console, go to RDS Console.
    ii. Select the DB cluster and click **Modify**.
    iii. Scroll to the **Database Authentication section**.
    iv. Enable **Password and IAM database authentication**.
    v. Click **Continue**.
    vi. Choose whether to **Apply Immediately** or during the next maintenance window, and then click **ModifyDBcluster.**
  - Create database user using admin login:
    - `CREATE USER <iamuser> WITH LOGIN;`
  - Grant below role to user
    - `GRANT rds_iam TO <iamuser>;`
    - `ALTER ROLE <iamuser> WITH LOGIN CREATEDB;`
    - `GRANT rds_superuser TO <iamuser>;`

For information on the steps to authenticate the Cohesity cluster to the Amazon Aurora for PostgreSQL in the Cohesity cluster UI, see Protect Amazon Aurora PostgreSQL Using Ingest Method.

Ports Used for Communication

For information on port numbers required, see AWS Cloud Database Workloads.

Incremental Backup Requirements

Cohesity supports full and differential incremental backups for Amazon Aurora PostgreSQL databases.

Incremental backups rely on PostgreSQL logical decoding and replication slots.

To enable logical decoding in Amazon Aurora PostgreSQL, configure the `rds.logical_replication` server parameter to 1 in a custom DB parameter group.

Perform the following steps in the AWS Console:

1. Create a custom DB cluster parameter group.
2. Set the `rds.logical_replication` server parameter to 1.
3. Associate the parameter group with the database instance or cluster.
4. Restart the aurora cluster writer instance to apply the changes.
5. Verify the configuration using the following SQL command:
   `SHOW wal_level;`

Permissions Required to Refresh Materialized Views

To refresh materialized views automatically or manually, the database user must have the required schema and object permissions.

For materialized view refresh:

```
GRANT USAGE ON SCHEMA schema_name TO refresh_user;

GRANT SELECT ON ALL TABLES IN SCHEMA schema_name TO refresh_user;

GRANT ALL ON materialized_view_name TO refresh_user;
```

For dependency analysis:

```
GRANT SELECT ON pg_class, pg_namespace, pg_depend, pg_rewrite TO refresh_
user;
```

Replace `` `schema_name` `` with your actual schema name.

Replace `` `refresh_user` `` with the database user configured in your Amazon Aurora PostgreSQL.

Replace `` `materialized_view_name` `` with each materialized view name.

> **Note:** If the database user does not have the required permissions, use manual refresh after recovery.

Considerations for AWS Aurora PostgreSQL Data Protection

Review and understand the following limitations and known issues before you protect Amazon Aurora PostgreSQL:

Backup Considerations

Cohesity supports full and differential incremental backups for Amazon Aurora PostgreSQL:

- **General Backup Considerations**
  - Replica server databases are not supported for full or incremental backups.
  - If a full backup fails due to temporary objects, delete the temporary objects manually and run the backup again.

- **Incremental Backup Considerations**
  - Incremental backups are not supported in the following scenarios:
  - Databases containing Large Objects.
  - The bytea_output server parameter is set to escape.
  - The backup user is a non-admin user.
  - Tables owned by other users do not have a primary key.
- **Database Rename Considerations for Incremental Backups**
  - Incremental backups rely on PostgreSQL replication slots to track database changes. Before renaming a protected database, identify the replication slot associated with the database:

    ```
    SELECT slot_name
    FROM pg_replication_slots
    WHERE database = '<database_name>';
    ```

    After renaming the database, drop the existing replication slot:

    ```
    SELECT pg_drop_replication_slot('<replication_slot_name>');
    ```

    Failure to remove stale replication slots may impact future incremental backup operations.

Restore Considerations

- Restore of Amazon Aurora PostgreSQL to an EC2 instance running PostgreSQL is not supported.
- Restoring a backup taken from a higher Aurora PostgreSQL database version to a lower version is not supported.
- Restoring security privileges is not supported.
- Restore does not fail if the owner or role does not exist on the destination.
- Post-restore, the database is associated with the role based on the credentials used for the destination cluster.
- ACLs and permissions for a database are restored only if an IAM-authenticated PostgreSQL user is attached to the cluster.
- Users must modify database ownership manually after restore.
- Point-in-time restore (PITR) is supported when recovering Amazon Aurora PostgreSQL at cluster-level.
- Cross-account restore of Amazon Aurora PostgreSQL protected using the AWS Snapshot protection type is not supported.

Protection Scope and Configuration Considerations

- Cohesity protects only user databases within Amazon Aurora PostgreSQL. System databases are not protected.
- You cannot perform ingest-based and snapshot-based protections within the same Protection Group.
- If databases of a cluster are discovered, you must select all databases of that cluster for snapshot-based protection.
- If you select clusters for which no databases are discovered, ingest-based protection cannot be performed.
- If a full backup fails because of temporary objects, delete the temporary objects manually and run the backup again.

Authentication and User Considerations

- Uppercase usernames are not supported for new users added after PostgreSQL server creation.
- SCRAM authentication configured on a database cluster is not supported.

Naming Restrictions

The following characters are not supported in database names during restore workflows:

`, @, , [,], !, #, %, ^, ., ,, &, *, (,), <, >, ?, /, |, }, {, ~, :, ', ", ;, +, =, -

Encryption Considerations

- PostgreSQL does not natively provide Transparent Data Encryption (TDE) in standard community deployments. However, PostgreSQL supports application-level and column-level encryption through cryptographic extensions and functions (for example, pgcrypto).
- In Amazon Aurora PostgreSQL, encryption at rest is typically provided through AWS-managed storage encryption mechanisms.
- During protection, backup, and recovery operations, encrypted column data is preserved in its encrypted form. After restore, applications can continue to access and decrypt the data using the appropriate encryption keys and application logic.
- Backup and restore operations do not modify or decrypt application-encrypted column data. You are responsible for maintaining access to the required encryption keys after recovery.
- When recovering encrypted Aurora cluster to the same AWS account but a different AWS region, you must create a KMS encryption key in the target region with the same alias name as the KMS key used to encrypt the source Aurora cluster.
- When performing cross-region cross-account recovery of an Amazon Aurora cluster encrypted with a customer-managed KMS key, ensure the following:

- The source account's KMS key alias is available in both the source and destination regions of the destination account.

Example:

If A1 is the source account with region R1, and A2 is the destination account with region R2, ensure the KMS key alias from A1 is available in both R1 and R2 of A2.

- The destination account has permission to access the source account's KMS key alias.
- The option group from the source account is available in both the source and destination regions of the destination account.

Backup Amazon Aurora PostgreSQL

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads including Amazon Aurora PostgreSQL associated with the specified AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **Aurora PostgreSQL** tab to view all discovered clusters.

Based on the registration level of your AWS source, the **Aurora PostgreSQL** tab displays the entity hierarchy for the Amazon Aurora PostgreSQL environment, allowing navigation from the organization or account level down to individual databases.

For information on registering AWS, see Register AWS Source.

Amazon Aurora PostgreSQL Entity Hierarchy

The **Aurora PostgreSQL** tab displays the entity hierarchy for the Aurora PostgreSQL, allowing you to navigate from the organization or account level down to individual databases.

The **Aurora PostgreSQL** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with Cohesity for Aurora PostgreSQL protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	Organizational Unit (OU) The organizational units within the AWS organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain Aurora PostgreSQL clusters.
AWS Account	The AWS account registered with Cohesity for Aurora protection. Applicable when the AWS source is registered at the account level.

Structure	Description
AWS Region	The AWS region where the Aurora resources reside.
Availability Zone	The availability zone within the region where the Aurora cluster is deployed.
Aurora Cluster	The specific Aurora PostgreSQL cluster within the selected availability zone.
Database	The individual databases hosted within the Aurora PostgreSQL cluster. The database will be listed in the entity hierarchy only if you have provided credentials to discover the databases of the Amazon Aurora PostgreSQL cluster.

Backup Scope by Aurora PostgreSQL Hierarchy

The Cohesity cluster allows you to back up Amazon Aurora PostgreSQL at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level.

You can choose to back up:

- **AWS Organization** – Protect all Amazon Aurora PostgreSQL clusters and databases across all member accounts in the registered AWS organization.
- **Organizational Unit (OU)** – Protect all Amazon Aurora PostgreSQL clusters and databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all Amazon Aurora PostgreSQL clusters and databases within a specific member account in the registered AWS organization.
- **AWS Account** – Protect all Amazon Aurora PostgreSQL clusters and databases within the registered standalone AWS account.
- **AWS Region** – Protect all Amazon Aurora PostgreSQL clusters and databases within a specific AWS Region.
- **Availability Zone** – Protect Amazon Aurora PostgreSQL clusters and databases within a specific Availability Zone.
- **Cluster** – Protect a specific Amazon Aurora PostgreSQL cluster and all databases hosted on it.
- **Database** – Protect individual databases within an Amazon Aurora PostgreSQL cluster by providing the database credentials.

Aurora PostgreSQL Backup Methods

Cohesity provides the following Protection Types for the protecting Amazon Aurora PostgreSQL:

- **AWS Snapshot method**: Cohesity uses this option to protect Amazon Aurora PostgreSQL clusters. For protecting the Amazon Aurora PostgreSQL clusters, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon Aurora PostgreSQL cluster.
- **Ingest method**: Cohesity uses this option to protect the Amazon Aurora PostgreSQL cluster at the database level by ingesting the backup data to the Cohesity cluster.

Supported Backup Types

Cohesity performs full and incremental backup of Amazon Aurora PostgreSQL. Incremental backups capture only the data changes since the last successful backup, helping reduce backup windows, optimize storage consumption, and improve backup efficiency.

Protect Amazon Aurora PostgreSQL Using Ingest Method

With the Ingest protection method, Cohesity cluster protects the Amazon Aurora PostgreSQL of your AWS source at the database level.

Authenticate Cohesity Cluster for Aurora PostgreSQL Protection

To protect the Amazon Aurora PostgreSQL of your AWS source at the database level using the Ingest method, you must provide the credentials of the Amazon Aurora PostgreSQL cluster while creating the protection. By providing the credentials, Cohesity discovers the databases as objects for protection.

Perform the following steps for authentication:

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AWS source and click into it.
3. Click the **Aurora PostgreSQL** tab.
4. Select the checkbox for the hierarchy level (AWS Organization, OUs, member account, account, region, availability zone, or cluster level), and then click **Database Credentials** above the checkboxes.

   The **Database Credentials for <cluster>** appears.
5. Select one of the following authentication methods for Cohesity to discover the databases of the Amazon Aurora PostgreSQL clusters:
   - **Credentials**: Enter the **Username**, and **Password** for the Amazon Aurora PostgreSQL clusters.
   - **IAM Authentication**: Enter the **Username** for IAM authentication.

     For information on the authentication requirements, see Authenticate for Ingest-based Protection.
6. Click **Save**.

   The databases of the Amazon Aurora PostgreSQL clusters are discovered.

Protect Amazon Aurora PostgreSQL Database

To protect your Amazon Aurora PostgreSQL using Ingest method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **Aurora PostgreSQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon Aurora PostgreSQL databases for protection using one of the following views:
   - **List View**: Use this option to select the Aurora PostgreSQL databases for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select databases for protection at various hierarchy levels.

     Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon Aurora PostgreSQL clusters.
     - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
       - **AWS Organization**
       - **Organizational Unit (OU)**
       - **Member Account**
       - **AWS Region**
       - **Availability Zone**
       - **Cluster**
     - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
       - **AWS Account**
       - **AWS Region**
       - **Availability Zone**
       - **Cluster**

     You can also exclude the Amazon Aurora PostgreSQL databases from auto-protection within an organization, OUs, memeber account, account, region, or availability zone by clicking the auto-protect icon ( ) next to the cluster. The

clusters excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the databases for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all databases with that tag for protection.

  Optionally,

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon Aurora PostgreSQL databases with this tag.
  - You can exclude the Amazon Aurora PostgreSQL databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude databases with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The Ingest-based policy is enabled by default. Create or select an existing policy from the **Policy (Cohesity Ingest)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	Batch Count	Enter the number of records to process in a single batch during a protection run. This value helps balance backup performance and resource utilization on the Aurora PostgreSQL. Higher values can improve throughput but may increase memory and database load. The default value is 1024.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Protect Amazon Aurora PostgreSQL Using AWS Snapshot Method

With the AWS Snapshot protection method, Cohesity cluster protects the Amazon Aurora PostgreSQL at the cluster level.

To protect your Amazon Aurora PostgreSQL cluster using AWS snapshot method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **Aurora PostgreSQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon Aurora PostgreSQL clusters for protection using one of the following views:
    - **List View**: Use this option to select the Aurora PostgreSQL clusters for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select clusters for protection at various hierarchy levels.

      Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new Aurora PostgreSQL clusters added to your AWS source.
        - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Organization**
            - **Organizational Unit (OU)**
            - **Member Account level**
            - **AWS Region**
            - **Availability Zone**
        - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Account**
            - **AWS Region**
            - **Availability Zone**

      You can also exclude the Amazon Aurora PostgreSQL clusters from auto-protection within an account, region, or availability zone by clicking the auto-protect icon ( ) next to the cluster. The clusters excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the Aurora PostgreSQL clusters for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all Aurora PostgreSQL clusters with that tag for protection.

  Optionally,

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon Aurora PostgreSQL clusters with this tag.
  - You can exclude the Amazon Aurora PostgreSQL clusters from auto-protection within a tag by clicking the auto-protect icon ( ) next to the cluster. The clusters excluded from auto-protection are displayed with the exclude ( ) icon.

  You can exclude clusters with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The AWS Snapshot-based policy is enabled by default. Create or select an existing policy from the **Policy (AWS Snapshot)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	Batch Count	Enter the number of records to process in a single batch during a protection run. This value helps balance backup performance and resource utilization on the Aurora PostgreSQL. Higher values can improve throughput but may increase memory and database load. The default value is 1024.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon Aurora PostgreSQL Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon Aurora PostgreSQL clusters or databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon Aurora PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Aurora PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Aurora PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Aurora PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.
7. Type **YES** to confirm that you want to delete the Protection Group.
8. Click **Delete Group**.

Recover Amazon Aurora PostgreSQL

Depending on the backup method you used for protecting Amazon Aurora PostgreSQL, you can:

Recover Amazon Aurora PostgreSQL at Cluster Level

You can recover the Amazon Aurora PostgreSQL at cluster level if you have protected the Amazon Aurora PostgreSQL using the AWS Snapshot Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon Aurora PostgreSQL Protection.

To recover Amazon Aurora PostgreSQL at cluster level:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** >**Aurora PostgreSQL**.
3. Select **Cluster** as the recovery task.
4. Search for the cluster by entering characters of the cluster name, or by specifying the wildcard character *.
5. From the list of search results, select the cluster to recover. By default, the latest snapshot of the cluster is selected for recovery. To recover from a different snapshot,

perform the following steps depending on the object selected:

From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for <*cluster_name*>** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*cluster_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. You can also select any point on the green solid line on the slider to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.
  b. Select the snapshots to be recovered.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.
- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.

6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the cluster and then provide the required information:
   - **Original Location**: Restore the cluster to the same region.

     > **Note:** If an Aurora cluster with the same name already exists, the restore will fail. Ensure the original cluster is not present or restore with a different name.

   - **New Location**: Restore the Aurora cluster to a different region. You can restore the Aurora cluster to the same or a different region within the same

AWS account as the original Aurora cluster.

Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source**: Select the AWS organization that contains the destination Aurora cluster. You can recover the Aurora cluster within the same organization as the source Aurora cluster.
  - **AWS Account**: Select the member account to which you want to restore the Aurora cluster.
  - **Region**: Select a region to restore the Aurora cluster.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered Aurora cluster.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered Aurora cluster.
  - **Availability Zone**: Select an availability zone in AWS to recover the Aurora cluster.
- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account to which you want to restore the Aurora cluster.
  - **Region**: Select a region to restore the Aurora cluster.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered Aurora cluster.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered Aurora cluster.
  - **Availability Zone**: Select an availability zone in AWS to recover the Aurora cluster.

To overwrite any existing Aurora cluster with the same name, enable **Overwrite existing Aurora cluster(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of Amazon Aurora PostgreSQL cluster is not supported.
> - If an Aurora cluster with the same name already exists, the restore will fail.

8. **Database Instance Identifier**: Specify the unique key that identifies the database cluster that will be recovered.
9. **DB Option Group**: Select an option group that contains the option you want to attach to the Aurora cluster that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
10. **DB Parameter Group**: Select the database parameter group to associate with the Aurora cluster.
11. **DB Port**: Specify the TCP/IP port that the Aurora cluster will use for application connections. The connection string of any application connecting to the Aurora cluster must specify the port number of the Aurora cluster. Both the security group applied to the Aurora cluster and your company's firewalls must allow connections to this port. The default port is 3306.
12. Configure the following Recovery Options:
    a. **Multi A-Z Deployment**: Enable this option if you want the Aurora cluster to be recovered to have a standby cluster deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the Aurora cluster to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the Aurora cluster to automatically upgrade when a new minor database engine
13. **Task name**: Optionally, change the default name of the recovery task
14. Click **Recover**.

Recover Amazon Aurora PostgreSQL at Database Level

You can recover the Amazon Aurora PostgreSQL at database level if you have protected the Amazon Aurora PostgreSQL using the Ingest Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon Aurora PostgreSQL Protection.

To recover Amazon Aurora PostgreSQL:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **Aurora PostgreSQL**.
3. Select **Database** as the recovery task.
4. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
5. From the list of search results, select the database to recover. By default, the latest snapshot of the database is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for <*database_name*>** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*database_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
     b. Select the snapshots to be recovered.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
     b. Select the snapshot to be recovered by clicking the corresponding blue dot.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:

- **Original Location**: Restore the database to the original Amazon Aurora PostgreSQL cluster.

  > **Note:** Restoring to the original location overwrites the existing database.

  Type **Yes** to confirm.

- **New Location**: Restore the database to a different Amazon Aurora PostgreSQL cluster. Cohesity supports cross-region and cross-account restores of Amazon Aurora PostgreSQL at the database level.

  Based on the registration level of the AWS source, select the recovery destination hierarchy:

  - If the AWS source is registered at the organization level:
    - **Source**: Select the AWS organization that contains the destination Amazon Aurora PostgreSQL cluster. You can recover data only within the same AWS Organization as the source.
    - **AWS Account**: Select the member account that contains the destination Amazon Aurora PostgreSQL cluster.
    - **Cluster**: Select the Amazon Aurora PostgreSQL cluster to which you want to restore the database.
  - If the AWS source is registered at the account level:
    - **AWS Account**: Select the registered AWS account that contains the destination Amazon Aurora PostgreSQL cluster.
    - **Cluster**: Select the Amazon Aurora PostgreSQL cluster to which you want to restore the database.

  a. To overwrite any existing database with the same name, enable **Overwrite existing database(s)** and type **Yes** to confirm.

  > **Note:** Restoring to a new location also replaces any existing database with the same name in the selected cluster.

8. In the **Recovery Options** section, configure the following options:

Options	Description
Refresh Materialized Views	Enable this option to automatically refresh PostgreSQL materialized views after recovery to ensure they reflect the latest restored data. Disabling this option may reduce recovery time, but materialized views can remain outdated until manually refreshed. To manually refresh , download the script from the recovery task details page, and then run it on the recovered PostgreSQL database. For more information, see What are Materialized Views.
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

9. Click **Recover**.

Amazon Aurora PostgreSQL Materialized View Refresh After Recovery

When recovering Amazon Aurora PostgreSQL databases, materialized views need to be refreshed to reflect the recovered data. This topic explains the available options for refreshing materialized views and how to choose the right approach for your environment.

What are Materialized Views

Materialized views are database objects that contain the results of a query, stored physically like a table. Unlike regular views, they need to be explicitly refreshed to reflect changes in their underlying base tables.

Why Refresh is Needed After Recovery

After a database recovery operation, materialized views contain stale data from the backup point. The views must be rebuilt to work with the newly recovered data. Without refresh, queries against materialized views will return outdated or incorrect results.

Materialized View Refresh Options

You can refresh the materialized views using one of the following options:

- Option 1: Automatic Refresh During Recovery
- Option 2: Manual Refresh After Recovery

Comparison of Refresh Options

Factor	Automatic Refresh	Manual Refresh
Database Size	Small to medium	Large
Refresh Duration	Short	Large
Recovery Time Objective	Flexible	Strict
Permissions Available	Yes	Strict
Scheduling Flexibility	Low	High
Administrative Control	Minimal	Greater

Option 1: Automatic Refresh During Recovery

During the recovery workflow, you can enable automatic materialized view refresh using the toggle available under **Recovery Options**.

When to Use

- Use this option in the following scenarios:
- Smaller databases where refresh completes quickly
- Immediate data consistency is critical
- Database credentials have required permissions to refresh materialized views
- Longer recovery duration is acceptable

Permission Verification

Before enabling automatic refresh, verify that the Amazon Aurora PostgreSQL database user has the required permissions:

- To check schema usage

```
SELECT has_schema_privilege('refresh_user', 'schema_name', 'USAGE');
```

- Check table select permissions

```
SELECT has_table_privilege('refresh_user', 'table_name', 'SELECT');
```

- Check materialized view permissions

```
SELECT has_table_privilege('refresh_user', 'materialized_view_name',
'INSERT, UPDATE, DELETE');
```

How to Enable

When initiating recovery from the Cohesity cluster UI, enable the Refresh materialized views automatically after recovery option under Recovery Options.

Option 2: Manual Refresh After Recovery

Skip automatic refresh during recovery and perform the materialized view refresh separately using the script available in the Cohesity UI.

When to Use

Use this option in the following scenarios:

- Large databases with materialized views that require significant time to refresh.
- AWS source credentials do not have sufficient permissions to refresh materialized views.
- Recovery Time Objective (RTO) requirements demand faster database availability.
- Materialized view refresh must be scheduled during off-peak hours.

Benefits of Manual Post-Recovery Refresh

- Faster Recovery Time: Your database becomes available immediately after recovery.
- Controlled Timing: Schedule the refresh during off-peak hours.
- Better Resource Management: Avoid resource contention during critical recovery operations.
- Flexibility: DB administrators can monitor and manage the refresh process independently.
- Permission Flexibility: Use a different user with appropriate permissions for refresh operations.

Permission Verification

Before refreshing the materialized view manually, verify that the Amazon Aurora PostgreSQL database user has the required permissions:

- To check schema usage

```
SELECT has_schema_privilege('refresh_user', 'schema_name', 'USAGE');
```

- Check table select permissions

```
SELECT has_table_privilege('refresh_user', 'table_name', 'SELECT');
```

- Check materialized view permissions

```
SELECT has_table_privilege('refresh_user', 'materialized_view_name',
'INSERT, UPDATE, DELETE');
```

How to Use

1. During recovery, leave the Refresh Materialized Views option disabled.
2. Complete the database recovery.
3. Download the refresh script from the Cohesity UI.
4. Run the script manually after recovery.

Download the refresh script

In the Cohesity UI:

1. Navigate to Recoveries.
2. Select the PostgreSQL recovery task.
3. Open the Options tab.
4. Locate the Materialized View Refresh section.
5. Download the refresh script.

Run the Refresh Script

1. Execute the downloaded script on the recovered database
2. The script automatically determines materialized view dependencies.
3. Materialized views are refreshed in dependency order.

Best Practices

- **Assess Your Data Size**: For databases with materialized views exceeding several GB, consider manual refresh.
- **Verify Permissions**: Before using automatic refresh, ensure DataProtect source credentials have required permissions.
- **Save the Script**: Download and save the refresh script from the recovery details page for future use.
- **Test in Non-Production**: Validate your chosen approach in test/dev environments first.
- **Schedule Appropriately**: For manual refresh, run during maintenance windows or low-traffic periods
- **Monitor Resources**: Materialized view refresh can be resource-intensive; ensure adequate system resources.
- **Regular Maintenance**: Consider implementing periodic refresh schedules (using tools like `pg_cron`) independent of recovery operations.
- **Document Your Views**: Maintain documentation of materialized views and their dependencies.

Amazon Relational Database Service for MySQL

Amazon Relational Database Service (RDS) for MySQL is a fully managed, cloud-based relational database service offered by Amazon Web Services (AWS). Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Amazon Relational Database Service for MySQL instances.

Backup Methods

You can protect your Amazon RDS for MySQL using one of the following backup methods:

- **AWS Snapshot methods**: Cohesity uses this option to protect Amazon RDS for MySQL instances. For protecting the Amazon RDS for MySQL instances, Cohesity leverages the AWS native snapshots and stores them in the same AWS account and region as the source Amazon RDS for MySQL instance.
- **Ingest method**: Cohesity uses this option to protect the Amazon RDS for MySQL instance at the database level by ingesting the backup data to the Cohesity cluster.

> **Important:**
> - Cohesity supports the protection of Amazon RDS for MySQL only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster) and Cohesity Cloud Edition (legacy) on AWS.
> - On Cohesity Cloud Edition on AWS, Cohesity supports the protection of Amazon RDS for MySQL at both the instance level (via AWS snapshot-based protection) and the database level (via ingest-based protection).
> - On Cohesity Cloud Edition (legacy) on AWS, Cohesity supports the backup and recovery of Amazon RDS for MySQL only at the instance level.

This topic provides information on how you can protect Amazon RDS for MySQL on Cohesity Cloud Edition on AWS.

Plan and Prepare for Amazon RDS for MySQL Protection

Before protecting Amazon RDS for MySQL with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon RDS for MySQL with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Authenticate for Ingest-based Protection

If you are planning to protect the Amazon RDS for MySQL using Ingest-based protection method, you must authenticate the Cohesity cluster to discover the databases of the Amazon RDS for MySQL instance using one of the following methods:

<table>
<tr><th>Authentication Method</th><th>Permissions</th><th>Example Command</th></tr>
<tr><td>Username and password of the Amazon RDS MySQL</td><td rowspan="2">SELECT
INSERT
UPDATE
DELETE
CREATE
DROP
RELOAD
PROCESS
REFERENCES
INDEX
ALTER
SHOW DATABASES
LOCK TABLES
CREATE VIEW
SHOW VIEW
CREATE ROUTINE
ALTER ROUTINE
SHOW ROUTINE
EVENT
TRIGGER
SET_USER_ID</td><td>Create the Database user:
<code>CREATE USER 'dbuser'@'%' IDENTIFIED BY 'StrongPassword!';</code>
Grant below permission to user:
<code>GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, PROCESS, REFERENCES, INDEX, ALTER, SHOW DATABASES, LOCK TABLES, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, SHOW ROUTINE, EVENT, TRIGGER, SET_USER_ID

FLUSH PRIVILEGES;</code></td></tr>
<tr><td>IAM-authenticated MySQL user</td><td>Create the Database user, using the master login:
<code>CREATE USER <iamuser> IDENTIFIED WITH AWSAuthenticationPlugin as 'RDS';</code>
Grant below permission to user:
<code>GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, PROCESS, REFERENCES, INDEX, ALTER, SHOW DATABASES, LOCK TABLES, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, SHOW_ROUTINE, EVENT,TRIGGER, SET_USER_ID
ON . TO `<iamuser>`@'%'
WITH GRANT OPTION;
FLUSH PRIVILEGES;</code></td></tr>
</table>

For information on the steps to authenticate the Cohesity cluster to the Amazon RDS for MySQL in the Cohesity cluster UI, see Protect Amazon RDS for MySQL Using Ingest Method.

Ports Used for Communication

For information on port numbers required, see AWS Cloud Database Workloads.

Considerations for AWS RDS for MySQL Data Protection

Review and understand the following considerations before you protect Amazon RDS for MySQL:

- Cohesity performs only full backup of Amazon RDS for MySQL.
- Cohesity protects only the user databases within the Amazon RDS for MySQL. Protection of system databases are not supported.
- The backup and recovery of Amazon RDS for MySQL instance at database level is supported only on Cohesity Cloud Edition on AWS.
- Restoring a backup taken from a higher RDS MySQL database version to a lower version is not supported.
- Point-in-time restore (PITR) is supported when recovering Amazon RDS for MySQL at instance-level.
- You cannot perform ingest-based and snapshot-based protections within the same Protection Group.
- If the databases of an instance are discovered, then to perform snapshot-based protection, you must select all the databases of that instance.
- If you select instances for which no databases are discovered, then you cannot perform ingest-based protection.
- Cross-account restore of Amazon RDS for MySQL instance protected using the AWS Snapshot protection type is not supported.
- When recovering encrypted Amazon RDS instance to the same AWS account but a different AWS region, you must create a KMS encryption key in the target region with the same alias name as the KMS key used to encrypt the source Amazon RDS instance.
- When performing cross-region cross-account recovery of an Amazon RDS instance encrypted with a customer-managed KMS key, ensure the following:
  - The source account's KMS key alias is available in both the source and destination regions of the destination account.

  Example:

  If A1 is the source account with region R1, and A2 is the destination account with region R2, ensure the KMS key alias from A1 is available in both R1 and R2 of A2.

  - The destination account has permission to access the source account's KMS key alias.

- The option group from the source account is available in both the source and destination regions of the destination account.

Backup Amazon RDS for MySQL

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads including Amazon RDS for MySQL associated with the registered AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **RDS MySQL** tab to view all discovered instances.

The **RDS MySQL** tab displays the entity hierarchy for the Amazon RDS for MySQL environment, allowing navigation from the organization or account level down to individual databases. For information on registering AWS, see Register AWS Source.

RDS MySQL Entity Hierarchy

The **RDS MySQL** tab displays the entity hierarchy for the RDS instance, allowing you to navigate from the organization or account level down to individual databases.

The **RDS MySQL** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with Cohesity for protecting Amazon RDS for MySQL. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The organizational units within the AWS Organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain Amazon RDS for MySQL.
AWS Account	The AWS account registered with Cohesity for RDS protection. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the RDS resources reside.
Availability Zone	The availability zone within the region where the RDS instance is deployed.
RDS Instance	The specific RDS MySQL instance within the selected availability zone.
Database	The individual databases hosted within the RDS MySQL instance.

Backup Scope by RDS MySQL Hierarchy

The Cohesity cluster allows you to back up Amazon RDS for MySQL at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level.

- **AWS Organization** – Protect all RDS for MySQL instances and databases across all member accounts in the registered AWS Organization.
- **Organizational Unit (OU)** – Protect all RDS for MySQL instances and databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all RDS for MySQL instances and databases within a specific member account in the registered AWS Organization.
- **AWS Account** – Protect all RDS for MySQL instances and databases within the registered standalone AWS account.
- **AWS Region** – Protect all RDS for MySQL instances and databases within a specific AWS region.
- **Availability Zone** – Protect RDS for MySQL instances and databases within a specific Availability Zone.
- **RDS Instance** – Protect a specific RDS for MySQL instance and all databases hosted on it.
- **Database** – Protect individual databases within an RDS for MySQL instance by providing the database credentials.

RDS for MySQL Backup Methods

Cohesity provides the following Protection Types for protecting Amazon RDS MySQL:

- **AWS Snapshot method**: Cohesity uses this option to protect Amazon RDS for MySQL instances. For protecting the Amazon RDS for MySQL instances, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon RDS for MySQL instance.
- **Ingest method**: Cohesity uses this option to protect the Amazon RDS for MySQL instance at the database level by ingesting the backup data to the Cohesity cluster.

Supported Backup Types

Cohesity performs only full backup of Amazon RDS for MySQL.

Protect Amazon RDS for MySQL Using Ingest Method

With the Ingest protection method, Cohesity cluster protects the Amazon RDS for MySQL at the database level.

Authenticate for RDS for MySQL Protection

To protect the Amazon RDS for MySQL at the database level using the Ingest method, you must provide the credentials of the Amazon RDS for MySQL instance while creating the protection. By providing the credentials, Cohesity discovers the databases as objects for protection.

Perform the following steps for authentication:

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AWS source and click into it.
3. Click the **RDS MySQL** tab.
4. Select the checkbox for the hierarchy level (AWS Organization, OUs, account, region, availability zone, or instance level), and then click **Database Credentials** above the checkboxes.

   The Database Credentials for <cluster> appears.
5. Select one of the following authentication methods for Cohesity to discover the databases of the Amazon RDS for MySQL instances:
   - **Credentials**: Enter the **Username**, and **Password** for the Amazon RDS for MySQL instances.
   - **IAM Authentication**: Enter the **Username** for IAM authentication.

   For information on the authentication requirements, see Authenticate for Ingest-based Protection.
6. Click **Save**.

The databases of the Amazon RDS for MySQL instances are discovered.

Protect Amazon RDS for MySQL Database

To protect your Amazon RDS for MySQL using Ingest method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS MySQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for MySQL databases for protection using one of the following views:
   - **List View**: Use this option to select the RDS for MySQL databases for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select databases for protection at various hierarchy levels.

     Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon RDS for MySQL instances.
     - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:

- **AWS Organization**
- **Organizational Unit (OU)**
- **Member Account**
- **AWS Region**
- **Availability Zone**
- **RDS Instance**

- If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
  - **AWS Account**
  - **AWS Region**
  - **Availability Zone**

    RDS Instance

  You can also exclude the Amazon RDS for MySQL databases from auto-protection within an Organization, OU, member account, account, region, or availability zone by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the databases for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all databases with that tag for protection.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for MySQL databases with this tag.
  - You can exclude the Amazon RDS for MySQL databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude databases with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.

7. The Ingest-based policy is enabled by default. Create or select an existing policy from the **Policy (Cohesity Ingest)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Protect Amazon RDS for MySQL Using AWS Snapshot Method

With the AWS Snapshot protection method, Cohesity cluster protects the Amazon RDS MySQL of your AWS account at the instance level.

To protect your Amazon RDS for MySQL instance using AWS snapshot method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS MySQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for MySQL instances for protection using one of the following views:
    - **List View**: Use this option to select the RDS for MySQL instances for protection from a flat, searchable list.

- **Hierarchy View**: Use this option to select instances for protection at various hierarchy levels.

  Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new RDS for MySQL instances added to your AWS source.

  - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
    - **AWS Organization**
    - **Organizational Unit (OU)**
    - **Member Account**
    - **AWS Region**
    - **Availability Zone**
  - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
    - **AWS Account**
    - **AWS Region**
    - **Availability Zone**

  You can also exclude the Amazon RDS for MySQL instances from auto-protection within an account, region, or availability zone by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the instances for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all instances with that tag for protection.

  Optionally,

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for MySQL instances with this tag.
  - You can exclude the Amazon RDS for MySQL instances from auto-protection within a tag by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude instances with a specific tag from auto-protection by

clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The AWS Snapshot-based policy is enabled by default. Create or select an existing policy from the **Policy (AWS Snapshot)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon RDS for MySQL Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon RDS for MySQL instances or databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon RDS for MySQL.

3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.

- Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.

5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon RDS for MySQL

Depending on the backup method you used for protecting Amazon RDS for MySQL, you can:

Recover Amazon RDS for MySQL at Instance Level

Recover Amazon RDS for MySQL at Database Level

Recover Amazon RDS for MySQL at Instance Level

You can recover the Amazon RDS for MySQL at instance level if you have protected the Amazon RDS for MySQL using the AWS Snapshot Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon RDS for MySQL Protection.

To recover Amazon RDS for MySQL:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS MySQL**.

1. Select **Instance** as the recovery task.
2. Search for the instance by entering characters of the instance name, or by specifying the wildcard character *.
3. From the list of search results, select the instance to recover. By default, the latest snapshot of the instance is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for <*instance_name*>** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:

     a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*instance_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified

range. You can also select any point on the green solid line on the slider to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.

b. Select the snapshots to be recovered.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

- To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

4. Click **Next: Recover Options**.
5. In the **Recover To** section, select any one of the following options based on where you want to restore the instance and then provide the required information:

- **Original Location**: Restore the instance to the same region.

> **Note:** If an RDS instance with the same name already exists, the restore will fail. Ensure the original instance is not present or restore with a different name.

Type **Yes** to confirm.

- **New Location**: Restore the instance to a different region. You can restore the instance to the same or a different region within the same account as the original RDS instance.

Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source**: Select the AWS organization that contains the destination instance. You can recover the instance within the same organization as the source instance.
  - **AWS Account**: Select the member account to which you want to restore the instance.

- **Region**: Select a region to restore the instance.
- **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
- **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
- **Availability Zone**: Select an availability zone in AWS to recover the RDS.

- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.

To overwrite any existing instance with the same name, enable **Overwrite existing instance(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of RDS instance protected using the AWS Snapshot protection type is not supported.
> - If an RDS instance with the same name already exists, the restore will fail.

6. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
7. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
8. **DB Parameter Group**: Select the database (DB) parameter group to associate with the DB instance.
9. **DB Port:** Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port. The

default port is 3306.

10. Configure the following **Recovery Options**:
    a. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.
11. **Task name**: Optionally, change the default name of the recovery task.
12. Click **Recover**.

Recover Amazon RDS for MySQL at Database Level

You can recover the Amazon RDS for MySQL at database level if you have protected the Amazon RDS for MySQL using the RDS Ingest Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon RDS for MySQL Protection.

To recover Amazon RDS for MySQL:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS MySQL**.
3. Select **Database** as the recovery task.
4. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
5. From the list of search results, select the database to recover. By default, the latest snapshot of the database is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for <*database_name*>** page appears. By default, the option

to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*database_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
  b. Select the snapshots to be recovered.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.
- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.

6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:
   - **Original Location**: Restore the database to the original Amazon RDS for MySQL instance.

     > **Note:** Restoring to the original location overwrites the existing database.

     Type **Yes** to confirm.
   - **New Location**: Restore the database to a different Amazon RDS for MySQL instance. Cohesity supports cross-region and cross-account restores of Amazon RDS for MySQL at the database level.

     Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source**: Select the AWS organization that contains the destination Amazon RDS for MySQL instance. You can recover data only within the same AWS Organization as the source.
  - **AWS Account**: Select the member account that contains the destination Amazon RDS for MySQL instance.
  - **Instance**: Select the Amazon RDS for MySQL instance to which you want to restore the database.
- If the AWS source is registered at the account level:
  - **AWS Account**: Select the registered AWS account that contains the destination Amazon RDS for MySQL instance.
  - **Instance**: Select the Amazon RDS for MySQL instance to which you want to restore the database.

To overwrite any existing database with the same name, enable **Overwrite existing database(s)** and type **Yes** to confirm.

> **Note:** Restoring to a new location also replaces any existing database with the same name in the selected instance.

8. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

9. Click **Recover**.

Amazon Relational Database Service for PostgreSQL

Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Amazon RDS for PostgreSQL instances.

Backup Methods

You can protect your Amazon RDS for PostgreSQL using one of the following backup methods:

- **AWS Snapshot methods**: Cohesity uses this option to protect Amazon RDS for PostgreSQL instances. For protecting the Amazon RDS for PostgreSQL instances, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon RDS for PostgreSQL instance.
- **Ingest method**: Cohesity uses this option to protect the Amazon RDS for PostgreSQL instance at the database level by ingesting the backup data to the Cohesity cluster.

> **Important:**
> - Cohesity supports the protection of Amazon RDS for PostgreSQL only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster) and Cohesity Cloud Edition (legacy) on AWS.
> - On Cohesity Cloud Edition on AWS, Cohesity supports the protection of Amazon RDS for PostgreSQL at both the instance level (via AWS snapshot-based protection) and the database level (via ingest-based protection).
> - On Cohesity Cloud Edition (legacy) on AWS, Cohesity supports the backup and recovery of Amazon RDS for PostgreSQL only at the instance level.

This topic provides information on how you can protect Amazon RDS for PostgreSQL on Cohesity Cloud Edition on AWS.

Plan and Prepare for Amazon RDS for PostgreSQL Protection

Before protecting Amazon RDS for PostgreSQL with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon RDS for PostgreSQL with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Authenticate for Ingest-based Protection

When creating a Protection Group to protect Amazon RDS for PostgreSQL using the Ingest-based protection method, you must authenticate the PostgreSQL client running on the Cohesity cluster to discover the databases of the Amazon RDS for PostgreSQL instance. This can be done using one of the following authentication methods:

- **Username** and **password** of the Amazon RDS for PostgreSQL instances.

  Ensure the user has the following permissions:

- `psql -h instance_fqdn -U postgres`
- `CREATE USER dbuser WITH PASSWORD '<password>' CREATEDB;`
- `GRANT rds_superuser TO dbuser;`

- **IAM-authenticated PostgreSQL user**.

  If you choose the IAM-authenticated PostgreSQL user method, ensure the following:

  - Enable IAM DB authentication on the RDS DB instance:
    i. From the AWS Management Console, go to RDS Console.
    ii. Select the DB instance and click **Modify**.
    iii. Scroll to the **Database Authentication** section.
    iv. Enable **Password and IAM database authentication**.
    v. Click **Continue**.
    vi. Choose **Apply Immediately**, and then click **ModifyDBinstance.**
  - Create database user using admin login:
    - `CREATE USER <iamuser> WITH LOGIN;`
  - Grant the following role to user:
    - `GRANT rds_iam TO <iamuser>;`
    - `ALTER ROLE <iamuser> WITH LOGIN CREATEDB;`
    - `GRANT rds_superuser TO <iamuser>;`

Incremental Backup Requirements

Cohesity supports full and differential incremental backups for Amazon RDS for PostgreSQL databases.

Incremental backups rely on PostgreSQL logical decoding and replication slots.

To enable logical decoding in Amazon RDS PostgreSQL and Aurora PostgreSQL, configure the `rds.logical_replication` server parameter to 1 in a custom DB parameter group.

Perform the following steps in the AWS Console:

1. Create a custom DB parameter group.
2. Set the `rds.logical_replication` server parameter to 1.
3. Associate the parameter group with the database instance or cluster.
4. Restart the RDS instance to apply the changes.
5. Verify the configuration using the following SQL command:

   ```
   SHOW wal_level;
   ```

Permissions Required to Refresh Materialized Views

To refresh materialized views automatically or manually, the database user must have the required schema and object permissions.

For materialized view refresh:

```
GRANT USAGE ON SCHEMA schema_name TO refresh_user;

GRANT SELECT ON ALL TABLES IN SCHEMA schema_name TO refresh_user;

GRANT ALL ON materialized_view_name TO refresh_user;
```

For dependency analysis:

```
GRANT SELECT ON pg_class, pg_namespace, pg_depend, pg_rewrite TO refresh_
user;
```

Replace `schema_name` with your actual schema name.

Replace `refresh_user` with the database user configured in your Amazon RDS for PostgreSQL.

Replace `materialized_view_name` with each materialized view name.

> **Note:** If the database user does not have the required permissions, use manual refresh after recovery.

Ports Used for Communication

For information on port numbers required, see AWS Cloud Database Workloads.

Considerations for AWS RDS for PostgreSQL Data Protection

Review and understand the following limitations and known issues before you protect Amazon RDS for PostgreSQL:

Backup Considerations

Cohesity supports full and differential incremental backups for Amazon RDS for PostgreSQL databases.

- **General Backup Considerations**
  - Replica server databases are not supported for full or incremental backups.
  - If a full backup fails due to temporary objects, delete the temporary objects manually and run the backup again.
- **Incremental Backup Considerations**
  - Incremental backups are not supported in the following scenarios:
  - Databases containing large objects.

- The bytea_output server parameter is set to escape.
- The backup user is a non-admin user.
- Tables owned by other users do not have a primary key.

- **Database Rename Considerations for Incremental Backups**
  - Incremental backups rely on PostgreSQL replication slots to track database changes. Before renaming a protected database, identify the replication slot associated with the database:

    ```
    SELECT slot_name

    FROM pg_replication_slots

    WHERE database = '<database_name>';
    ```

    After renaming the database, drop the existing replication slot:

    ```
    SELECT pg_drop_replication_slot('<replication_slot_name>');
    ```

  - Failure to remove stale replication slots may impact future incremental backup operations.

Restore Considerations

- After restoring an incremental backup, the last value of the sequence may not be consistent.
- Restore of Amazon RDS for PostgreSQL to an EC2 instance running PostgreSQL is not supported.
- Restoring a backup taken from a higher RDS PostgreSQL database version to a lower version is not supported.
- Restoring security privileges is not supported.
- Point-in-time restore (PITR) is supported when recovering Amazon RDS for PostgreSQL at instance-level.
- Restore does not fail if the owner or role does not exist on the destination.
- Post-restore, the database is associated with the role based on the credentials used for the destination instance.
- ACLs and permissions for a database are restored only if an IAM-authenticated PostgreSQL user is attached to the instance.
- Users must modify database ownership manually after restore.
- Cross-account restore of Amazon RDS for PostgreSQL instance protected using the AWS Snapshot protection type is not supported.

Protection Scope and Configuration Considerations

- Cohesity protects only user databases within Amazon RDS for PostgreSQL. System databases are not protected.
- You cannot perform ingest-based and snapshot-based protections within the same Protection Group.
- If databases of an instance are discovered, you must select all databases of that instance for snapshot-based protection.
- If you select instances for which no databases are discovered, ingest-based protection cannot be performed.

Authentication and User Considerations

- Uppercase usernames are not supported for new users added after PostgreSQL server creation.
- SCRAM authentication configured on a database instance is not supported.

Naming Restrictions

The following characters are not supported in database names during restore workflows:

`, @, \, [,], !, #, %, ^, ., ,, &, *, (,), <, >, ?, /, |, }, {, ~, :, ', ", ;, +, =, -

Encryption Considerations

- PostgreSQL does not natively provide Transparent Data Encryption (TDE) in standard community deployments. However, PostgreSQL supports application-level and column-level encryption through cryptographic extensions and functions (for example, pgcrypto).
- In Amazon RDS for PostgreSQL, encryption at rest is typically provided through AWS-managed storage encryption mechanisms.
- During protection, backup, and recovery operations, encrypted column data is preserved in its encrypted form. After restore, applications can continue to access and decrypt the data using the appropriate encryption keys and application logic.
- Backup and restore operations do not modify or decrypt application-encrypted column data. You are responsible for maintaining access to the required encryption keys after recovery.
- When recovering encrypted RDS intance to the same AWS account but a different AWS region, you must create a KMS encryption key in the target region with the same alias name as the KMS key used to encrypt the source RDS intance .
- When performing cross-region cross-account recovery of an Amazon RDS intance encrypted with a customer-managed KMS key, ensure the following:
    - The source account’s KMS key alias is available in both the source and destination regions of the destination account.

Example:

If A1 is the source account with region R1, and A2 is the destination account with region R2, ensure the KMS key alias from A1 is available in both R1 and R2 of A2.

- The destination account has permission to access the source account's KMS key alias.
- The option group from the source account is available in both the source and destination regions of the destination account.

Backup Amazon RDS for PostgreSQL

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads including Amazon RDS for PostgreSQL associated with the specified AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS account source, each workload type is displayed on its own tab. Select the **RDS PostgreSQL** tab to view all discovered instances.

The **RDS PostgreSQL** tab displays the entity hierarchy for the Amazon RDS for PostgreSQL environment, allowing you to navigate from the account level down to individual databases. For information on registering AWS, see Register AWS Source.

The **RDS PostgreSQL** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS Organization registered with Cohesity for RDS PostgreSQL protection. Applicable when the AWS source is registered at the organization level.
Organization Unit	The organizational units within the AWS Organization that contain member accounts.
Member Account	The AWS member accounts within the Organization that contain RDS PostgreSQL instances.
AWS Account	The AWS account registered with Cohesity for RDS protection. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the RDS resources reside.
Availability Zone	The availability zone within the region where the RDS instance is deployed.
RDS Instance	The specific RDS PostgreSQL instance within the selected availability zone.
Database	The individual databases hosted within the RDS PostgreSQL instance. The database will be listed in the entity hierarchy only if you have provided credentials to discover the databases of the Amazon RDS for PostgreSQL instance.

Backup Scope by RDS PostgreSQL Hierarchy

The Cohesity cluster allows you to back up Amazon RDS for PostgreSQL at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level.

You can choose to back up:

- **AWS Organization** – Protect all RDS for PostgreSQL instances and databases across all member accounts in the registered AWS Organization.
- **Organizational Unit (OU)** – Protect all RDS for PostgreSQL instances and databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all RDS for PostgreSQL instances and databases within a specific member account in the registered AWS organization.
- **AWS Account** – Protect all RDS for PostgreSQL instances and databases within the registered standalone AWS account.
- **AWS Region** – Protect all RDS for PostgreSQL instances and databases within a specific AWS region.
- **Availability Zone** – Protect RDS for PostgreSQL instances and databases within a specific Availability Zone.
- **RDS Instance** – Protect a specific RDS for PostgreSQL instance and all databases hosted on it.
- **Database** – Protect individual databases within an RDS for PostgreSQL instance by providing the database credentials.

RDS for PostgreSQL Backup Methods

Cohesity provides the following Protection Types for protecting the Amazon RDS PostgreSQL:

- **AWS Snapshot method**: Cohesity uses this option to protect Amazon RDS for PostgreSQL instances. For protecting the Amazon RDS for PostgreSQL instances, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon RDS for PostgreSQL instance.
- **Ingest method**: Cohesity uses this option to protect the Amazon RDS for PostgreSQL instance at the database level by ingesting the backup data to the Cohesity cluster.

Supported Backup Types

Cohesity performs full and incremental backup of Amazon RDS for PostgreSQL. Incremental backups capture only the data changes since the last successful backup, helping reduce backup windows, optimize storage consumption, and improve backup efficiency.

Protect Amazon RDS for PostgreSQL Using Ingest Method

With the Ingest protection method, Cohesity cluster protects the Amazon RDS PostgreSQL of your AWS source at the database level.

Authenticate for RDS for PostgreSQL Protection

To protect the Amazon RDS PostgreSQL at the database level using the Ingest method, you must provide the credentials of the Amazon RDS PostgreSQL instance while creating the protection. By providing the credentials, Cohesity discovers the databases as objects for protection.

Perform the following steps for authentication:

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AWS source and click into it.
3. Click the **RDS PostgreSQL** tab.
4. Select the checkbox for the hierarchy level (AWS Organization, OUs, member account, account, region, availability zone, or instance level), and then click **Database Credentials** above the checkboxes.

   The **Database Credentials for <instance>** appears.
5. Select one of the following authentication methods for Cohesity to discover the databases of the Amazon RDS for PostgreSQL instances:
   - **Credentials**: Enter the **Username**, and **Password** for the Amazon RDS for PostgreSQL instances.
   - **IAM Authentication**: Enter the **Username** for IAM authentication.

     For information on the authentication requirements, see Authenticate for Ingest-based Protection.
6. Click **Save**.

   The databases of the Amazon RDS for PostgreSQL instances are discovered.

Protect Amazon RDS for PostgreSQL Database

To protect your Amazon RDS for PostgreSQL using Ingest method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS PostgreSQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for PostgreSQL databases for protection using one of the following views:

- **List View**: Use this option to select the RDS for PostgreSQL databases for protection from a flat, searchable list.
- **Hierarchy View**: Use this option to select databases for protection at various hierarchy levels.

  Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon RDS for PostgreSQL instances.

  - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
    - **AWS Organization**
    - **Organizational Unit (OU)**
    - **Member Account**
    - **AWS Region**
    - **Availability Zone**
    - **RDS Instance**
  - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
    - **AWS Account**
    - **AWS Region**
    - **Availability Zone**
    - **RDS Instance**

  You can also exclude the Amazon RDS for PostgreSQL databases from auto-protection within an account, region, availability zone, or instance by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
- **Tags**: Use this option to select the databases for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all databases with that tag for protection.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for PostgreSQL databases with this tag.

- You can exclude the Amazon RDS for PostgreSQL databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
- You can exclude databases with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The Ingest-based policy is enabled by default. Create or select an existing policy from the **Policy (Cohesity Ingest)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	Batch Count	Enter the number of records to process in a single batch during a protection run. This value helps balance backup performance and resource utilization on the PostgreSQL instance. Higher values can improve throughput but may increase memory and database load. The default value is **1024**.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Protect Amazon RDS for PostgreSQL Using AWS Snapshot Method

With the AWS Snapshot protection method, Cohesity cluster protects the Amazon RDS PostgreSQL of your AWS account at the instance level.

To protect your Amazon RDS for PostgreSQL instance using AWS snapshot method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS PostgreSQL**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for PostgreSQL instances for protection using one of the following views:
    - **List View**: Use this option to select the RDS for PostgreSQL instances for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select instances for protection at various hierarchy levels.

      Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new RDS for PostgreSQL instances added to your AWS source.
        - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Organization**
            - **Organizational Unit (OU)**
            - **Member Account level**
            - **AWS Region**
            - **Availability Zone**
        - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Account**
            - **AWS Region**
            - **Availability Zone**

      You can also exclude the Amazon RDS for PostgreSQL instances from auto-protection within an account, region, or availability zone by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the instances for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all instances with that tag for protection.

  Optionally,

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for PostgreSQL instances with this tag.
  - You can exclude the Amazon RDS for PostgreSQL instances from auto-protection within a tag by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude instances with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The AWS Snapshot-based policy is enabled by default. Create or select an existing policy from the **Policy (AWS Snapshot)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	Batch Count	Enter the number of records to process in a single batch during a protection run. This value helps balance backup performance and resource utilization on the PostgreSQL instance. Higher values can improve throughput but may increase memory and database load. The default value is 1024.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon RDS for PostgreSQL Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon RDS for PostgreSQL instances or databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon RDS for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon RDS for PostgreSQL

Depending on the backup method you used for protecting Amazon RDS for PostgreSQL, you can:

- Recover Amazon RDS for PostgreSQL at Instance Level
- Recover Amazon RDS for PostgreSQL at Database Level

Recover Amazon RDS for PostgreSQL at Instance Level

You can recover the Amazon RDS for PostgreSQL at instance level if you have protected the Amazon RDS for PostgreSQL using the AWS Snapshot Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon RDS for PostgreSQL Protection.

To recover Amazon RDS for PostgreSQL instance:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS PostgreSQL**.
1. Select **Instance** as the recovery task.
2. Search for the instance by entering characters of the instance name, or by specifying the wildcard character *.
3. From the list of search results, select the instance to recover. By default, the latest snapshot of the instance is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for *<instance_name>*** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the **Edit recovery point for *<instance_name>*** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. You can also select any point on the green solid line on the slider to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.
  b. Select the snapshots to be recovered.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.
- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.

4. Click **Next: Recover Options**.
5. In the **Recover To** section, select any one of the following options based on where you want to restore the instance and then provide the required information:
   - **Original Location**: Restore the instance to the same region.

     > **Note:** If an RDS instance with the same name already exists, the restore will fail. Ensure the original instance is not present or restore with a different name.

     Type **Yes** to confirm.
   - **New Location**: Restore the instance to a different region. You can restore the instance to the same or a different region within the same account as the

original RDS instance.

Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source**: Select the AWS organization that contains the destination instance. You can recover the instance within the same organization as the source instance.
  - **AWS Account**: Select the member account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.
- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.

To overwrite any existing instance with the same name, enable **Overwrite existing instance(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of RDS instance protected using the AWS Snapshot protection type is not supported.
> - If an RDS instance with the same name already exists, the restore will fail.

6. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
7. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
8. **DB Parameter Group**: Select the database (DB) parameter group to associate with the DB instance.
9. **DB Port:** Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port. The default port is 3306.
10. Configure the following **Recovery Options**:
    a. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.
11. **Task name**: Optionally, change the default name of the recovery task.
12. Click **Recover**.

Recover Amazon RDS for PostgreSQL at Database Level

You can recover the Amazon RDS for PostgreSQL at database level if you have protected the Amazon RDS for PostgreSQL using the RDS Ingest Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon RDS for PostgreSQL Protection.

To recover Amazon RDS for PostgreSQL:

1. Navigate to **Data Protection** > **Recoveries**.
2. Select **Database** as the recovery task.
3. Click **Recover** at the top right of the page and then select **AWS** > **RDS PostgreSQL**.
4. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
5. From the list of search results, select the database to recover. By default, the latest snapshot of the database is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for <*database_name*>** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*database_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
     b. Select the snapshots to be recovered.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
     b. Select the snapshot to be recovered by clicking the corresponding blue dot.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:
   - **Original Location**: Restore the database to the original Amazon RDS for PostgreSQL instance.

> **Note:** Restoring to the original location overwrites the existing database.

Type **Yes** to confirm.

- **New Location**: Restore the database to a different Amazon RDS for PostgreSQL instance. Cohesity supports cross-region and cross-account restores of Amazon RDS for PostgreSQL at the database level.

  Based on the registration level of the AWS source, select the recovery destination hierarchy:

  - If the AWS source is registered at the organization level:
    - **Source**: Select the AWS Organization that contains the destination Amazon RDS for PostgreSQL instance. You can recover data only within the same AWS Organization as the source.
    - **AWS Account**: Select the member account that contains the destination Amazon RDS for PostgreSQL instance.
    - **Instance**: Select the Amazon RDS for PostgreSQL instance to which you want to restore the database.
  - If the AWS source is registered at the account level:
    - **AWS Account**: Select the registered AWS account that contains the destination Amazon RDS for PostgreSQL instance.
    - **Instance**: Select the Amazon RDS for PostgreSQL instance to which you want to restore the database.

  To overwrite any existing database with the same name, enable **Overwrite existing database(s)** and type **Yes** to confirm.

  > **Note:** Restoring to a new location also replaces any existing database with the same name in the selected instance.

8. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

9. Click **Recover**.

Amazon RDS for PostgreSQL Materialized View Refresh After Recovery

When recovering Amazon RDS for PostgreSQL databases, materialized views need to be refreshed to reflect the recovered data. This topic explains the available options for refreshing materialized views and how to choose the right approach for your environment.

What are Materialized Views

Materialized views are database objects that contain the results of a query, stored physically like a table. Unlike regular views, they need to be explicitly refreshed to reflect changes in their underlying base tables.

Why Refresh is Needed After Recovery

After a database recovery operation, materialized views contain stale data from the backup point. The views must be rebuilt to work with the newly recovered data. Without refresh, queries against materialized views will return outdated or incorrect results.

Materialized View Refresh Options

You can refresh the materialized views using one of the following options:

- Option 1: Automatic Refresh During Recovery
- Option 2: Manual Refresh After Recovery

Comparison of Refresh Options

Factor	Automatic Refresh	Manual Refresh
Database Size	Small to medium	Large
Refresh Duration	Short	Large
Recovery Time Objective	Flexible	Strict
Permissions Available	Yes	Strict
Scheduling Flexibility	Low	High
Administrative Control	Minimal	Greater

Option 1: Automatic Refresh During Recovery

During the recovery workflow, you can enable automatic materialized view refresh using the toggle available under **Recovery Options**.

When to Use

- Use this option in the following scenarios:
- Smaller databases where refresh completes quickly
- Immediate data consistency is critical
- Database credentials have required permissions to refresh materialized views
- Longer recovery duration is acceptable

Permission Verification

Before enabling automatic refresh, verify that the Amazon RDS for PostgreSQL database user has the required permissions:

- To check schema usage

```
SELECT has_schema_privilege('refresh_user', 'schema_name', 'USAGE');
```

- Check table select permissions

```
SELECT has_table_privilege('refresh_user', 'table_name', 'SELECT');
```

- Check materialized view permissions

```
SELECT has_table_privilege('refresh_user', 'materialized_view_name',
'INSERT, UPDATE, DELETE');
```

How to Enable

When initiating recovery from the Cohesity cluster UI, enable the Refresh materialized views automatically after recovery option under Recovery Options.

Option 2: Manual Refresh After Recovery

Skip automatic refresh during recovery and perform the materialized view refresh separately using the script available in the Cohesity UI.

When to Use

Use this option in the following scenarios:

- Large databases with materialized views that require significant time to refresh.
- AWS source credentials do not have sufficient permissions to refresh materialized views.
- Recovery Time Objective (RTO) requirements demand faster database availability.
- Materialized view refresh must be scheduled during off-peak hours.

Benefits of Manual Post-Recovery Refresh

- Faster Recovery Time: Your database becomes available immediately after recovery.
- Controlled Timing: Schedule the refresh during off-peak hours.
- Better Resource Management: Avoid resource contention during critical recovery operations.
- Flexibility: DB administrators can monitor and manage the refresh process independently.
- Permission Flexibility: Use a different user with appropriate permissions for refresh operations.

Permission Verification

Before refreshing the materialized view manually, verify that the Amazon RDS for PostgreSQL database user has the required permissions:

- To check schema usage

```
SELECT has_schema_privilege('refresh_user', 'schema_name', 'USAGE');
```

- Check table select permissions

```
SELECT has_table_privilege('refresh_user', 'table_name', 'SELECT');
```

- Check materialized view permissions

```
SELECT has_table_privilege('refresh_user', 'materialized_view_name',
'INSERT, UPDATE, DELETE');
```

How to Use

1. During recovery, leave the Refresh Materialized Views option disabled.
2. Complete the database recovery.
3. Download the refresh script from the Cohesity UI.
4. Run the script manually after recovery.

Download the refresh script

In the Cohesity UI:

1. Navigate to **Recoveries**.
2. Select the PostgreSQL recovery task.
3. Open the **Options** tab.
4. Locate the Materialized View Refresh section.
5. Download the refresh script.

Run the Refresh Script

1. Execute the downloaded script on the recovered database
2. The script automatically determines materialized view dependencies.
3. Materialized views are refreshed in dependency order.

Best Practices

- **Assess Your Data Size**: For databases with materialized views exceeding several GB, consider manual refresh.
- **Verify Permissions**: Before using automatic refresh, ensure DataProtect source credentials have required permissions.
- **Save the Script**: Download and save the refresh script from the recovery details page for future use.
- **Test in Non-Production**: Validate your chosen approach in test/dev environments first.
- **Schedule Appropriately**: For manual refresh, run during maintenance windows or low-traffic periods
- **Monitor Resources**: Materialized view refresh can be resource-intensive; ensure adequate system resources.
- **Regular Maintenance**: Consider implementing periodic refresh schedules (using tools like `pg_cron`) independent of recovery operations.
- **Document Your Views**: Maintain documentation of materialized views and their dependencies.

Amazon Relational Database Service for Oracle

Amazon Relational Database Service (RDS) for Oracle is a fully managed service offered by Amazon Web Services (AWS) that simplifies the deployment, operation, and scaling of Oracle Database in the cloud. Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Amazon RDS for Oracle.

Backup Methods

You can protect your Amazon RDS for Oracle using one of the following backup methods:

- **AWS Snapshot methods**: For protecting the Amazon RDS for Oracle, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon RDS for Oracle.
- **Ingest method**: Cohesity uses this option to protect Amazon RDS for Oracle by ingesting the backup data into the Cohesity cluster. During a restore operation,

Cohesity does not rely on connectivity to AWS RDS Service. Instead, restoration requires connectivity only to the Cohesity cluster where the backups reside and the EC2 target VM.

Important:

- Cohesity supports the protection of Amazon RDS for Oracle only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster) and Cohesity Cloud Edition (legacy) on AWS.
- Cohesity supports both AWS snapshot-based protection and ingest-based protection of Amazon RDS for Oracle on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster).
- On Cohesity Cloud Edition (legacy) on AWS (formerly known as AWS Cloud Edition cluster), Cohesity supports only ingest-based protection of Amazon RDS for Oracle.
- The ingest-based protection of Amazon RDS for Oracle on Cohesity Cloud Edition on AWS is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

This topic provides information on how you can protect Amazon RDS for Oracle on Cohesity Cloud Edition on AWS.

Plan and Prepare for Amazon RDS for Oracle Protection

Before protecting Amazon RDS for Oracle with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon RDS for Oracle with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Authenticate for Ingest-based Protection

To protect Amazon RDS for Oracle using the Ingest-based protection method, you must authenticate by providing the username and password of an Amazon RDS for Oracle instance user. These credentials are required for Cohesity to connect to the Oracle instance and initiate backup operations.

For CDB-enabled databases, Cohesity must connect to one of the PDBs. In this case, the credentials must be provided in the following format:

Username: `{PDB username}@{PDB name}`

Password: Corresponding PDBs password

Example:

- Consider a CDB named **CDB** that contains two PDBs: **PDB1** and **PDB2**. The credentials for these PDBs are as follows:
- **PDB1**: `pdb1admin/Mynameis1pwd`
- **PDB2**: `pdb2admin/Mynameis2pwd`

When adding credentials for the CDB database (**CDB**) during protection, you can choose either of the following:

- `Pdb1admin@PDB1`: In this case use the password `Mynameis1pwd`
- `Pdb2admin@PDB2`: In this case use the password `Mynameis2pwd`

For information on the steps to authenticate the Cohesity cluster to the Amazon RDS for Oracle in the Cohesity cluster UI, see Protect Amazon RDS for Oracle Using Ingest Method.

AWS S3 Staging Bucket Requirement

For ingest-based backup of Amazon RDS Oracle, Cohesity uses an Amazon S3 staging bucket as temporary storage.You can either allow Cohesity to automatically create the staging bucket when you initiate a backup or recovery task, or you can manually create the staging bucket in your AWS account before starting the process.

You must use the same bucket for both backup and recovery. If you choose to manually create the bucket, ensure to create the AWS S3 bucket of type, S3 Standard in the same AWS region as the database selected for backup and restore.

You can specify only one custom bucket name per Protection Group. Because an S3 bucket belongs to a single AWS region, a custom bucket is supported only when all databases in the Protection Group are in the same region. If your Protection Group spans multiple regions, use the **Auto Create** option instead.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Ports Used for Communication

For information on port numbers required, seeAWS Cloud Database Workloads.

(Optional) Encrypt Data to be Protected

Optionally, if you want to encrypt the data to be protected, you can enable the NNE (Native Network Encryption) option for RDS Oracle.

Considerations for AWS RDS for Oracle Data Protection

Review and understand the following limitations and known issues before you protect Amazon RDS for Oracle:

- Cohesity supports only full backups for Amazon RDS for Oracle instances. Even if you choose an incremental-based policy, Cohesity will always perform a full backup, regardless of the selected policy type.
- Incremental and log backups are not supported.
- Do not select a log-backup–based policy, as log backups are not supported. The backup will fail if you select a log-based policy.
- Tune the section size and the number of parallel streams based on the database size for optimal performance.
- Group databases of similar size into the same Protection Group to improve overall performance.
- Cross-account restore of Amazon RDS for Oracle protected using the AWS Snapshot protection type is not supported.
- When you use the **Specify Bucket to Use** option, you can specify only one bucket name per Protection Group. Because an S3 bucket resides in a single AWS region, this option is not supported in scenarios where databases from multiple regions or accounts might be protected together, such as:
  - Auto-protect at the AWS Account, OU, or Organization level.
  - Tag-based auto-protection.
  - A Protection Group that includes databases from multiple AWS regions.
  - AWS Organization registration spanning multiple accounts with databases.
  - For these scenarios, use the **Auto Create** option to allow Cohesity to automatically create the appropriate staging bucket per region.
- When recovering (snapshot-based) encrypted RDS instance to the same AWS account but a different AWS region, you must create a KMS encryption key in the target region with the same alias name as the KMS key used to encrypt the source RDS instance.
- When performing cross-region cross-account recovery (snapshot-based) of an Amazon RDS instance encrypted with a customer-managed KMS key, ensure the following:
  - The source account's KMS key alias is available in both the source and destination regions of the destination account.

  Example:

  If A1 is the source account with region R1, and A2 is the destination account with region R2, ensure the KMS key alias from A1 is available in both R1 and R2 of A2.

- The destination account has permission to access the source account's KMS key alias.
- The option group from the source account is available in both the source and destination regions of the destination account.

Backup Amazon RDS for Oracle

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads including Amazon RDS for Oracle associated with the specified AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **RDS Oracle** tab to view all discovered instances.

For information on registering AWS, see Register AWS Source.

Based on the registration level of your AWS source, the **RDS Oracle** tab displays the entity hierarchy for the Amazon RDS for Oracle, allowing navigation from the organization or account level down to individual instances.

RDS for Oracle Entity Hierarchy

The **RDS Oracle** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS Organization registered with Cohesity for RDS for Oracle protection. Applicable when the AWS source is registered at the organization level.
Organization Units (OUs)	The organizational units within the AWS Organization contain member accounts.
Member Account	The AWS member accounts within the organization that contain RDS for Oracle instances.
AWS Account	The AWS account registered with Cohesity for RDS protection. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the RDS resources reside.

Structure	Description
RDS for Oracle Instance	The specific RDS for Oracle instance within the selected region. You can mouse hover over each database to view the following details: • **Size**: The logical size of the database. • **AWS Region**: The region where the RDS for Oracle instance resides. • **Multi-tenant**: Whether multi-tenancy is enabled on the database. Multi-tenancy enabled implies that the database is Oracle container Database (CDB). • **Pluggable DBs**: The number of pluggable databases on the selected RDS for Oracle. • **Protection Details**: Whether the database is protected or unprotected.

Backup Scope by RDS for Oracle Hierarchy

The Cohesity cluster allows you to back up Amazon RDS for Oracle at different levels of the entity hierarchy. You can choose to back up:

- **AWS Organization** – Protect all RDS for Oracle instances across all member accounts in the registered AWS organization.
- **Organizational Unit (OU)** – Protect all RDS for Oracle instances within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all RDS for Oracle instances within a specific member account in the registered AWS organization.
- **AWS Account** – Backup all the Amazon RDS for Oracle instances within the registered AWS account.
- **AWS Region** – Backup all the Amazon RDS for Oracle instances within a specific region of the registered account.
- **RDS for Oracle instance** – Backup individual Amazon RDS for Oracle instance.

Backup Process

AWS Snapshot Method

When you initiate a protection run using the AWS Snapshot method, Cohesity triggers a native AWS RDS snapshot of the selected Amazon RDS for Oracle instance. The snapshot is created and stored directly within the same AWS account and region as the source instance, leveraging AWS's built-in snapshot infrastructure.

Ingest Method

When you initiate a protection run in the Cohesity cluster to back up Amazon RDS for Oracle, Cohesity first exports the selected instance data to an Amazon S3 staging bucket.

Once the data is successfully exported to the S3 bucket, Cohesity ingests the staged data from the S3 bucket into the Cohesity cluster. Once the ingestion is successful and the data

is fully backed up to the cluster, Cohesity automatically deletes the exported data from the staging bucket to free up storage and minimize S3 costs.

Supported Backup Types

Cohesity supports only full backups for Amazon RDS for Oracle instances. Incremental and log backups are not supported.

Protect Amazon RDS for Oracle Using Ingest Method

In Cohesity, Amazon RDS for Oracle instances are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

With the Ingest protection method, Cohesity protects Amazon RDS for Oracle by ingesting the backup data into the Cohesity cluster.

Prerequisites

Choose Staging Bucket

Before initiating a backup of Amazon RDS for Oracle instances, ensure that an Amazon S3 staging bucket is available in the AWS account where the Amazon RDS for Oracle instances reside. This bucket is used to temporarily store exported data during the backup process. You can manually specify a new or existing S3 bucket in your AWS account, or allow Cohesity to automatically create one when the backup process begins.

AWS Source Registered at the Account Level

When creating a Protection Group, you can manually specify an S3 bucket only if all selected RDS for Oracle instances belong to the same AWS account and region.

If all selected RDS for Oracle instances are in the same AWS account and region:

You can manually specify a newly created or existing S3 bucket, or allow Cohesity to automatically create one. If you are manually specifying an S3 bucket, ensure the S3 bucket is in the same AWS account and region as the selected RDS for Oracle instances.

If the selected RDS for Oracle instances span multiple regions:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

AWS Source Registered at the Organization Level

When protecting RDS for Oracle instances across member accounts, you can manually specify an S3 bucket only if all selected RDS for Oracle instances for protection belong to the same member account and region as the specified S3 bucket.

If all selected RDS for Oracle instances are in the same AWS member account and region:

You can manually specify a newly created or existing S3 bucket, or allow Cohesity to automatically create one. If you are manually specifying an S3 bucket, ensure the S3 bucket is in the same member account and region as the selected RDS for Oracle instances.

If the selected RDS for Oracle instances span multiple AWS member accounts or regions:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Authenticate for RDS for Oracle Protection

To protect the RDS for Oracle by ingesting the backup data to the Cohesity cluster, the Cohesity cluster must authenticate to the RDS for Oracle.

Perform the following steps for authentication:

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AWS source and click into it.
3. Click the **RDS Oracle** tab.
4. Select the checkbox for the hierarchy level (AWS Organization, OUs, member account, account, region, or instance level), and then click **instance Credentials** above the checkboxes.

   The **instance Credentials for <instance>** appears.
5. Provide the Username and Enter Password of the Amazon RDS for Oracle. For information on the authentication requirements, see Authenticate for Ingest-based Protection.
6. Click **Save**.

Protect Amazon RDS for Oracle Instance Using Ingest Method

Cohesity uses the Ingest Method option to protect the Amazon RDS for Oracle by ingesting the backup data into the Cohesity cluster.

To protect Amazon RDS for Oracle instances:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS Oracle**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for Oracle instance for protection using one of the following views:
   - **List View**: Use this option to select the RDS for Oracle instances for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select instances for protection at various hierarchy levels.

     Optionally:

     You can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new RDS for Oracle instances added to the AWS source.

     - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
       - **AWS Organization**
       - **Organizational Unit (OU)**
       - **Member Account**
       - **AWS Region**
     - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
       - **AWS Account**
       - **AWS Region**

     You can also exclude the Amazon RDS for Oracle instances from auto-protection within an account or region by clicking the auto-protect icon ( ) next to the RDS for Oracle instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.
   - **Tags**: Use this option to select the instances for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all instances with that tag for protection.

     Optionally:

- You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for Oracle instances with this tag.
- You can exclude the Amazon RDS for Oracle instances from auto-protection within a tag by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.
- You can exclude instances with a specific tag from auto-protection by clicking the exclude icon (⊘). Tags excluded from auto-protection are displayed with the exclude icon.

> **Note:** If you enable auto-protect, the Protection Group may include databases from multiple AWS regions. In this case, you must use the **Auto Create** option for the staging bucket. Using **Specify Bucket to Use** with a multi-region selection will cause the backup to fail.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary AWS S3 staging bucket in the same region as your RDS for Oracle instances selected for backup. The bucket follows the naming format: `cohe-<source-id>-<region-name>`.

     Where:
     - `source-id` is the AWS account ID.
     - `region-name` is the AWS region where the database reside.
   - **Specify Bucket to Use**: Enter the name of an existing AWS S3 bucket in the same region and account as the databases being protected. If the Protection Group includes databases from multiple AWS regions, do not use this option, the backup will fail. Use **Auto Create** instead.

> **Note:** The **Specify Bucket to Use** option is not supported when the Protection Group includes databases from multiple AWS regions. This includes Protection Groups using auto-protect at the AWS Account, OU, or Organization level, tag-based auto-protection, or manual selection of databases across regions. In these cases, select **Auto Create**. If you select **Specify Bucket to Use** in these scenarios, the backup will fail.

If the selected databases span multiple AWS accounts or regions, the option to manually specify an S3 bucket isn't available. In such cases, Cohesity automatically creates the staging bucket. For more information, see Choose Staging Bucket.

After each protection run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

8. In the Policy section, select **Ingest Method**, and then select a policy from the drop-down.

> **Note:**
> - You must not select a log backup-based policy as log backup is not supported. Backup will fail if you select a log backup-based policy.
> - If you select an incremental-based policy, Cohesity will still perform a full backup.

9. Select **More Options** to change or configure any of the Advanced Settings.
10. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Parallel Streams**	Specify the number of parallel streams to be used during backup. The default value is 4.
	Section Size	Define the section size for data transfer during the backup operation. We recommend using 40 GB as the section size, which is set as the default value.
	Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Protect Amazon RDS for Oracle Using AWS Snapshot Method

In Cohesity, Amazon RDS for Oracle instances are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

With the AWS Snapshot protection method, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon RDS for Oracle.

To protect your Amazon RDS for Oracle instance using AWS snapshot method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS Oracle**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for Oracle instances for protection using one of the following views:
    - **List View**: Use this option to select the RDS for Oracle instances for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select instances for protection at various hierarchy levels.

      Optionally,

      You can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new RDS for Oracle instances added to the AWS source.

        - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Organization**
            - **Organizational Unit (OU)**
            - **Member Account**
            - **AWS Region**
        - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Account**
            - **AWS Region**

You can also exclude the Amazon RDS for Oracle instances from auto-protection within an account or region by clicking the auto-protect icon ( ) next to the RDS for Oracle instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the instances for protection based on tags.

  Optionally,

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for Oracle instances with this tag.
  - You can exclude the Amazon RDS for Oracle instances from auto-protection within a tag by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude instances with a specific tag from auto-protection by clicking the exclude icon (⊘). Tags excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The AWS Snapshot-based policy is enabled by default. Create or select an existing policy from the **Policy (AWS Snapshot)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Parallel Streams**	This option is not applicable for snapshot-based protection.
	Section Size	This option is not applicable for snapshot-based protection.
	Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon RDS for Oracle Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon RDS for Oracle instances or databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon RDS for Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
    - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover RDS for Oracle

You can recover RDS for Oracle instance using one of the following methods:

- Recover Amazon RDS for Oracle Using Snapshot Method
- Recover Amazon RDS for Oracle Using Ingest Method

Recover Amazon RDS for Oracle Using Ingest Method

If you choose to recover the Amazon RDS for Oracle using the ingest method, then you can recover the Amazon RDS for Oracle instance as a standalone instance on an EC2 instance of the same AWS account or a different AWS account.

Prerequisites

Before recovering an Amazon RDS for Oracle instance as a standalone database on an EC2 instance, ensure the following:

- **Oracle Software Installation**: The target EC2 instance must have Oracle software installed. The Oracle software version must match the version of the source RDS for Oracle instance.
    - If the target Oracle version matches the source version, no additional configuration is required.
    - If the target Oracle version is newer than the source, you may need to manually set the COMPATIBLE parameter in the pfile during recovery. The value should match the source database's compatibility version.
- **Database Instance**: The database instance you are recovering must not already exist on the target EC2 instance.

- **Free Space**: The EC2 instance must have sufficient free space to restore the database.
- **Cohesity Agent**: The Cohesity Agent must be running on the EC2 instance (this applies only to Linux-based Oracle databases).
- **OS Authentication**: If the Cohesity Agent is running as a non-root, non-oracle user, OS authentication must be enabled for that user.

Recover Amazon RDS for Oracle

To recover Amazon RDS for Oracle:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS Oracle**.
3. Select **Ingest** as the **Recovery Task**.
4. Search for the instance by entering the instance name or a partial name, or use the wildcard character (*) to view all available instances.
5. From the list of search results, select the instance to recover. By default, the latest snapshot of the instance is selected for recovery. To recover from a different snapshot, perform the following steps depending on the instance selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for *<instance_name>*** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the **Edit recovery point for *<instance_name>*** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
     b. Select the snapshots to be recovered.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

6. Click **Next: Recover Options**.
7. In the **Recovery To** section, provide the following information:

   Based on the registration level of the AWS source, select the recovery destination hierarchy:

   - If the AWS source is registered at the organization level:
     - **Source**: Select the destination AWS Organization. You can recover data only within the same AWS Organization as the source.
     - **Account**: Select the destination account.
     - **EC2 Instance:** Type the EC2 instance name to which you want to recover the RDS for Oracle as a standalone database.
   - If the AWS source is registered at the account level:
     - **Account**: Select the destination account.
     - **EC2 Instance**: Type the EC2 instance name to which you want to recover the RDS for Oracle as a standalone database.

   a. **Restore Database Files to**: Specify the path to an existing empty directory. The newly created database files will reside in this path.

   > **Note:** You can restore to an ASM path using this option. Specify just the ASM volume name instead of the entire path. Example: The restore path for ASM volume data1 is +data1. If you specify the entire path, +data1/dbname the restore job will fail.

   b. **Oracle Home**: The ORACLE_HOME value for the host where the database is restored.

   c. **Base Directory**: The directory for the restored database.

   d. **Target Database Name**: The name for the newly restored database.

8. In the **Recovery Options** section, configure the following options:

   a. **Number of Channels**: Specify the number of channels to perform the restore job. The system will run these channels in parallel to improve the restore performance.

b. **Use NOFILENAMECHECK**: Set this option to **YES** to force RMAN to skip checking whether target datafile names match the source datafile names.

> **Important:** Use NOFILENAMECHECK with caution, as this restore can inadvertently clobber the datafiles if restored back to the source host which already resides on the host.

> **Note:** Some restore use case require restoring the datafiles and/or online logfiles using the same names and file paths as the source database, but to a different host (i.e. not restoring back to the source database host).

For example, using this restore to create a standby database of the source database with exactly the same directory structure on a different host. In this case, it is required to set the 'Use NOFILENAMECHECK' option for this alternate restore.

c. **Enable Archive Log mode for the Database**: Set this option ON if you want to enable the redo log archiving on the recovered database.

d. **BCT File Path**: Specify the Block Change Tracking (BCT) file path to determine where the block change tracking file will be created. If not provided, BCT will not be enabled for the restored database.

e. **Number of Temporary Files to Use**: Enter the number of temp files to rename from the source database.

f. **Redo Log Group Settings**: Provide the number, size, and prefix of redo log groups, along with the redo log group member path for the recovery database. Redo log groups store all changes made to the database and protect against instance failure.

g. **Task Name**: Update the default name of the recovery task.

9. Click **Recover**.

Recover Amazon RDS for Oracle Using Snapshot Method

If you choose to recover the Amazon RDS for Oracle using the snapshot method, then you can recover RDS for Oracle instance to the original location or to a different location.

To recover Amazon RDS for Oracle:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS Oracle**.
3. Select **Snapshot** as the **Recovery Task**.

4. Search for the instance by entering characters of the instance name, or by specifying the wildcard character *.
5. From the list of search results, select the instance to recover. By default, the latest snapshot of the instance is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for *<instance_name>*** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the **Edit recovery point for *<instance_name>*** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. You can also select any point on the green solid line on the slider to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.
     b. Select the snapshots to be recovered.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
     b. Select the snapshot to be recovered by clicking the corresponding blue dot.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the instance and then provide the required information:

- **Original Location**: Restore the instance to the same region.

  > **Note:** If an RDS instance with the same name already exists, the restore will fail. Ensure the original instance is not present or restore with a different name.

  Type **Yes** to confirm.
- **New Location**: Restore the instance to a different region. You can restore the instance to the same or a different region within the same account as the original RDS instance.

  Based on the registration level of the AWS source, select the recovery destination hierarchy:

  - If the AWS source is registered at the organization level:
    - **Source**: Select the AWS organization that contains the destination instance. You can recover the instance within the same organization as the source instance.
    - **AWS Account**: Select the member account to which you want to restore the instance.
    - **Region**: Select a region to restore the instance.
    - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
    - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
    - **Availability Zone**: Select an availability zone in AWS to recover the RDS.
  - If the AWS source is registered at the account level:
    - **Source**: Select the registered AWS account to which you want to restore the instance.
    - **Region**: Select a region to restore the instance.
    - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
    - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
    - **Availability Zone**: Select an availability zone in AWS to recover the RDS.

To overwrite any existing instance with the same name, enable **Overwrite existing instance(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of RDS instance protected using the AWS Snapshot protection type is not supported.
> - If an RDS instance with the same name already exists, the restore will fail.

8. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
9. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
10. **DB Parameter Group**: Select the database (DB) parameter group to associate with the DB instance.
11. **DB Port:** Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port. The default port is 3306.
12. Configure the following **Recovery Options**:
    a. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.
13. **Task name**: Optionally, change the default name of the recovery task.
14. Click **Recover**.

Amazon Relational Database Service for Microsoft SQL Server

Amazon Relational Database Service (RDS) for SQL Server is a fully managed relational database service offered by Amazon Web Services (AWS) that simplifies the deployment, operation, and scaling of Microsoft SQL Server databases in the cloud. Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for the backup, recovery, and data management of Amazon Relational Database Service for Microsoft SQL Server (MS SQL) databases, including databases enabled with Transparent Data Encryption (TDE).

Backup Method

Cohesity uses the Ingest method to protect Amazon RDS for MS SQL Server databases. With this approach, Cohesity exports database-level data from Amazon RDS for MS SQL to an Amazon S3 staging bucket, then ingests the data into the Cohesity cluster for backup. In addition to full backups, Cohesity supports cumulative incremental backups, which capture all changes since the most recent full backup. Incremental backups always depend on the last successful full backup.

> **Important:**
>
> - Cohesity supports the protection of Amazon RDS for MS SQL only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster) and Cohesity Cloud Edition (legacy) on AWS .
> - On Cohesity Cloud Edition on AWS, Cohesity supports the protection of Amazon RDS for MS SQL at both the instance level (via AWS snapshot-based protection) and the database level (via ingest-based protection).
> - On Cohesity Cloud Edition (legacy) on AWS, Cohesity supports the backup and recovery of Amazon RDS for MS SQL only at the instance level.

Plan and Prepare for Amazon RDS for MS SQL Server Protection

Before protecting Amazon RDS for MS SQL Server with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon RDS for MS SQL Server with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Ports Used for Communication

For information on port numbers required, see AWS Cloud Database Workloads.

Authentication Requirement

When creating a Protection Group to protect the Amazon RDS for MS SQL Server at database level, you must authenticate the Cohesity cluster to discover the databases of the Amazon RDS for MS SQL Server instance using the credentials of the RDS for MS SQL Server instance. You can provide either the credentials of:

- Admin user, or
- Custom database user

If you are providing the Custom user credentials of RDS for MS SQL Server instance, then ensure the followings:

- The SQL user has the following permissions on the database:

```
-- Create a login on the SQL Server

CREATE LOGIN dbuser WITH PASSWORD = '';

-- Create a database user mapped to the login

CREATE USER [dbuser] FOR LOGIN [dbuser];

-- Add the user to the db_owner role

ALTER ROLE [db_owner] ADD MEMBER [dbuser];
```

- The custom user is not a part of any database deny role such as `db_denydatareader`, `db_denydatawriter`. These roles prevent discovery and backup operations.

For information on the steps to authenticate the Cohesity cluster to the Amazon RDS for MS SQL in the Cohesity cluster UI, see Protect Amazon RDS for MS SQL Server Using Ingest Method.

AWS S3 Staging Bucket Requirement

For ingest-based backup and recovery of Amazon RDS for Microsoft SQL Server databases, Cohesity uses an Amazon S3 staging bucket as temporary storage. You can either allow Cohesity to automatically create the staging bucket when you initiate a backup or recovery task, or you can manually create the staging bucket in your AWS account before starting the process.

You must use the same bucket for both backup and recovery. If you choose to manually create the bucket, ensure to create the AWS S3 bucket of type, S3 Standard in the same AWS region as the database selected for backup and restore.

You can specify only one custom bucket name per Protection Group. Because an S3 bucket belongs to a single AWS region, a custom bucket is supported only when all databases in the Protection Group are in the same region. If your Protection Group spans multiple regions, use the **Auto Create** option instead.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Configure Option Group

An Option group is required as the RDS for MS SQL Server backup and restore operations use the SQLSERVER_BACKUP_RESTORE option, which enables native backup and restore integration with Amazon S3.

Perform the following steps to configure Option Group:

1. Create an Option Group.
2. Add SQLSERVER_BACKUP_RESTORE Option to the Option group.
3. Associate the Option group to RDS Instances.

Create an Option Group

To create Option group:

1. Log in to the Amazon Management Console and navigate to RDS.
2. From the left menu, select **Option groups**.
3. Click **Create Option Group**.
4. Provide the following information:
    a. Name
    b. Description
    c. Engine
    d. Major Engine Version
    e. Click **Create**.

Add SQLSERVER_BACKUP_RESTORE Option to Option Group

The SQLSERVER_BACKUP_RESTORE option needs to be associated with the Option group to perform backup and restore of RDS for MS SQL Server.

To add SQLSERVER_BACKUP_RESTORE option to the Option Group:

1. From the AWS Management console, open the Option Group created.
2. Click **Add Option**.
3. Select **SQLSERVER_BACKUP_RESTORE**.

4. Select **CohesityBackupS3StagingRole a**s IAM role to grant write access to staging S3 bucket.
5. Save the changes.

Associate the Option Group to RDS Instances

To perform the backup and recovery of an RDS for MS SQL Server instance, it must be associated with the Option group with SQLSERVER_BACKUP_RESTORE option attached.

To associate the Option group to RDS Instances you want to protect:

1. From the AWS Management console.
2. Navigate to **Databases** in RDS.
3. Select your SQL Server instance.
4. Modify the instance.
5. Under **Option group**, select the Option Group you created.
6. Apply the changes.

Considerations for Amazon RDS for MS SQL Data Protection

Review and understand the following limitations and known issues before you protect Amazon RDS for MS SQL Server databases:

Backup Capabilities and Behavior

- Cohesity supports full and cumulative incremental backups for Amazon RDS for Microsoft SQL Server databases.
- Incremental backups are cumulative and depend on the most recent full backup.
- If the last full backup is unavailable, incremental restore is not possible.
- For Multi-AZ RDS instances, automated backups cannot be disabled and may cause more frequent fallback to full backups.

Backup Prerequisites and Requirements

- The RDS for MS SQL Server instance must be in an available state for Cohesity to initiate a backup.
- Backups cannot be initiated if the server is offline or unavailable.
- The staging buckets used for temporarily storing data during backup and recovery must be located in the same AWS region as:
    - The database being protected, and
    - The selected recovery target.

Protection Policy Constraints

- You cannot perform ingest-based and snapshot-based protections within the same protection group.
- If you select databases for protection, you cannot choose a snapshot-based policy.
- If you select instances with no databases discovered, you cannot choose an ingest-based policy.

Restore Limitations and Scope

- Restoring a backup taken from a higher RDS for MS SQL Server database version to a lower version is not supported.
- Restoring to another AWS region within the same AWS account is supported.
- Restoring to a different AWS account is not supported.
- Incremental restore is supported only to single-AZ RDS instances.
- Point-in-time restore (PITR) is supported when recovering Amazon RDS for MS SQL at instance-level.
- Cross-account restore of Amazon RDS for MS SQL instance protected using the AWS Snapshot protection type is not supported.

Transparent Data Encryption (TDE) Considerations

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- Cohesity backs up only databases, not TDE certificates.
- You are responsible for managing TDE certificates in Amazon RDS.
- To restore a TDE-enabled database to a different MS SQL Server instance:
  - First back up the TDE certificate from the source instance.
  - Then restore the certificate to the target instance before performing the database restore from the Cohesity cluster UI.

For information on the backup and restore of TDE certificate, see the AWS documentation.

KMS Encryption and Cross-Region/Cross-Account Recovery

- When recovering encrypted RDS instances to the same AWS account but a different AWS region, you must create a KMS encryption key in the target region with the same alias name as the KMS key used to encrypt the source RDS instance.
- When performing cross-region cross-account recovery of an Amazon RDS instance encrypted with a customer-managed KMS key, ensure the following:

- The source account's KMS key alias is available in both the source and destination regions of the destination account.

Example:

If A1 is the source account with region R1, and A2 is the destination account with region R2, ensure the KMS key alias from A1 is available in both R1 and R2 of A2.

- The destination account has permission to access the source account's KMS key alias.
- The option group from the source account is available in both the source and destination regions of the destination account.

Custom Staging Bucket and Multi-Region Considerations

When you use the **Specify Bucket to Use** option, you can specify only one bucket name per Protection Group. Because an S3 bucket resides in a single AWS region, this option is not supported in scenarios where databases from multiple regions or accounts might be protected together, such as:

- Auto-protect at the AWS Account, OU, or Organization level.
- Tag-based auto-protection.
- A Protection Group that includes databases from multiple AWS regions.
- AWS Organization registration spanning multiple accounts with databases.
- For these scenarios, use the **Auto Create** option to allow Cohesity to automatically create the appropriate staging bucket per region.

Attach Permissions to IAM Role for RDS for MS SQL Protection

After completing source registration by running the CloudFormation Template (CFT), the following IAM roles and profiles are created:

- *CohesityInstanceRole*
- *CohesityInstanceProfile*
- *CohesitySourceRegistrationRole*

The *CohesitySourceRegistrationRole* contains the permissions defined in the CFT that allow Cohesity to discover AWS resources in the entity hierarchy within the product UI.

However, to perform backup and recovery operations for RDS for MS SQL, you must attach the permissions listed in the Plan and Prepare for Amazon RDS for MS SQL Server Protection section to the *CohesityInstanceRole*.

Attach Permissions

Perform the following steps to attach the permissions to the *CohesityInstanceRole*.

1. In the AWS Management Console, navigate to **IAM** > **Roles**.
2. Search for and open the role *CohesityInstanceRole*.

3. Go to the **Permissions** tab and select **Add permissions** > **Create inline policy**.
4. Switch to the **JSON** tab.
5. Paste the required RDS for MS SQL permissions JSON policy and save it. For the list of permissions, see Plan and Prepare for Amazon RDS for MS SQL Server Protection.

Backup Amazon RDS for MS SQL Server

Depending on the services selected while registering the source, the Cohesity cluster automatically discovers all supported workload with the registered AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **RDS MSSQL** tab to view all discovered instances. For information on registering AWS, see Register AWS Source.

Based on the registration level of your AWS source, the **RDS MSSQL** tab displays the entity hierarchy for the Amazon RDS for MS SQL Server environment, allowing navigation from the organization or account level down to individual databases.

Amazon RDS for Microsoft SQL Server Hierarchy

The Amazon RDS for Microsoft SQL Server entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with Cohesity for RDS for MS SQL protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The organizational units within the AWS organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain RDS for MS SQL Server instances.
AWS Account	The AWS account registered with Cohesity for protecting Amazon RDS for Microsoft SQL Server databases. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the RDS resources reside.
Availability Zone	The availability zone within the region where the RDS instance is deployed.
RDS MS SQL Server Instance	The specific RDS for MS SQL Server instance within the selected availability zone.
RDS MS SQL Server Database	The individual databases hosted within the RDS for MS SQL Server instance. The database will be listed in the entity hierarchy only if you have provided credentials to discover the databases of the Amazon RDS for MS SQL Server instance.

Backup Scope by RDS for Microsoft SQL Server Hierarchy

The Cohesity cluster allows you to back up Amazon RDS for MS SQL Server at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level.

- **AWS Organization** – Protect all RDS for MS SQL Server instances and databases across all member accounts in the registered AWS organization.
- **Organizational Unit (OU)** – Protect all RDS for MS SQL Server instances and databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all RDS for MS SQL Server instances and databases within a specific member account in the registered AWS organization.
- **AWS Account** – Backup all the Amazon RDS for MS SQL Server databases within the registered AWS account.
- **AWS Region** – Backup all the Amazon RDS for MS SQL Server databases within a specific region of the registered account.
- **Availability Zone** – Protect Amazon RDS for MS SQL Server databases within a specific availability zone.
- **RDS MS SQL Server Instance** – Backup all Amazon RDS for MS SQL Server databases within a specific database instance.
- **RDS MS SQL Server Databases** – Backup individual Amazon RDS for MS SQL Server database.

Backup Process

AWS Snapshot Method

When you initiate a protection run using the AWS Snapshot method, Cohesity triggers a native AWS RDS snapshot of the selected Amazon RDS for MS SQL Server instance. The snapshot is created and stored directly within the same AWS account and region as the source instance, leveraging AWS's built-in snapshot infrastructure.

Ingest Method

When you initiate a protection run in the Cohesity cluster to back up Amazon RDS for MS SQL Server databases, Cohesity first exports the selected database data to an Amazon S3 staging bucket.

Once the data is successfully exported to the S3 bucket, Cohesity ingests the staged data from the S3 bucket into the Cohesity cluster. Once the ingestion is successful and the data is fully backed up to the Cohesity cluster, Cohesity automatically deletes the exported data from the staging bucket to free up storage and minimize S3 costs.

Supported Backup Types

Cohesity supports full and cumulative incremental backups. The first protection run after the initial backup configuration or upgrade is always a full backup. The subsequent

protection runs perform cumulative incremental (differential) backups. Each incremental backup contains all changes since the most recent full backup.

Incremental backups fall back to full backup if:

- An automated full backup or a snapshot (initiated from the MS SQL Server databases) was taken after the last successful full backup initiated by Cohesity.

  > **Note:** It is recommended to disable automated backups to prevent frequent fallbacks to full backups.

- No prior full backup exists.

> **Note:** Cumulative incremental backup support is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Protect Amazon RDS for MS SQL Server Using Ingest Method

With the Ingest protection method, Cohesity cluster protects the Amazon RDS for MS SQL Server of your AWS source at the database level.

In Cohesity, Amazon RDS for MS SQL Server databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

> **Note:** The backup and recovery of Transparent Data Encryption (TDE)–enabled Amazon RDS for Microsoft SQL Server databases is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Prerequisites

Choose Staging Bucket

Before initiating a backup of Amazon Web Services Amazon RDS for MS SQL Server at database, an Amazon S3 staging bucket must be available in the AWS account where the RDS for MS SQL Server resources reside. This bucket is used to temporarily store exported data during the backup process. You can manually create or select an existing S3 bucket in your AWS account, or allow Cohesity to automatically create one when the backup process begins.

AWS Source Registered at the Account Level

You can manually specify a new or existing S3 bucket, or allow Cohesity to automatically create one.

If all selected databases are in the same AWS account and region:

You can manually specify a newly created or existing S3 bucket, or allow Cohesity to automatically create one.

If the selected databases span multiple regions:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

AWS Source Registered at the Organization Level

You can manually specify an S3 staging bucket only if all selected databases for backup or recovery belong to the same AWS member account and region as the specified S3 bucket.

If all selected databases are in the same AWS member account and region:

You can manually specify a new or existing S3 bucket, or allow Cohesity to automatically create one.

If the selected databases span multiple AWS member accounts or regions:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Authenticate RDS for MS SQL Server Protection

To protect the RDS for MS SQL Server at database level, the Cohesity cluster must authenticate to the MS SQL Server instance. Once authentication is successful, Cohesity discovers the databases of the MS SQL Server instance for protection.

Perform the following steps for authentication:

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AWS account and click into it.
3. Click the **RDS MSSQL** tab.
4. Select the checkbox for the hierarchy level (AWS Organization, OU, member account, account, region, Availability zone or cluster level), and then click **Database**

Credentials above the checkboxes.

The **Database Credentials for <instance>** appears.

5. Provide the **Username** and **Enter Password** of the Amazon RDS for MS SQL Server instances. For more information, see Plan and Prepare for Amazon RDS for MS SQL Server Protection.
6. Click **Save**.
7. Click the **Refresh** icon for the discovered databases to appear on the UI.

The databases corresponding to the selected instance or hierarchy level are discovered.

Protect Amazon RDS for MS SQL Server Database

To protect Amazon RDS for MS SQL Server databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **AWS** > **RDS MSSQL**.
2. In the **New Protection** page, click **Add Objects** and select your registered source, or click **Register Source** to add a new one.
3. Select the RDS for MS SQL Server databases using one of the following views or option:
    - **List View**: Use this option to select the RDS for MS SQL Server databases for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select RDS for MS SQL Server databases for protection at various hierarchy levels.

      Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon RDS for MS SQL Server instances.

        - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Organization**
            - **Organizational Unit (OU)**
            - **Member Account level**
            - **AWS Region**
            - **Availability Zone**
            - **RDS Instance**

- If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
  - **AWS Account**
  - **AWS Region**
  - **Availability Zone**
  - **RDS Instance**

  You can exclude the RDS for MS SQL Server databases from auto-protection within an account, region, availability zone, or instance by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the databases for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all databases with that tag for protection.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added RDS for MS SQL Server databases with this tag.
  - You can exclude the RDS for MS SQL Server databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude databases with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

> **Note:** If you enable auto-protect, the Protection Group may include databases from multiple AWS regions. In this case, you must use the **Auto Create** option for the staging bucket. Using **Specify Bucket to Use** with a multi-region selection will cause the backup to fail.

4. Click **Continue.**
5. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
6. In the **Staging Bucket** section, choose one of the following options:

- **Auto Create**: Cohesity automatically creates a temporary AWS S3 staging bucket in the same region as your RDS for MS SQL Server databases selected for backup. The bucket follows the naming format: `cohe-<source-id>-<region-name>`.

  Where,

  - `source-id` is the AWS account ID.
  - `region-name` is the AWS region where the database reside.
- **Specify Bucket to Use**: Enter the name of an existing AWS S3 bucket in the same region and account as the databases being protected. If the Protection Group includes databases from multiple AWS regions, do not use this option, the backup will fail. Use **Auto Create** instead.

  > **Note:** The **Specify Bucket to Use** option is not supported when the Protection Group includes databases from multiple AWS regions. This includes Protection Groups using auto-protect at the AWS Account, OU, or Organization level, tag-based auto-protection, or manual selection of databases across regions. In these cases, select **Auto Create**. If you select **Specify Bucket to Use** in these scenarios, the backup will fail.

  If the selected databases span multiple AWS accounts or Regions, the option to manually specify an S3 bucket isn't available. In such cases, Cohesity automatically creates the staging bucket.

  After each protection run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

7. The Ingest-based policy is enabled by default. Create or select an existing policy from the **Policy (Cohesity Ingest)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Protect Amazon RDS for MS SQL Server Using AWS Snapshot Method

With the AWS Snapshot protection method, Cohesity cluster protects the Amazon RDS for MS SQL Server of your AWS source at the instance level.

In Cohesity, Amazon RDS for MS SQL Server instances are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

To protect your Amazon RDS for MS SQL Server instance using AWS snapshot method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS MS SQL**.

3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for MS SQL Server instances for protection using one of the following views:
    - **List View**: Use this option to select the RDS for MS SQL Server instances for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select instances for protection at various hierarchy levels.

      Optionally,

      You can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new RDS for MS SQL Server instances added to the AWS source.

        - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Organization**
            - **Organizational Unit (OU)**
            - **Member Account**
            - **AWS Region**
        - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Account**
            - **AWS Region**

      You can also exclude the Amazon RDS for MS SQL Server instances from auto-protection within an account or region by clicking the auto-protect icon ( ) next to the RDS for MS SQL Server instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.
    - **Tags**: Use this option to select the instances for protection based on tags.

      Optionally,

        - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for MS SQL Server instances with this tag.
        - You can exclude the Amazon RDS for MS SQL Server instances from auto-protection within a tag by clicking the auto-protect icon ( ) next to the

instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.

- You can exclude instances with a specific tag from auto-protection by clicking the exclude icon (⊘). Tags excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. The AWS Snapshot-based policy is enabled by default. Create or select an existing policy from the **Policy (AWS Snapshot)** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Parallel Streams**	This option is not applicable for snapshot-based protection.
	Section Size	This option is not applicable for snapshot-based protection.
	Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon RDS for MS SQL Server Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon RDS for MS SQL Server database to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon RDS for MS SQL Server database.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MS SQL Server database.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MS SQL Server database.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MS SQL Server database.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon RDS for MS SQL Server

Depending on the backup method you used for protecting Amazon RDS for MS SQL, you can:

- Recover Amazon RDS for MS SQL Server at Instance Level
- Recover Amazon RDS for MS SQL Server at Database Level

Recover Amazon RDS for MS SQL Server at Instance Level

You can recover the Amazon RDS for MS SQL Server at instance level if you have protected the Amazon RDS for MS SQL using the AWS Snapshot Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon RDS for MS SQL Server Protection section.

To recover Amazon RDS for MS SQL Server at instance level :

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS MS SQL**.
1. Select **Instance** as the recovery task.
2. Search for the instance by entering characters of the instance name, or by specifying the wildcard character *.
3. From the list of search results, select the instance to recover. By default, the latest snapshot of the instance is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for *<instance_name>*** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the **Edit recovery point for *<instance_name>*** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. You can also select any point on the green solid line on the slider to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.
  b. Select the snapshots to be recovered.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.
- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.

4. Click **Next: Recover Options**.
5. In the **Recover To** section, select any one of the following options based on where you want to restore the instance and then provide the required information:
   - **Original Location**: Restore the instance to the same region.

     > **Note:** If an RDS instance with the same name already exists, the restore will fail. Ensure the original instance is not present or restore with a different name.

     Type **Yes** to confirm.
   - **New Location**: Restore the instance to a different region. You can restore the instance to the same or a different region within the same account as the

original RDS instance.

Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source**: Select the AWS organization that contains the destination instance. You can recover the instance within the same organization as the source instance.
  - **AWS Account**: Select the member account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.
- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.

To overwrite any existing instance with the same name, enable **Overwrite existing instance(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of RDS instance protected using the AWS Snapshot protection type is not supported.
> - If an RDS instance with the same name already exists, the restore will fail.

6. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
7. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
8. **DB Parameter Group**: Select the database (DB) parameter group to associate with the DB instance.
9. **DB Port:** Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port. The default port is 3306.
10. Configure the following **Recovery Options**:
    a. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.
11. **Task name**: Optionally, change the default name of the recovery task.
12. Click **Recover**.

Recover Amazon RDS for MS SQL Server at Database Level

You can recover the Amazon RDS for MS SQL Server at database level if you have protected the Amazon RDS for MS SQL using the RDS Ingest Protection Type.

> **Note:** We recommend that you review the Plan and Prepare for Amazon RDS for MS SQL Server Protection section.

Recover TDE-Enabled Databases

Before you recover a TDE-enabled database to another Amazon RDS for SQL Server instance using Cohesity, you must:

1. Back up the TDE certificate from the source instance to Amazon S3.
2. Restore the TDE certificate to the target instance. For more information, see the Amazon Documentation.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Recovery Process

When you initiate a recovery, Cohesity exports the selected database data from the Cohesity cluster to an Amazon S3 staging bucket. You can manually create or select an existing S3 bucket in your AWS account, or allow Cohesity to automatically create one when the recovery process begins.

After the data is successfully exported to the staging bucket, Amazon RDS reads the exported files from the bucket and imports them into the specified Amazon RDS for MS SQL database instance. Once the import completes, Cohesity automatically deletes the exported files from the staging bucket to free up storage and minimize S3 costs.

If you choose Auto Create:

Cohesity automatically creates a staging bucket in the same AWS account and region as the target RDS for MS SQL Server database instance.

If you choose to specify a bucket:

You must specify an S3 bucket that you created or that already exists. The S3 bucket must be in the same AWS region as the target RDS for MS SQL database instance.

Recover Amazon RDS for MS SQL Server Database

To recover Amazon RDS for MS SQL Server database:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS MSSQL**.
3. Select **Database** as the recovery task.
4. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
5. From the list of search results, select the database to recover.

- By default, the latest snapshot of the database is selected for recovery.
- To recover from a different snapshot, hover over the snapshot and click the **Edit** icon. The **Edit Recovery Point for <*database_name*>** page appears.
- By default, snapshots are displayed in the List mode. You can switch to the Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

To select a snapshot from the **List** mode:

a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*database_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

b. Select the snapshots to be recovered.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:
   - **Original Location**: Restore the database to the original Amazon RDS for MS SQL database instance.

     > **Note:** Restoring to the original location will overwrite any existing database with the same name.

     Type **Yes** to confirm.
   - **New Location**: Restore the database to an Amazon RDS for MS SQL database instance within the same AWS account. Restoring to another region within the same AWS account is supported. However, restoring to a different AWS account is not supported.

     Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source:** Select the AWS organization that contains the destination Amazon RDS for MS SQL Server instance. You can recover data only within the same AWS Organization as the source.
  - **Account**: Select the AWS member account to which you want to restore the database.
  - **Instance**: Select the Amazon RDS for MS SQL database instance to which you want to restore the database.
- If the AWS source is registered at the account level:
  - **Source**: Select the AWS account to which you want to restore the database.
  - **Instance**: Select the Amazon RDS for MS SQL database instance to which you want to restore the database.

To overwrite any existing database with the same name, enable **Overwrite existing database(s)** and type **Yes** to confirm.

> **Note:** Restoring to a new location also replaces any existing database with the same name in the selected instance.

> **Warning:** Restoring to the original location by overwriting the existing database may fail with the following error: `Operation PRE_RESTORE failed, but no errors reported by nbdbpaasagent`.
> **Resolution**: Rerun the restore taks or restore it with a different name to the original location.

8. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary AWS S3 staging bucket in the same region as your RDS for MS SQL Server databases selected for recovery. The bucket follows the naming format: `cohe-<source-id>-<region-name>`.

     Where:
     - `source-id` is the AWS account ID.
     - `region-name` is the target database AWS region.
   - **Specify Bucket to Use**: Enter the name of the AWS S3 staging bucket you

created for temporarily storing the data. Ensure the bucket is in the same region as the target RDS for MS SQL database instance.

After the recovery task, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.

> **Note:** Using a bucket in a different region than the database may incur additional storage costs.

9. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

10. Click **Recover**.

Amazon Relational Database Service Custom for Oracle

Cohesity supports protecting Relational Database Service (RDS) Custom for Oracle databases using two methods:

- **Ingest method (via Cohesity Oracle Adapter)**: Performs database-level protection by installing the Cohesity Oracle Adapter directly on the RDS Custom instance. This method enables RMAN-based, application-consistent backups and restores with external storage like EC2 VMs as target storage.
- **Snapshot method (via native AWS snapshots)**: Uses AWS's native snapshot capability to capture and restore the RDS Custom instance at the storage layer.

The setup and workflows differ for each method. Choose the method that best suits your protection strategy and operational requirements.

Cohesity enables protection of Amazon RDS Custom for Oracle databases through two supported approaches — ingest-based protection using the Cohesity Oracle Adapter and snapshot-based protection using AWS native snapshots.

With the ingest method, you can use a Protection Group to perform database-level protection for Amazon RDS Custom for Oracle instances. The Cohesity Oracle Adapter, installed directly on the RDS Custom instance, integrates natively with Oracle Recovery Manager (RMAN) to deliver application-consistent backup and recovery for single-instance RDS custom oracle databases.

The snapshot method, on the other hand, leverages AWS native snapshot capabilities to back up and restore the RDS Custom instance at the infrastructure level, offering a fast and storage-efficient alternative for customers preferring snapshot-based workflows.

> **Note:** Cohesity supports the backup and recovery of Amazon RDS Custom for Oracle on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster).

To learn more, see the following topics:

- Plan and Prepare for Amazon RDS Custom for Oracle
- Manage Protection Policy
- Backup Amazon RDS Custom for Oracle

Plan and Prepare for Amazon RDS Custom for Oracle

This topic provides information on minimum permissions, prerequisites, and considerations to protect Amazon RDS Custom for Oracle using both **Ingest** and **Snapshot** methods.

- Ingest method (via Cohesity Oracle Adapter)
- Snapshot method (via native AWS snapshots)

Ingest method (via Cohesity Oracle Adapter)

To protect Amazon Custom RDS for Oracle with Cohesity using the Ingest method, you must register your custom RDS EC2 instance as a source in the Cohesity cluster and perform backup and recovery operations. Ensure you meet the prerequisites and understand the considerations before registering for the AWS account.

Prerequisite

- The Cohesity Oracle adapter has to be installed on the Custom RDS EC2 instance.
- **Non-overlapping backup windows**: Schedule Amazon RDS Custom backup and maintenance windows so they do not overlap with Cohesity backup jobs.
- **Archive log retention**: Configure RDS Custom archive log retention to ensure that archived logs are retained long enough for at least one Cohesity backup cycle to complete successfully.

  For example, create the following configuration file on your RDS Custom instance:

```
sudo mkdir -p /opt/aws/rdscustomagent/config
sudo vi /opt/aws/rdscustomagent/config/redo_logs_custom_
configuration.json
Add the following content:
{
"archivedLogRetentionHours": "60"
}
```

This ensures archived logs are retained for 60 hours, preventing them from being purged before Cohesity backup jobs complete.

For more details, see Configuring archived log retention for Amazon RDS Custom.

- **Required service endpoints**: The following table lists the service endpoints required for the backup and recovery of AWS Custom RDS databases:

AWS Resource Type	Service Endpoints
IAM	iam.amazonaws.com
KMS	kms.<region>.amazonaws.com
RDS	rds.<region>.amazonaws.com
EC2 (for SBT traffic & agent operations)	ec2.<region>.amazonaws.com
CloudWatch Logs (for monitoring agent activity)	logs.<region>.amazonaws.com

Considerations

- Ingest Backup and recovery of Amazon RDS Custom databases are not supported on Physical clusters or Virtual Edition clusters. To back up and recover Amazon RDS Custom databases, the cluster must be a Next-Gen Cloud Edition (NGCE) cluster.
- Cohesity uses the Oracle Adapter installed on the RDS Custom instance to perform database-level ingest backups and restores. Backup and recovery operations are managed through Oracle RMAN, with Cohesity acting as the backup target.
- Cohesity supports full, incremental, and archivelog backups for RDS Custom databases.
- You can restore backups to the original RDS Custom instance or to an alternate location (such as another RDS Custom instance or an on-premises Oracle environment).
- Cohesity does not support recovery to different accounts with different regions, but recovery to different accounts with the same region, the same account with the same region, and the same account with different regions are supported.
- Point-in-time recovery (PITR) is supported for RDS Custom databases if archive log retention is configured to meet Cohesity backup requirements and AWS automated backups are enabled. PITR allows you to select a specific time within the retention period and restore the database to that point in time.
- Cohesity supports full, incremental, and archived log backups via Oracle RMAN using the Cohesity Adapter.

- Recovery of an RDS Custom database encrypted with the default AWS KMS Customer Managed Key (CMK) to a different account is not supported because AWS does not allow sharing snapshots encrypted with the default KMS key.
- You cannot restore an RDS Custom database to a location (same or alternate) if the target already has a database with the same name. AWS prohibits multiple RDS Custom database instances with identical names.
- Workaround: Assign a unique name to the RDS Custom database instance before performing recovery.
- When creating a protection group for Linux sources, choose either Imagecopy or Backupset (SBT) as the backup type; these types are not interchangeable.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Snapshot method (via native AWS snapshots)

To protect Amazon Custom RDS for Oracle with Cohesity, you must register your AWS account as a source in the Cohesity cluster and perform backup and recovery operations. Ensure you meet the prerequisites and understand the considerations before registering for the AWS account.

Prerequisites

- Non-overlapping backup windows: Schedule Amazon RDS Custom backup and maintenance windows so they do not overlap with Cohesity backup jobs.
- **Archive log retention**: Configure RDS Custom archive log retention to ensure that archived logs are retained long enough for at least one Cohesity backup cycle to complete successfully.

  For example, create the following configuration file on your RDS Custom instance:

  ```
  sudo mkdir -p /opt/aws/rdscustomagent/config
  sudo vi /opt/aws/rdscustomagent/config/redo_logs_custom_
  configuration.json
  Add the following content:
  {
  "archivedLogRetentionHours": "60"
  }
  ```

  This ensures archived logs are retained for 60 hours, preventing them from being purged before Cohesity backup jobs complete.

  For more details, see Configuring archived log retention for Amazon RDS Custom.
- Required Permissions - IAM requires the user permissions listed below to execute.

To register an AWS account with the Cohesity cluster, you need to run the CloudFormation template on the AWS console. The IAM user you use must have the required permissions to run the CloudFormation Template and to create and view the stack. The following table lists the permissions the IAM user must have to run the CloudFormation template.

> **Note:** Ensure to add these permissions manually.

Hypervisor	Cluster Type
cloudformation:CreateStack	To create a new stack.
cloudformation:DeleteStack	To delete a stack.
cloudformation:ValidateTemplate	To validate the CFT before deployment.
cloudformation:DescribeStacks	To view the status of stacks.
cloudformation:ListStacks	To view the status of stacks.
cloudformation:GetTemplate	To fetch template for an existing stack.
cloudformation:DescribeStackEvents	To view stack creation progress/errors.
cloudformation:DescribeStackResources	To see resources created in the stack.

- The following table lists the permissions that are added to the policy of the CloudFormation template.

> **Note:** You do not need to add these permissions manually, as they are automatically added when you run the CFT.

Resource	Permissions	Reason
ebs	ebs: GetSnapshotBlock ebs:ListChangedBlocks ebs:ListSnapshotBlocks	These permissions are required for EBS direct APIs to read & write data from/to EBS snapshots.

Resource	Permissions	Reason
ec2	ec2:AllocateAddress ec2:AssociateSubnetCidrBlock ec2:AssociateVpcCidrBlock ec2:AttachInternetGateway ec2:AttachNetworkInterface ec2:AttachVolume ec2:AttachVpnGateway ec2:AuthorizeSecurityGroupIngress ec2:AuthorizeSecurityGroupEgress ec2:CancelImportTask ec2:CopySnapshot ec2:CreateFleet ec2:CreateImage ec2:CreateInstanceExportTask ec2:CreateKeyPair ec2:CreateNetworkInterface ec2:CreateSecurityGroup ec2:CreateSnapshot ec2:CreateTags ec2:CreateVolume ec2:DeleteNetworkInterface ec2:DeleteSnapshot ec2:DeleteTags ec2:DeleteVolume ec2:DeregisterImage ec2:DescribeAvailabilityZones ec2:DescribeExportTasks ec2:DescribeFleetHistory ec2:DescribeFleetInstances	Required for AWS source registration, and discover the resources present in the account, which will be used for backups. Also needed for recovery to provide list of options to choose from.

Resource	Permissions	Reason
	ec2:DescribeFleets ec2:DescribeImages ec2:DescribeImportImageTasks ec2:DescribeImportSnapshotTasks ec2:DescribeInstanceAttribute ec2:DescribeInstanceCreditSpecifications ec2:DescribeInstances ec2:DescribeInstanceStatus ec2:DescribeKeyPairs ec2:DescribeNetworkInterfaces ec2:DescribeRegions ec2:DescribeReservedInstances ec2:DescribeReservedInstancesListings ec2:DescribeReservedInstancesModifications ec2:DescribeReservedInstancesOfferings ec2:DescribeSecurityGroupReferences ec2:DescribeSecurityGroups ec2:DescribeSnapshotAttribute ec2:DescribeSnapshots ec2:DescribeSubnets ec2:DescribeTags ec2:DescribeVolumeAttribute ec2:DescribeVolumes ec2:DescribeVolumeStatus ec2:DescribeVpcAttribute ec2:DescribeVpcClassicLink ec2:DescribeVpcClassicLinkDnsSupport ec2:DescribeVpcEndpointConnectionNotifications ec2:DescribeVpcEndpointConnections	

Resource	Permissions	Reason
	ec2:DescribeVpcEndpoints ec2:DescribeVpcEndpointServiceConfigurations ec2:DescribeVpcEndpointServicePermissions ec2:DescribeVpcEndpointServices ec2:DescribeVpcPeeringConnections ec2:DescribeVpcs ec2:DetachNetworkInterface ec2:DetachVolume ec2:ImportImage ec2:ImportInstance ec2:ImportKeyPair ec2:ImportSnapshot ec2:ImportVolume ec2:ModifyFleet ec2:ModifyInstanceAttribute ec2:ModifyNetworkInterfaceAttribute ec2:ModifySnapshotAttribute ec2:ModifyVolume ec2:ModifyVolumeAttribute ec2:RegisterImage ec2:RunInstances ec2:StartInstances ec2:StopInstances ec2:TerminateInstances	

Resource	Permissions	Reason
iam	iam:AddRoleToInstanceProfile iam:AttachRolePolicy iam:CreatePolicy iam:CreateRole iam:DeleteRole iam:DeleteRolePolicy iam:GetAccessKeyLastUsed iam:GetInstanceProfile iam:GetPolicy iam:GetRole iam:GetRolePolicy iam:GetUser iam:GetUserPolicy iam:ListAccessKeys iam:ListPolicies iam:ListRolePolicies iam:ListRoles iam:ListRoleTags iam:PassRole iam:PutRolePolicy iam:TagRole iam:UntagRole iam:UpdateRole iam:UpdateUser	

Resource	Permissions	Reason
kms	kms:CreateGrant kms:Decrypt kms:DescribeKey kms:Encrypt kms:GenerateDataKey kms:ListAliases kms:ListKeys kms:ReEncryptFrom kms:ReEncryptTo	KMS permissions are needed to read data of an encrypted database at the time of backup, as well as write encrypted data to the recovered database. Describe permissions are needed so we can list & identify keys associated with database instances.

Resource	Permissions	Reason
rds	rds:AddTagsToResource rds:CopyDBClusterSnapshot rds:CopyDBSnapshot rds:CreateDBClusterSnapshot rds:CreateDBInstance rds:CreateDBSnapshot rds:DeleteDBClusterSnapshot rds:DeleteDBSnapshot rds:DescribeDBClusterSnapshots rds:DescribeDBClusters rds:DescribeDBInstances rds:DescribeDBParameterGroups rds:DescribeDBSnapshots rds:DescribeDBSubnetGroups rds:DescribeOptionGroups rds:DescribeSourceRegions rds:ModifyDBClusterSnapshotAttribute rds:ModifyDBSnapshotAttribute rds:RestoreDBClusterFromSnapshot rds:RestoreDBInstanceFromDBSnapshot rds:RestoreDBInstanceToPointInTime	These permissions are required to register the AWS account on Cohesity with the IAM role which got created by the Cloud Formation template. Once the source is registered with Cohesity, described permissions are needed so Cohesity can identify resources present in the account, which will be used for backups as well as at the time of recovery we use this information to provide a list of options to choose from. Cohesity creates snapshots of the Amazon RDS & Aurora instances while backing up and storing the different database instance attributes and tags. While recovering the database instance, Cohesity creates DB instance/cluster, it also attaches the original tags. Cohesity requires the delete snapshots permissions to delete the expired/old snapshots it creates during the backup. We need to modify snapshot attributes permission so that we can share the snapshot across accounts if cross-account recovery is attempted. RestoreDBInstanceToPointInTime and RestoreDBClusterToPointInTime is needed to do the point in time recoveries.

Resource	Permissions	Reason
s3	s3:AbortMultipartUpload s3:CreateBucket s3:DeleteBucket s3:DeleteObject s3:DeleteObjectTagging s3:DeleteObjectVersion s3:DeleteObjectVersionTagging s3:GetBucketAcl s3:GetBucketLocation s3:GetBucketNotification s3:GetBucketObjectLockConfiguration s3:GetBucketOwnershipControls s3:GetBucketPolicy s3:GetBucketTagging s3:GetBucketVersioning s3:GetInventoryConfiguration s3:GetLifecycleConfiguration s3:GetObject s3:GetObjectAcl s3:GetObjectAttributes s3:GetObjectTagging s3:GetObjectTorrent s3:GetObjectVersion s3:GetObjectVersionAcl s3:AbortMultipartUpload s3:CreateBucket s3:DeleteBucket s3:DeleteObject s3:DeleteObjectTagging	These permissions are required for the backup and recovery of Amazon S3 objects.

Resource	Permissions	Reason
	s3:DeleteObjectVersion s3:DeleteObjectVersionTagging s3:GetBucketAcl s3:GetBucketLocation s3:GetBucketNotification s3:GetBucketObjectLockConfiguration s3:GetBucketOwnershipControls s3:GetBucketPolicy s3:GetBucketTagging s3:GetBucketVersioning s3:GetInventoryConfiguration s3:GetLifecycleConfiguration s3:GetObject s3:GetObjectAcl s3:GetObjectAttributes s3:GetObjectTagging s3:GetObjectTorrent s3:GetObjectVersion s3:GetObjectVersionAcl s3:GetObjectVersionAttributes s3:GetObjectVersionTagging s3:GetObjectVersionTorrent s3:HeadObject s3:ListAllMyBuckets s3:ListBucket s3:ListBucketMultipartUploads s3:ListMultipartUploadParts s3:PutBucketAcl s3:PutBucketNotification	

Resource	Permissions	Reason
	s3:PutBucketPolicy s3:PutBucketPublicAccessBlock s3:PutBucketTagging s3:PutInventoryConfiguration s3:PutObject s3:PutObjectAcl s3:PutObjectRetention s3:PutObjectTagging s3:PutObjectVersionAcl s3:PutObjectVersionTagging s3:RestoreObject	
ssm	ssm:GetCommandInvocation ssm:SendCommand ssm:UpdateInstanceInformation ssmmessages:*	
sqs	sqs:CreateQueue sqs:DeleteQueue sqs:DeleteMessage sqs:GetQueueUrl sqs:PurgeQueue sqs:ReceiveMessage sqs:GetQueueAttributes sqs:SetQueueAttributes sqs:TagQueue sqs:ListQueues	These permissions are required for capturing the incremental changes on the Amazon S3 buckets.

Resource	Permissions	Reason
event	events:PutRule events:DeleteRule events:PutTargets events:RemoveTargets events:TagResource events:ListRules	These permissions are required for capturing the incremental changes on the S3 buckets
glue	glue:CreateJob glue:DeleteJob glue:StartJobRun glue:GetJobRun glue:UpdateJob glue:TagResource	These permissions are required for sorting the inventory report. The sorted inventory report is then used to back up the Amazon S3 objects to the Cohesity cluster.

Resource	Permissions	Reason
dynamodb	dynamodb:CreateTable dynamodb:DescribeContinuousBackups dynamodb:DescribeExport dynamodb:DescribeImport dynamodb:DescribeTable dynamodb:ExportTableToPointInTime dynamodb:ImportTable dynamodb:ListTables dynamodb:ListTagsOfResource dynamodb:TagResource dynamodb:UpdateContinuousBackups dynamodb:UpdateTable dynamodb:BatchWriteItem dynamodb:DeleteItem dynamodb:GetItem dynamodb:PutItem dynamodb:Query dynamodb:RestoreTableToPointInTime dynamodb:Scan dynamodb:UpdateItem	These permissions are required for backing up and recovering Amazon DynamoDB tables.
CloudWatch Logs	logs:DescribeLogGroups logs:CreateLogGroup logs:CreateLogStream logs:DescribeLogStreams logs:GetLogEvents logs:PutLogEvents logs:PutRetentionPolicy	

Ports Used for Communication

For information on port numbers required, see Manage Firewall Ports.

Considerations

- To back up and recover Amazon RDS Custom databases, the cluster must be a Next-Gen Cloud Edition (NGCE) cluster.
- Cohesity uses the Oracle Adapter installed on the RDS Custom instance to perform database-level ingest backups and restores.
- Cohesity supports full, incremental, and archivelog backups for RDS Custom databases.
- You can restore backups to the original RDS Custom instance or to an alternate location (such as another RDS Custom instance or an on-premises Oracle environment).
- Cohesity does not support recovery to different accounts with different regions, but recovery to different accounts with the same region, the same account with the same region, and the same account with different regions is supported.
- Point-in-time recovery (PITR) is supported for RDS Custom databases if archive log retention is configured to meet Cohesity backup requirements and AWS automated backups are enabled. PITR allows you to select a specific time within the retention period and restore the database to that point in time.
- Recovery of an RDS Custom database encrypted with the default AWS KMS Customer Managed Key (CMK) to a different account is not supported because AWS does not allow sharing snapshots encrypted with the default KMS key.
- You cannot restore an RDS Custom database to a location (same or alternate) if the target already has a database with the same name. AWS prohibits multiple RDS Custom database instances with identical names.
- Workaround: Assign a unique name to the RDS Custom database instance before performing recovery.
- If you are protecting an Amazon RDS Custom instance within a private VPC, then you must:
- Register the AWS cloud source using an IAM role with permissions for RDS Custom, EC2, KMS, CloudWatch Logs, and STS. For more information, see Register an AWS Cloud Source using IAM Role.
- Ensure the following service endpoints are accessible from the Cohesity cluster and the RDS Custom EC2 instance:
  - `iam.amazonaws.com`
  - `sts.amazonaws.com (AWS Security Token Service)`
  - `kms.<region>.amazonaws.com`
  - `rds-custom.<region>.amazonaws.com`

- `ec2.<region>.amazonaws.com`
- `logs.<region>.amazonaws.com (CloudWatch Logs)`

- Enable AWS Security Token Service (STS) in the default AWS region where the Cohesity cluster is deployed.

Register and Manage Amazon RDS Custom for Oracle

Cohesity supports both ingest-based and snapshot-based backup of Custom Relational Database Service (RDS) databases for Oracle:

- Register Amazon RDS Custom for Oracle using the Ingest Method
- Register Amazon RDS Custom for Oracle using the Snapshot Method

Register Amazon RDS Custom for Oracle Using Ingest Method

To protect your Amazon RDS Custom for Oracle instances using the ingest method, follow the steps below:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** > **Databases** > **Oracle Source**.

   > **Note:** A physical server refers to a bare metal server or a VM running the supported versions of the OS and Oracle Database with the Cohesity agent installed.

3. In the **Select Source Type** drop-down, select **Oracle Standalone Source**.
4. **Enter Host Address**, for an Oracle Standalone server, specify the Hostname, Server IP, or FQDN of the Oracle server.
5. **Authentication Type**, specify the Authentication Type to register the Oracle source.

   When connecting to the Oracle databases, Cohesity can use either the Oracle OS Authentication or Oracle DB Authentication method. If you use DB authentication, provide the database username and password.
6. Click **Register**.
7. Select one of the following options from the **Cohesity Network Interface**:
   - **Auto Select**: If you select this option, the system automatically selects the interface group to use for the VLAN.
   - **Interface Group**: If you select this option, you must select the interface group to use for the source.

> **Note:** If you select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for all the objects in the protection job.

Register Amazon RDS Custom for Oracle Using Snapshot Method

To protect your Amazon RDS Custom for Oracle instances, you need to register your AWS account as a source on the Cohesity cluster. After registering your AWS account as a source, you can:

- Update the AWS source configuration
- Refresh the source details
- Unregister the AWS source from the Cohesity cluster

Register AWS Account

Before registering AWS account as a source, make sure you have completed all prerequisite tasks and review the considerations.

To register AWS account as a source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click **Register** at the top-right of the page and then select **Cloud Platforms** > **AWS**.

   The **Register AWS Source** dialog box appears.
3. Enter the **AWS Account ID**.
4. Select the AWS Services you want to register:
   - EC2
   - RDS
   - S3
5. (Applies only for Amazon S3 service) Cohesity utilizes the Amazon S3 inventory report to protect the Amazon S3 bucket. Under **Inventory Report Location**, provide the following information to create the inventory report:
   - **S3 buckets ARN**: Enter the ARN of the Amazon S3 bucket where you want to create the inventory report. The ARN you provide must be of only those Amazon S3 buckets that are in the same AWS account and cloud region as the Amazon S3 you want to protect.
   - **Prefix**: Add a prefix value to the name of the inventory report that will be created.

- The prefix can be of any character and can also include white spaces.
- For example, Report-Source Bucket.
- The prefix should not begin or end with a forward slash (/) .
- The prefix should not contain consecutive forward slashes.
- You must not upload any files in the prefix of the inventory report.

The inventory report will be created at `<Prefix>/<Path_to_inventoy_report>` when you initiate the protection. Cohesity creates the path to the inventory report on AWS based on the cluster, Amazon S3 bucket, and protection.

6. Click **Next** to generate a CloudFormation Template (CFT), which you will use to complete the AWS source registration.
7. Click **Download** to download the CloudFormation Template.
8. On running the CFT, IAM roles and policies are created.
9. Once the roles and policies are created successfully, the **Register AWS Source** dialog box will indicate the account authentication status as **Confirmed**.
10. Once account authentication is successful, click **Register**.

    If authentication fails, contact Cohesity Support.

Run CloudFormation Template

To complete the registration of the AWS account with the Cohesity cluster, you must run the Cohesity's CloudFormation Template on the AWS account you are registering. You can download CFT from the Cohesity cluster UI as you register the AWS source.

When you execute the CFT, it creates an AWS IAM role and the policies for accessing all IAM roles and KMS keys in your account.

Executing the CFT Template

1. In the **AWS Management Console**, navigate to **Services** and search for **CloudFormation**.
2. Click **Create stack** > **With new resources** to create a stack under AWS cloudFormation using the downloaded CloudFormation Template.
    a. In the **Prerequisite - Prepare template** page, select **Upload a template** file.
    b. Click **Choose File** to browse and upload the JSON Cloudformation template. Click **Next**.
    c. Enter a **Stack name** and specify the AWS Account ID where the Cloud Edition cluster is running.
    d. Select **source_registration** for creating IAM role for registering AWS source.

e. Click **Next**.

f. Acknowledge the creation of IAM resources by selecting the checkbox, then click **Create Stack**.

g. Once the stack creation is successful, view the **Resources** tab to check the created IAM roles.

h. The following IAM roles and profiles are created:

- CohesityInstanceRole
- CohesityInstanceProfile
- CohesitySourceRegistrationRole

Attaching the IAM Role to Cohesity Nodes

To complete registration, attach the **CohesityInstanceProfile** role to all nodes in the cluster:

1. Navigate to **Services** > **EC2** > **Instances**.
2. Select the cluster created. For a multi node, the IAM role must be attached on all nodes.
3. Select **Actions** > **Instance Settings** > **Attach/Replace IAM Role**.
4. In the **IAM Role** field, enter the name/search for the Instance Profile that was created.

   Example: *CohesityInstanceProfile*
5. Click **Apply**.

   The IAM role is successfully attached to that VM.
6. Click **Close**.

> **Note:** If any new node is added to the cluster, then this role has to be attached to every newly added node in the cluster as well.

Complete AWS Source Registration

Once the IAM role is attached, complete the AWS source registration on the cluster. For more information, to register an AWS Cloud Source using IAM Role.

Update AWS Account Source Configuration

You can edit the registered AWS account to add or remove the Amazon services protected by the Cohesity from your AWS account.

To edit the AWS account source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AWS account source and select Edit.
3. In the **Edit AWS Source** page, select or unselect the Amazon services.
4. Click **Update**.

Refresh AWS Account Source

You can refresh the AWS account source configuration and fetch the latest changes on your AWS account.

To refresh the AWS account source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the AWS account source and select **Refresh**.

Unregister AWS Account Source

If you plan to stop taking the backup of your AWS source, you can unregister the AWS source from the Cohesity Cluster.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the AWS source.

To unregister the AWS source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AWS account source and select **Unregister**.

Backup Amazon RDS Custom for Oracle

Cohesity supports both ingest-based and snapshot-based backup of Custom Relational Database Service (RDS) databases for Oracle:

- Add or Edit a Protection Group for Amazon RDS Custom for Oracle using Ingest Method
- Add or Edit a Protection Group for Amazon RDS Custom for Oracle using Snapshot Method

Protect Amazon RDS Custom for Oracle Using Ingest Method

Cohesity supports ingest-based backup of Custom Relational Database Service (RDS) databases for Oracle.

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might

find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit an Oracle Source
- Create or Edit a Standard Policy

To create a Protection Group for Custom RDS for Oracle:

1. Navigate to **Data Protection > Protection** and do one of the following:
   - To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
   - To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. If creating a Protection Group for **Custom RDS**, select **Databases** > **Oracle Database**.
3. **Name**: Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
4. **Source**: Select the database you want to protect. You can also register a new source with the Cohesity cluster so you can protect it.
5. Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box indicates the object is selected to be backed up.
6. Click the arrow adjacent to the Object (**Expand All**) to expand the entire Object hierarchy. Click the arrow again (**Collapse All**) to collapse the entire Object hierarchy.
7. The graphical representations in the table are described below.

   When creating a new Protection Group, a shield next to the Object name indicates the Object has already been protected by another Protection Group.

   Click the icon to see details about the object.

   If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.
8. Click **Save Selection**.
9. **Policy**: Select an existing protection policy or create a new one by selecting **New Policy**.

10. **Backup Type**: When creating a protection group for Linux sources, choose either Imagecopy or Backupset (SBT) as the backup type. Backup type selection is possible while creation of the protection job. They are not interchangeable.
    a. **Image Copy**: Use this option if you prefer NFS-based backups and require forever incremental backup capability. This is selected by Default.
    b. **Backup Sets**: Use this option to enable encryption in transit using gRPC and to avoid NFS client dependencies.

    > **Note:**
    > - You must create or select a policy with the **Periodic Full Backup** option set when using Oracle SBT.
    > - This is an Early Access feature. Contact your Cohesity account team to enable the feature.

11. **Start Time**: Indicates when the Protection Group should run. The current time is displayed by default, but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.

12. **End Date**: Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion, even if it completes after this date.
13. **Cancel Runs at Quiet Time Start**: Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default, this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
    a. **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    b. **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    c. **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.
14. **Priority**: Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to

process them, those with High priority are given the highest priority to execute, Medium the second-highest priority, and Low the lowest priority.

15. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

    The connection to the SMTP Server must be enabled to send email notifications. You can enable the SMTP Server connection in the configuration settings page (**Cohesity UI > Settings > Summary > Configure > Enable SMTP Server**).

16. **SLA**: The service-level agreement (SLA) defines the expected time to complete either a full backup or an incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.

    a. **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.

    b. **Full**: Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).

17. **Description**: Optional. Click the plus sign and enter a description for the Protection Group.

18. Click **Protect**.

The Protection page is displayed and includes the new Protection Group.

Protect Amazon RDS Custom for Oracle Using Snapshot Method

Cohesity supports snapshot-based backup of Custom Relational Database Service (RDS) databases for Oracle.

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit a Hypervisor Source
- Create or Edit a Standard Policy
- Create or Edit Storage Domains

To create a Protection Group for Custom RDS for Oracle:

1. Navigate to **Data Protection > Protection** > **Databases** > **Amazon RDS** and do one of the following:

    1. To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.

   2. To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. **Source**: Select the database you want to protect. You can also register a new source with the Cohesity cluster so you can protect it.
3. Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box indicates the object is selected to be backed up.
4. Click the arrow adjacent to the Object (**Expand All**) to expand the entire Object hierarchy. Click the arrow again (**Collapse All**) to collapse the entire Object hierarchy.
5. The graphical representations in the table are described below.

   When creating a new Protection Group, a shield next to the Object name indicates the Object has already been protected by another Protection Group.

   Click the icon to see details about the object.

   If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.
6. Click **Save Selection**.
7. **Policy**: Select an existing protection policy or create a new one by selecting **New Policy**.

   > **Note:** You must create or select a policy with the **Periodic Full Backup** option set.

8. **Start Time**: Indicates when the Protection Group should run. The current time is displayed by default, but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

   The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
9. **SLA**: The service-level agreement (SLA) defines the expected time to complete either a full backup or an incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
   a. **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the

differences (changed blocks) since the last protection run.

b. **Full**: Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).

10. From the Backup Options (Optional), configure the following **Additional Settings**:
    a. **Pause Future Runs**: Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
    b. **End Date**: Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion, even if it completes after this date.
    c. **Alerts**: Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
        i. **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
        i. **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
        ii. **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
    d. **Priority**: Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
    e. **Description**: Optional. Click the plus sign and enter a description for the Protection Group.
11. Click **Protect**.

The Protection page is displayed and includes the new Protection Group.

Recover Amazon RDS Custom for Oracle

Cohesity supports both ingest-based and snapshot-based recoveries of Custom Relational Database Service (RDS) databases for Oracle:

- Recover Amazon RDS Custom for Oracle Using Ingest Method
- Recover Amazon RDS Custom for Oracle Using Snapshot Method

Recover Amazon RDS Custom for Oracle Using Ingest Method

Cohesity provides the ability to recover the Custom Relational Database Service (RDS) databases for Oracle to its original location or an alternate location. For more information, see:

- Recover Custom RDS to the Original location
- Recover Custom RDS to an Alternate Location
- Create Cohesity View with DB Files - Custom RDS database

Recover Custom RDS to the Original location

To recover the Custom RDS to the original location:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Databases** > **Oracle Database**.
3. In the Recover page, search for the Oracle database instance that needs to be restored or use a wildcard name to search for a specific database instance.
    a. **CDB**: Search the Oracle CDB that you want to restore.
        1. Select the entire CDB (with all the PDBs) or a CDB and a subset of PDBs associated with that CDB for restore.
        2. Select the CDB and click **Next**.
    b. **PDB**: Search for the Oracle PDB that you want to restore.
        1. You can either search for the parent CDB associated with the PDB or directly search for the PDB.
        2. Select the PDB and click **Next**.
    c. **Oracle Database**: Search for the Oracle Database you want to restore.
4. Accept the latest Recovery Point or click the change icon and select a different recovery point.
    a. **Snapshots**: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
    b. **Point-in-Time**: Use the slide control to adjust the date and time and click **Save**.
5. Click **Next: Recover Options**.
6. From the **Recover To** options, select **Original Location** to recover the Custom RDS database to the original location.

    > **Note:** For PDB overwrite restore, PDB can be rolled back only within the current CDBs set.

7. From the Recovery Options (Optional), configure the following **Additional Settings**:

Field Name	Description
Leave Database in Recovery Mode:	Select the checkbox if you do not want the recovered database to be opened.
Shell Environment	Toggle on the Shell Environment to configure your shell environment that executes the Cohesity restore workflow using this setting. For example, when restoring the Oracle EBS database, the restore environment often needs to define a different TNS_ADMIN value to use a different sqlnet.ora file in the same Oracle Home shared with other databases. You can use the Shell Environment to define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your restored target database. For TDE support, the wallet location for your restore target database may depend on a shell variable. This is the case when your restored database wallet location is specified either by the ENCRYPTION_ WALLET_LOCATION or WALLET_LOCATION keyword in your sqlnet.ora file. A typical value for these parameters will reference a shell variable. For Example, ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE) (METHOD_DATA=(DIRECTORY=/etc/ORACLE/WALLETS/$ORACLE_UNIQUE_NAME/))) In this example, you will need to set ORACLE_UNIQUE_NAME in the SHELL ENVIRONMENT setting in the GUI to set up the Cohesity restore shell environment for the Oracle restore to find the Keystore.
Cluster Interface	Select one of the following options to assign the interface group for the recovery job. 1. **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration. 2. **Interface Group**: Select the interface group to use for the recovery. If you select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for the recovery job.
Task Name	(Optional) Provide the name for the recovery task. Click **Recover**.

The recovery process is initiated. You can monitor the recovery task from the Recoveries page.

Recover Custom RDS to the Alternate location

To recover the Custom RDS to the alternate location:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click Recover at the top right of the page and then select **Databases** > **Oracle Database**.

3. In the Recover page, search for the Oracle database instance that needs to be restored or use a wildcard name to search for a specific database instance.
    a. **CDB**: Search the Oracle CDB that you want to restore.
        1. Select the entire CDB (with all the PDBs) or a CDB and a subset of PDBs associated with that CDB for restore.
        2. Select the CDB and click **Next**.
    b. **PDB**: Search for the Oracle PDB that you want to restore.
        1. You can either search for the parent CDB associated with the PDB or directly search for the PDB.
        2. Select the PDB and click **Next**.
    c. **Oracle Database**: Search for the Oracle Database you want to restore.
4. Accept the latest Recovery Point or click the change icon and select a different recovery point.
    a. **Snapshots**: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
    b. **Point-in-Time**: Use the slide control to adjust the date and time and click **Save**.
5. Click **Next: Recover Options**.
6. From the **Recover To** options, select **Alternate CDB/PDB** to recover the Custom RDS database to the new location. From the drop-down menu for target source, select a registered RDS custom Instance to which you want to restore the Custom RDS database:
    1. **Oracle Host**: The host to which to restore the database.
    2. **Configure Channels**: Click these options to configure the following:
        a. **System selects active node**: If you select this option, Cohesity will select the channels for the recovery job.
        b. Select specific nodes:
            i. **Hostname**: Select the hostname to which you want to perform the restore.
            ii. **Node**: Select the node to which you want to perform the restore. This option is valid only in the case of a multi-node cluster.

                In the case of PDB alternate restore it is mandatory to select the node as we have to provide the target CDB instance name.
            iii. **Channels**: Specify the number of channels to perform the restore

job. The system will run these channels in parallel to improve the restore performance.

7. In the **Recovery Option**, enter the following:
   a. **Restore Database Files to**: Specify the path to an existing empty directory. The newly created database files will reside in this path.
   b. **Oracle Home**: The ORACLE_HOME value for the host where the database is restored.
   c. **Base Directory**: The directory for the restored database.
   d. **Target Database Name**: The name for the alternate CDB database.

      If you are restoring a PDB, you must provide the target CDB on which the PDB is restored.

      > **Note:** When recovering to an alternate location, make sure a database with the same name does not exist on the Oracle server instance. Otherwise, the recovery task will fail.

8. In the Recovery Options (Optional), configure the following **Additional Settings**:

Field Name	Description

Pfile Parameters	**Pfile Parameters**: The Pfile parameters option allows you to see the target database parameters before the restore job is initiated. The Pfile inherits some source database parameters to the target database and allows you to edit the target database parameters based on your requirement by specifying the additional parameters as a `key=value` pair.
	
	The Pfile lists the following parameters:
	
	• **Default Parameters**: Default parameters of the target database and you can edit the default
	• **Inherited Parameters**: Parameters from the source database and you can edit the inherited parameters, if required.
	• **Restricted Parameters**: You cannot edit the restricted parameters in the Pfile. The following are the restricted parameters:
	  • INSTANCE_NAME
	  • DB_NAME
	  • DB_UNIQUE_NAME
	  • DB_BLOCK_SIZE
	  • MAX_STRING_SIZE
	  • COMPATIBLE
	
	If you edit the following values in the **Recover Oracle** GUI, You must click the **Reset** button on the top-right corner to reset the Pfile and get the updated parameter values in the Pfile box.
	  • Restore database files to
	  • Oracle Home
	  • Oracle Base
	  • Database Name
	
	Cohesity Platform will create a target database with the specified parameters in the Pfile. For example, if `SGA_TARGET` is specified as a part of the Pfile parameters, the restore operation will use the value specified in the Pfile during the restore.

Rename Database	Rename the database, if required.

Field Name	Description
Enable Archive Log mode for the Database	Select the checkbox if you want to enable the redo log archiving on the recovered database.
BCT File Path	The BCT file path specifies the location where the block change tracking file will be created. If not provided, BCT is not enabled for the restored database.
Number of Temporary Files to Use	Number of temp files to rename from the source database.
Redo Log Group Settings	The Redo log group configuration stores all changes made to the database as they occur. Every instance of an Oracle Database has an associated redo log to protect the database in case of an instance failure. You can provide the number, size, and prefix of redo log groups along with the redo log group member path for the recovery database.
Leave Database in Recovery Mode	Select the checkbox if you do not want the recovered database to be opened.
Shell Environment	Toggle on the Shell Environment to configure your shell environment that executes the Cohesity restore workflow using this setting. For example, when restoring the Oracle EBS database, the restore environment often needs to define a different TNS_ADMIN value to use a different sqlnet.ora file in the same Oracle Home shared with other databases. You can use Shell Environment to define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your restored target database. For TDE support, the wallet location for your restore target database may depend on a shell variable. This is the case when your restored database wallet location is specified either by the ENCRYPTION_WALLET_LOCATION or WALLET_LOCATION keyword in your sqlnet.ora file. A typical value for these parameters will reference a shell variable. For Example, ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE) (METHOD_DATA=(DIRECTORY=/etc/ORACLE/WALLETS/$ORACLE_UNIQUE_NAME/))) In this example, you will need to set ORACLE_UNIQUE_NAME in the SHELL ENVIRONMENT setting in the GUI to set up the Cohesity restore shell environment for the Oracle restore to find the Keystore.

Field Name	Description
Cluster Interface	Select one of the following options to assign the interface group for the recovery job. 1. **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration. 2. **Interface Group**: Select the interface group to use for the recovery. If you select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for the recovery job.
Use NOFILENAMECHECK	Using NOFILENAMECHECK forces RMAN to skip checking whether target datafile names are the same as source datafile names.
Customize Destination Paths with SET NEWNAME	Use the SET NEWNAME command to define a new location for individual datafiles, tempfiles, tablespaces, or the whole database. If not specified, Cohesity uses DB_FILE_NAME_CONVERT, if not empty. Otherwise, it uses your input to restore the destination location. Use SET NEWNAME specifications to determine the location of your restored datafiles and logfiles You can specify DB_FILE_NAME_CONVERT together with this option to augment paths that are not affected by DB_FILE_NAME_CONVERT. Syntax Example: SET NEWNAME FOR DATAFILE 1 TO 'u01/home/test_1.dbf';
Task Name	(Optional) Provide the name for the recovery task.

9. Click **Recover**.

The recovery process is initiated. You can monitor the recovery task from the Recoveries page.

Create Cohesity View with DB Files - Custom RDS database

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click Recover at the top right of the page and then select **Databases** > **Oracle Database**.
3. In the Recover page, search for the Oracle database instance that needs to be restored or use a wildcard name to search for a specific database instance.

a. **CDB**: Search the Oracle CDB that you want to restore.
    1. Select the entire CDB (with all the PDBs) or a CDB and a subset of PDBs associated with that CDB for restore.
    2. Select the CDB and click **Next**.
b. **PDB**: Search for the Oracle PDB that you want to restore.
    1. You can either search for the parent CDB associated with the PDB or directly search for the PDB.
    2. Select the PDB and click **Next**.
c. **Oracle Database**: Search for the Oracle Database you want to restore.

4. Accept the latest Recovery Point or click the change icon and select a different recovery point.
    a. **Snapshots**: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
    b. **Point-in-Time**: Use the slide control to adjust the date and time and click **Save**.
5. Click **Next: Recover Options**.
6. From the **Recover To** options, select **Create Cohesity View with DB Files** to recover the Custom RDS database to the new location. From the drop-down menu for target source, select a registered Instance to which you want to create a Cohesity View with DB Files:
    1. **Oracle Host**: The host to which to restore the database.
    2. **Configure Channels**: Click these options to configure the following:
        a. **System selects active node**: If you select this option, Cohesity will select the channels for the recovery job.
        b. **Select specific nodes**:
            i. **Hostname**: Select the hostname to which you want to perform the restore.
            ii. **Node**: Select the node to which you want to perform the restore. This option is valid only in the case of a multi-node cluster.

            In the case of PDB alternate restore it is mandatory to select the node as we have to provide the target CDB instance name.
            iii. **Channels**: Specify the number of channels to perform the restore job. The system will run these channels in parallel to improve the restore performance.
7. In the **Recovery Option**, enter the following:

1. **View Name Identifier**: A view will be created with this name as identity.
2. **Cluster Interface**: Select one of the following options to assign the interface group for the recovery job.
    a. **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration.
    b. **Interface Group**: Select the interface group to use for the recovery.

    If you select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for the recovery job.
    c. **Task Name**: (Optional) Provide the name for the recovery task.

8. Click **Recover**.

The recovery process is initiated. You can monitor the recovery task from the Recoveries page.

Recover Amazon RDS Custom for Oracle Using Snapshot Method

Cohesity provides the ability to recover the Custom Relational Database Service (RDS) databases for Oracle to its original location or an alternate location. For more information, see:

- Recover Custom RDS to the Original location
- Recover Custom RDS to an Alternate Location

Recover Custom RDS to the Original location

To recover the Custom RDS to the original location:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Databases** > **Amazon RDS**.

1. Select **Instance** as the recovery task.
2. Search for the instance by entering characters of the instance name, or by specifying the wildcard character *.
3. From the list of search results, select the instance to recover. By default, the latest snapshot of the instance is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

    From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for *<instance_name>*** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the

option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*instance_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. You can also select any point on the green solid line on the slider to restore to a specific point in time. Selecting an invalid time from the slider automatically selects the closest available snapshot.
  b. Select the snapshots to be recovered.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.
- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster.
  d. Click **Select Recovery Point** and then go to step 5.

4. Click **Next: Recover Options**.
5. In the **Recover To** section, select any one of the following options based on where you want to restore the instance and then provide the required information:
   - **Original Location**: Restore the instance to the same region.

     > **Note:** If an RDS instance with the same name already exists, the restore will fail. Ensure the original instance is not present or restore with a different name.

     Type **Yes** to confirm.
   - **New Location**: Restore the instance to a different region. You can restore the instance to the same or a different region within the same account as the original RDS instance.

     Based on the registration level of the AWS source, select the recovery destination hierarchy:

- If the AWS source is registered at the organization level:
  - **Source**: Select the AWS organization that contains the destination instance. You can recover the instance within the same organization as the source instance.
  - **AWS Account**: Select the member account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.
- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.

To overwrite any existing instance with the same name, enable **Overwrite existing instance(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of RDS instance protected using the AWS Snapshot protection type is not supported.
> - If an RDS instance with the same name already exists, the restore will fail.

6. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
7. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups

compatible with the selected engine, a default group will be created at launch.

8. **DB Parameter Group**: Select the database (DB) parameter group to associate with the DB instance.
9. **DB Port:** Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port. The default port is 3306.
10. Configure the following **Recovery Options**:
    a. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.
11. **Task name**: Optionally, change the default name of the recovery task.
12. Click **Recover**.

Amazon Relational Database Service Custom for SQL Server

Cohesity supports protecting Amazon Relational Database Service (RDS) Custom for Microsoft SQL Server. This managed database service provided by AWS offers a balance between the automation of Amazon RDS and the flexibility of self-managing a database on Amazon EC2.

> **Note:** Cohesity supports the backup and recovery of Amazon RDS Custom for SQL Server on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster).

This chapter includes the following topics:

Planning and Preparation for Amazon RDS Custom for SQL Server

The following topics provide guidance that helps you get ready to implement MS SQL Server:

Requirements to Access Amazon RDS Custom for SQL Server

Cohesity DataProtect SQL Server agent must be installed on the corresponding EC2 instance that has the Amazon RDS Custom instance. You can then perform the backup and recovery operations as you would for a standalone SQL Server instance.

You must access Amazon RDS Custom to connect to the underlying EC2 instance and then to the database. Use remote access to the EC2 instance and then client tools to connect to the database.

1. **Prepare the EC2 Instance for Remote Desktop:**
    i. Ensure the VPC security group allows inbound RDP traffic (port 3389) from your source IP.
    ii. Create an RDS instance. For more information, see Creating an Amazon RDS DB Instance.
    The EC2 instance will also be created.
    iii. Configure the Windows firewall on the EC2 instance to allow remote desktop access.
    iv. On the EC2 instance, connect with the **Session Manager**.
    v. Run the following PowerShell command to open firewall:

    ```
    Set-NetFirewallRule -DisplayName "Remote Desktop - User Mode (TCP-In)" -Direction Inbound -LocalAddress Any -Profile Any
    ```

    vi. Retrieve the administrator password for the EC2 instance from AWS Secrets Manager using the key pair information.
    vii. Go to the **Secret Manager** and click on the required **secret name**.
    viii. Under the **Secret Value** section, select **Retrieve Secret Value**.
    ix. Copy the secret value in plain text.
2. **Connect to the EC2 Instance via RDP:**
    i. Use an RDP client (like Remote Desktop Connection on Windows) to connect to the EC2 instance using its private IP address and the administrator credentials. A jump server in a public subnet is often used for this.
    ii. On the EC2 instance, connect with the **RDP Client**.
    iii. Select **Get Password** at the bottom of the page.
    iv. Paste the secret value into the **Private Key Contents** field and click **Decrypt Password**.

v. Copy the decrypted password.

vi. Go back to the jump server and RDP into the EC2 instance with the administrator credentials and decrypted password.

3. **Connect to SQL Server:**

   i. From the RDP session on the EC2 instance, use SQL Server Management Studio (SSMS) or other SQL client tools to connect to the SQL Server database endpoint.

   ii. Use this endpoint to register your Microsoft SQL Server as a source on Cohesity DataProtect.

Requirements for Amazon RDS Custom for SQL Server Protection

Supported Software for Microsoft SQL Server Protection

For more information, see Supported Software.

Agent Minimum Permissions for Microsoft SQL Server Protection

You can use either the host LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

It is recommended to run the Cohesity Agent service with an Active Directory domain user account that is a member of the local administrator of the SQL Server host. The AD domain user account must be a member of the SQL sysadmin server role. The user account must have log-on rights to the SQL Server host in the local security policy of the SQL Server host.

If you do not use the LOCAL SYSTEM account, ensure the following for the chosen account:

- The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
- The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.
- The account must have the sysadmin role in the MS SQL Server instance.

Ports Used for Communication

- On physical servers or VMs with an ephemeral or installed agent, open the ports 445, 11113, 11117, and 50051. For more information, see Manage Firewall Ports.
- If the Windows Firewall is used:
  - Inbound rules:
    - Add a rule to accept SQL Server traffic and TCP connections on local port 1433.
    - Set Remote Port to All Ports.
  - Outbound rules (for MS SQL Server 2016 running on Windows 2016):

    Update the "Block network access for R local user accounts in SQL server instance MSSQLSERVER" rule by going to General tab > Action window > select "Allow the connection".

Prerequisites for Microsoft SQL Server Protection

- Verify MS SQL Server services are running.
- The **Readable Secondary** field in the Availability Group Properties of the SQL Server Availability Group (AAG) must be set to **Yes** if you have selected **Secondary Only** as the **Backup Preference setting for AAG**.
- MS SQL database and instance names must meet the Microsoft SQL Server naming requirements. Names containing illegal Windows file characters, such as double quotes, pipe symbols, greater than and less than signs (" | > <), a space at the beginning or end of the database name (applicable only for VDI-based backup) are not supported. For more information, see Naming Files, Paths, and Namespaces.
- On the Server's Windows system, Cohesity recommends to set the Power Plan to High performance.

Credentials and Privileges for Microsoft SQL Server Protection

There are three accounts you must consider when installing the Cohesity Agent:

- **Installation Account**: The account you use to log in to the host and run the installer.
- **Service Account**: The account under which the Cohesity Agent service runs on the SQL Server host. (Selected at the time of installation.)
- **SQL Server login account**: The SQL Server account by which the Cohesity Agent has access to the databases. (Configured after installation.)

Installation Account

When installing the Cohesity Agent on a Microsoft Windows host, you must log onto the SQL Server host with an account that has local administrator privileges. This is because the Windows agent installer installs the Cohesity Agent service. Microsoft requires the installation account to have administrator privileges on the host.

Choose the Service Account

The installer gives you the option between two types of service accounts:

- The LocalSystem account.
- A user account. Assigning privileges to the service account will differ depending on which type of account you choose. Cohesity recommends using either the LocalSystem account or an Active Directory Domain User account. See the chart below for a comparison of service account types.

Comparison of Service Account Types

SERVICE ACCOUNT	TYPE	DEFINITION	RECOMMENDATION
LocalSystem (Default)	LocalSystem Account (Built-in Host Account)	The LocalSystem account is a predefined local account on the host computer.	Recommended: The Local System account is a predefined local account used by the service control manager. It has extensive privileges on the local computer, and acts as the computer on the network. Its token includes the NT AUTHORITY\SYSTEM and BUILTIN\Administrators SIDs; these accounts have access to most system objects.
User Account	Domain User Account (Active Directory)	A user whose username and password are stored in Active Directory.	Recommended: A domain user account enables the service to take full advantage of the service security features of Windows and Microsoft Active Directory Domain Services. The service has whatever local and network access is granted to the account, or to any groups of which the account is a member.
	Local User Account (Host account)	In Windows, a local user is one whose username and encrypted password are stored locally, only on the computer itself.	Not recommended: When you log in as a local user, the computer checks its own list of users and its own password file to see if you are allowed to log in to the computer. The computer itself then applies all the permissions and restrictions that are assigned to you for that computer.

> **Note:** If you are unsure which account to use, choose the LocalSystem account. This account already has most of the required permissions. The agent service account can be changed later if needed.

Microsoft SQL Server on Linux OS

Cohesity provides backup and recovery solutions for Microsoft SQL Server databases running on the Linux Operating System.

Requirements

Operating System	Database	Notes
Linux OS RHEL 8 and higher versions **Note:** This is an Early Access feature. Contact Cohesity Support to enable the feature.	Microsoft SQL Server 2019	• Only VDI-based backups are currently supported on Linux OS. • High Availability Microsoft SQL Server configurations like FCI and AG are not supported.

Considerations

- Cross-platform database restores between Windows and Linux are currently not supported.
- For VDI-based backups to work, Microsoft requires the backup application account (Cohesity Agent's service account) and the Microsoft SQL Server service account to be added to each other's primary groups using the usermod command.

  For example:

  sudo usermod -a -G mssql cohesityagent

  sudo usermod -a -G cohesityagent mssql

  where, mssql is the user account of the mssql-server service and cohesityagent is the user account of the Cohesity Agent service.

  Following this, the Microsoft SQL Server and the Cohesity Agent services must be restarted for the changes to take effect.

  > **Note:** The usermod -a -G command may not work for domain accounts.

- Currently, only password-based SQL Server authentication is supported and Active Directory-based authentication is not supported for Linux.

  For supported features of SQL Server 2019 on Linux, see Editions and supported features of SQL Server 2019 on Linux.

Register SQL Server on Linux OS

- In **DataProtect as a Service**, select **Sources**.
- Click **+Register Source > Microsoft SQL > Start Registration**.
- In the **Register Microsoft SQL Server** page, select the **SaaS Connection** and click **Continue**.
- In the **Hostname or IP Address** field, enter the IP address of the server, FQDN of the server, or VIP of the SQL FCI. Cohesity recommends you to provide the FQDN of the server.
- Select the **Server Type** as Linux.
- Provide the username and password for the **Database Authentication**.
- Click **Complete**.

Supported Database Types for Microsoft SQL Server

Before proceeding to deploy the Cohesity agent for SQL data protection, we recommend you correctly identify the type of SQL database you want to protect. Below is a table of SQL database types that are supported by the Cohesity SQL agent, and the recommended protection methods.

To understand the concept of each method, see About Microsoft SQL Server Backup.

Database Types	Definition	Adapter-based Protection Method	
		File-based Protection	VDI-based Protection
Stand Alone	A SQL Server User database.	✅	✅
AG	A database that belongs to an availability group (AG). For each availability database, the availability group maintains a single read-write copy (the primary replica) and one to four read-only copies (secondary replicas). For details, see Always On availability groups: a high-availability and disaster-recovery solution in the Microsoft documentation.	✅	✅
FCI	A database that belongs to a Failover Cluster Instance (FCI), which is a single instance of SQL Server that is installed across Windows Server Failover Clustering (WSFC) nodes.	✅	✅

Database Types	Definition	Adapter-based Protection Method	
		File-based Protection	VDI-based Protection
CDC	A change data capture (CDC) enabled SQL database records activity and is applied to an internal table. This makes the details of the changes available in an easily consumed relational format. For details, see About Change Data Capture (SQL Server) in the Microsoft documentation.	✓	✓
TDE	A Transparent Data Encryption (TDE) enabled SQL database has its data files encrypted, known as encrypting data at rest. TDE performs real-time I/O encryption and decryption of the data and log files. For details, see Transparent Data Encryption (TDE) in the Microsoft documentation. **NOTE**: Cohesity does not manage the keys associated with the TDE database.	✓	✓
Filestream	A Filestream enabled SQL Database is a database that contains a table that stores its data outside the database on the host file system. SQL Filestream leverages the OS file system's ability to handle large objects like images, PDFs, Binary, Bmp, Docs, and other unstructured data. For details, see Filestream (SQL Server) in the Microsoft documentation.	✗	✓
SQL System	SQL Server maintains a set of four system-level databases (*master, model, msdb, tempdb*), which are essential for the operation of a server instance. System databases must be backed up after every significant update. The system databases that you must always back up include msdb, master, and model. System databases are not supported for log backups. For details, see Backup & restore: system databases (SQL Server) in the Microsoft documentation.	✓	✓

Now that you can identify the type of SQL database you want to protect, you can register it as a Cohesity source.

Microsoft SQL Server Dashboard

Access the MS SQL Dashboard by going to **Dashboard** and from the drop-down, select **MS SQL**.

The page's upper portion displays information such as metrics, status and recovery history. In the page's lower portion, you can start the different actions from the Hosts tab and the Databases tab. The MS SQL Dashboard displays local MS SQL databases only. To restore or clone replicated or CloudRetrieved databases, use their respective restore or clone workflows.

From the Hosts tab, mouse over the actions menu for the following actions (some actions are available for physical servers only):

- **Protect Databases** ( )—Create a Protection Group using this server.
- **Unregister MS SQL** ( )—Unregister the MS SQL application. The server remains registered with the Cohesity cluster.
- **Edit** ( )—Configure the settings of a registered MS SQL server.
- **Unregister** ( )—Unregister this server from the Cohesity cluster.
- **Refresh** ( )—Manually force a refresh of the object hierarchy.
- **Upgrade Agents** ( )—Upgrade the Cohesity agent on the server if an upgrade is available.

From the Databases tab, mouse over the actions menu for the following actions (some actions are not available for system databases):

- **Protect** ( )—Create a Protection Group using this database.
- **Run Now** ( )—Start a run of this Protection Group. You have the option to back up one or more objects.

  If the policy associated with the Protection Group is configured with replication and archive targets, you can select those targets. If the policy supports multiple backup types, you can select an incremental, full or log backup.

  VM MS SQL Protection Groups do not support selecting individual databases for run now; only the entire VM is selectable.
- **Pause Future Runs/Resume Future Runs** ( )
  - Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the job is currently executing a run, the current run continues to execute and only future runs of this job are paused.

    - Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this job.

- **Cancel this Protection Group Run** ( )—Cancel the current run of this Protection Group. (This option is only available when a Protection Group run is executing.) This cancellation only affects the current run of the Protection Group and not future runs.
- **Clone** ( )—Make a copy of the database. The clone database is a unique copy of the original database taken from the backup.
- **Recover** ( )—Restore a database. Follow the instructions provided in Recover Amazon RDS Custom for SQL Server.
- **Migrate** ( )—Migrate a database. Follow the instructions provided in Migrate Microsoft SQL Server Databases.

Database Recover and Clone Actions

In the Databases tab click a protected database to display available recover and clone actions.

To recover or clone the latest snapshot, click the arrow next to the **Protect** button at the top right corner and select an action.

By default, the latest snapshot is selected. You can optionally choose to restore from a different snapshot by clicking Edit ( ).

You can also provide a specific time to restore to if transaction logs are available.

Click a date on the calendar to select a snapshot from that date. The default point selected will be the latest log transaction or the latest full or incremental, whichever is the latest. The slider also displays available points in time and displays snapshots as circles. Manually select a point on the slider if you want to restore to a specific point in time. Selecting an invalid time from the slider automatically changes to the closest available snapshot. If your manual selection is within 10 minutes of a snapshot, your selection automatically snaps to it. You can also select a specific point in time by entering it in the time fields.

After selecting the desired recover point, click **Recover**. You can choose to clone the selected recover point by clicking the arrow adjacent to recover and selecting **Clone**.

In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

Register and Manage Amazon RDS Custom for SQL Server

To start protecting SQL Server databases, you need to register the SQL Server as a source in Cohesity. The following topics provide guidance that helps you register and manage SQL Server:

Auto-discover Microsoft SQL Server Topology when Registering Microsoft SQL Server

To register an MS SQL Server, MS SQL FCI, or an MS SQL AG of a Windows cluster, you can provide the IP address of the server, FQDN of the server, or VIP of the SQL FCI. The Cohesity cluster discovers the entire MS SQL topology on the Windows cluster and displays the relevant subset based on the MS SQL address you provide while registering. From the topology list, you can select the servers you want to register with the Cohesity cluster. To display the entire MS SQL topology, use the Windows Cluster IP/FQDN.

Prerequisites

Before you register the MS SQL, ensure to meet the following:

- MS SQL prerequisites. For more information, see Requirements for Amazon RDS Custom for SQL Server Protection.
- Download and Install the Agent.

Download and Install the Agent

Install the Cohesity agent on each SQL server that you want Cohesity to protect.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation using the following steps:
4. **Volume CBT (Changed Block Tracker):** Install this component to leverage Cohesity's built-in CBT technology. Without this component, only full backups can be performed. Installing this component requires a one-time reboot to load the Cohesity CBT driver. Until a reboot, only full backups can be performed.

   File System CBT (Changed Block Tracker): Install this component to perform MS SQL file-based protection.

   > **Note:** Installing the Volume CBT or File System CBT is not necessary for MS SQL VDI-based protection.

5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

   - The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.

- The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.

6. Wait for the installation to finish.

   In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.
7. The agent starts automatically.
8. Reboot the machine if you have installed the Volume CBT component.
9. Repeat the agent installation process on each SQL server you want to protect. This includes any standalone MS SQL servers and Microsoft SQL Server nodes with AGs.

Register MS SQL Server

1. Select **Data Protection > Sources**.
2. Click **Register** > **Databases** > **MS SQL Server (early access: auto-detect)**.

   The Register an MS SQL Server page appears.

   > **Note:** Ensure that the agent is installed on all the nodes that are part of the Always on Availability Group (AG) and Failover Cluster Instance (FCI). For more details, Download and Install the Agent.
3. In the **Hostname or IP Address** field, enter the IP address of the server, FQDN of the server, or VIP of the SQL FCI. Cohesity recommends you to provide the FQDN of the server.
4. Select the **Server Type**: Windows or Linux.

   For Linux, provide the username and password for the Database Authentication.

   > **Note:** This is an Early Access feature. Contact Cohesity Support to enable the feature.
5. Click **Register**.

   The Cohesity cluster auto-discovers the entire MS SQL topology on the cluster. Based on the MS SQL address you provide while registering, the Cohesity cluster displays the relevant subset of that MS SQL server. The Cohesity cluster also displays which version of the Cohesity agent is running on each of the servers.

   For AG and FCI, ensure to register all the nodes that are part of the AG and FCI.
6. From the topology list, select the entities you want to register with the Cohesity cluster and select **Complete Registration**.

   The selected entities are registered with the Cohesity cluster and are displayed on the source page of the Cohesity cluster UI.

AG is registered as a source, and the databases are listed under AG on the Sources page.

Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs

This topic describes how to set up standalone MS SQL servers and MS SQL Server Always On Availability Groups (AGs) to be protected by a Cohesity cluster.

> **Note:** The **Readable Secondary** field in the Availability Group Properties of the SQL Server Availability Group (AG) must be set to "Yes" or "Read-intent only".

The following are the different methods to register your physical or virtual MS SQL Servers allowing you to protect your MS SQL server databases:

- **Registering a SQL Server using the Cohesity SQL Agent**:

  This option is the preferred method to protect your MS SQL server databases on physical or virtual hardware. This option can be used to protect any SQL server configurations such as MSQL standalone, FCI, and AG. For more information, see Set Up MS SQL on Virtual or Physical Servers using the Agent.

  To set up an MS SQL FCI, follow the steps described in Set Up Microsoft SQL Server FCI.

- **Registering a Hypervisor for SQL Server**:

  Cohesity provides an additional option to protect your MS SQL server at the hypervisor level. While this option is available, it is not the preferred method to protect your SQL server databases. Following are the two available registration methods:

  - **Persistent Agent**—For more information, see Set Up MS SQL on VMs as a VMware vCenter application with the Installed Agent.
  - **Auto-Deployed Agent (non-Persistent)**—For more information, see Set Up MS SQL on VMs as a VMware vCenter application with the Auto-deployed Agent.

Ensure Prerequisites are Met

Ensure the prerequisites listed in Prerequisites for Microsoft SQL Server Protection are met.

Set Up MS SQL on Virtual or Physical Servers using the Agent

This option can be used for both physical or virtual servers that are in standalone or AG configurations. Cohesity recommends you use this method to protect the MS SQL Servers with the Cohesity cluster. For more information, see Set Up MS SQL on Virtual or Physical Servers using the Agent.

Download and Install the Agent

Install the Cohesity agent on each SQL server that you want Cohesity to protect.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation using the following steps:
4. **Volume CBT (Changed Block Tracker):** Install this component to leverage Cohesity's built-in CBT technology. Without this component, only full backups can be performed. Installing this component requires a one-time reboot to load the Cohesity CBT driver. Until a reboot, only full backups can be performed.

   File System CBT (Changed Block Tracker): Install this component to perform MS SQL file-based protection.

   > **Note:** Installing the Volume CBT or File System CBT is not necessary for MS SQL VDI-based protection.

5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

   - The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
   - The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.

6. Wait for the installation to finish.

   In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.
7. The agent starts automatically.
8. Reboot the machine if you have installed the Volume CBT component.
9. Repeat the agent installation process on each SQL server you want to protect. This includes any standalone MS SQL servers and Microsoft SQL Server nodes with AGs.

Register a Physical MS SQL Server

> **Important:** If the server contains one or more databases that belong to an Always On Availability Group (AAG), all remaining servers within the AAG must also be registered with the Cohesity cluster as MS SQL Server in order to be backed up. Navigate to **Data Protection > Sources > click on the required source > All Objects > Show AAG Nodes** to display the registered AAG nodes and any warnings about unregistered servers. If a replica is added to an existing AAG, the replica is automatically backed up by the next Protection Group run that protects that AAG, even if the AAG node list does not display it.

1. Select **Data Protection > Sources**.
2. Click **Register** > **Physical Server**.
3. Specify the hostname (FQDN) or server IP of the server to register. Cohesity recommends using the FQDN.
4. Select the **Server Type**: Windows or Linux.

   For Linux, provide the username and password for the Database Authentication.

   > **Note:** This is an Early Access feature. Contact Cohesity Support to enable the feature.

5. Click **Register**.

   When registration finishes, a new entry appears in the physical section of the Sources page.

6. Next to the newly added physical server, mouse over the actions menu (⋮) and click **Register as MS SQL Server**.

   When registration finishes successfully, the MS SQL Server information is displayed when you click the icon.

> **Note:**
> You can throttle the CPU capacity on the MS SQL physical servers (running on Windows 2012R2 and later) you registered with the Cohesity cluster. Throttling optimizes the CPU capacity utilized when a Protection Group with source-side deduplication is run on the server. However, for MS SQL Protection Group, source-side deduplication is not enabled by default. You must contact Cohesity Support for enabling the source-side deduplication for a Protection Group. CPU throttling does not affect the MS SQL Server if source-side deduplication in a SQL Protection Group is not enabled. For more information on enabling CPU throttling, see Network Bandwidth and CPU Throttling for Physical Server.

> **Note:**
> - AGs discovered on the registered standalone physical hosts or failover cluster are registered as sources internally.
> - After registration, databases belonging to an AG will appear only under the AG listed in Sources and not under the registered SQL host (which is an AG replica). SQL host lists only the local databases.

Set Up MS SQL on VMs as a VMware vCenter application with the Installed Agent

This method must be used only when you are registering SQL Server at the hypervisor-level. You can use this option for standalone SQL or SQL servers in AG running on the VMware vCenter.

Download and Install the Agent

Before registering the MSSQL server at the hypervisor-level, you must first install the Cohesity agent on each SQL server that you want Cohesity to protect.

You can optionally use the auto-deployed Cohesity agent, which does not require agent installation. Cohesity, however, recommends installing the agent because this avoids additional configuration such as disabling Windows User Account Control (UAC). If you still want to use the auto-deployed agent, follow the steps described in Set Up MS SQL on VMs as a VMware vCenter application with the Auto-deployed Agent.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.

3. As an administrator with local system privileges on that server, run the executable and complete the installation wizard.
4. You do not need to install any components. Backups leverage VMware VADP CBT or Hyper-V 2016 RCT.
5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

   - The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
   - The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.
6. Wait for the installation to finish.

   In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.
7. The agent starts automatically.
8. Repeat the agent installation process on each SQL server you want to protect. This includes any standalone MS SQL servers and MS SQL Server nodes with AGs.

Register MS SQL

> **Important:** If the server contains one or more databases that belong to an Always On Availability Group (AAG), all remaining servers within the AAG must also be registered with the Cohesity cluster as MS SQL Server in order to be backed up. Navigate to **Data Protection > Sources > click on the required source > All Objects > Show AAG Nodes** to display the registered AAG nodes and any warnings about unregistered servers. If a replica is added to an existing AAG, the replica is automatically backed up by the next Protection Group run that protects that AAG, even if the AAG node list does not display it.

1. Select **Data Protection > Sources**.
2. Navigate the hierarchy to the server that is running MS SQL.
3. Mouse over the server's actions menu (⋮) and click **Register Applications**.
4. Select **Installed Agent**.

   If the agent is not already installed, registration fails.

5. Select MS SQL.
6. Click **Register**.

   The agent starts automatically.

   The server's MS SQL and agent information display when you click the icon.

Set Up MS SQL on VMs as a VMware vCenter application with the Auto-deployed Agent

> **Important:** If the server contains one or more databases that belong to an Always On Availability Group (AAG), all remaining servers within the AAG must also be registered with the Cohesity cluster as MS SQL Server in order to be backed up. Navigate to **Data Protection > Sources > click on the required source > All Objects > Show AAG Nodes** to display the registered AAG nodes and any warnings about unregistered servers. If a replica is added to an existing AAG, the replica is automatically backed up by the next Protection Group run that protects that AAG, even if the AAG node list does not display it.

1. Select **Data Protection > Sources**.
2. Navigate the hierarchy to the server that is running MS SQL.
3. Mouse over the server's actions menu (⋮) and click **Register Applications**.
4. Select **Auto-Deployed Agent**.

   Enter the Windows administrator **Username** and **Password** for the server that is running MS SQL. Windows administrator privileges are required. Ensure that the user account used for auto-deploying the agent is added to all the MS SQL server's SQL instances' **Security > Logins** with Windows Integrated Authentication.

   Disable the UAC property Admin Approval Mode

   1. Click **Start**, type secpol.msc in the Search programs and files box, and then press Enter.
   2. If the User Account Control dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
   3. In the console tree, expand Local Policies, and click **Security Options**.
   4. In the details pane, scroll down and double-click the **Group Policy** setting that you want to change. This opens the UAC policy settings that a local administrator can modify.
   5. Disable **User Account Control: Run all administrators in Admin Approval Mode** for the Administrator account.
   6. Ensure that **User Account Control: Admin Approval Mode for Built-in Administrator Account** is set to Disabled.

7. Ensure that **User Account Control: Behavior of elevation prompt for administrators in Admin Approval Mode** is set to Elevate without Prompting.
8. Restart/reboot the Windows server for the changes to take effect.
9. Provide the Administrator/Domain Administrator user that has Admin Approval Mode disabled.

5. Select MS SQL.
6. Click **Register**.

The server's MS SQL information displays when you click the icon.

Upgrade the Cohesity Agent and Add the File System CBT Component

If necessary, you can upgrade the agent and add the File System Changed Block Tracking (CBT) manually. Start the upgrade and select the option to include the File System CBT component in the upgrade.

Troubleshoot

- The log backup of AG MS SQL Server fails with the following error if there is a break in the log chain.

  ```
  Log chain break error: Discovered a break in the log chain for
  <Database Names>
  ```

  Resolution: To resolve this issue, ensure that no other third-party applications are running a log backup, and then perform a full backup run to reset the log chain.

- The log backup of the AG MS SQL Server fails with the following error if a database is added or removed from the AG MS SQL Server.

  ```
  AG relationship error: Discovered a AG relationship error for database
  <Database Names>
  ```

  Resolution: To resolve this error, perform a full backup run.

For more troubleshooting information, see the following KB article:

- Collecting troubleshooting information for MS SQL issues

Log in to the Cohesity Support Portal to see more KB articles.

Set Up Microsoft SQL Server FCI

This topic describes how to set up an MS SQL Failover Cluster Instance (FCI) to be protected by a Cohesity cluster. To set up a standalone MS SQL Server or MS SQL AG, follow the steps described in Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs.

MS SQL FCI - Overview

An MS SQL FCI is an MS SQL instance installed on a Windows server that is part of a Windows Server Failover Cluster (WSFC or simply Windows Cluster for short). The SQL instance physically exists on (and is accessible from) only one node of the WSFC at any given time. If the FCI's current node goes down, the FCI moves to another node in the WSFC. The FCI is addressed by a Virtual IP (VIP) which points to the physical IP of the node on which it currently resides.

Set Up MS SQL FCI Overview

MS SQL FCI setup has the following steps:

1. Ensure Prerequisites Are Met
2. Download and Install the Agent on Each Node
3. Register an MS SQL FCI as a Source
4. Register an MS SQL FCI as MS SQL

Ensure Prerequisites Are Met

Ensure the prerequisites listed in Prerequisites for Microsoft SQL Server Protection are met.

Download and Install the Agent on Each Node

Before registering an MS SQL Server FCI, you must download and install the Cohesity agent on each node within the MS SQL FCI.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation using the following steps:
4. **Volume CBT (Changed Block Tracker):** Install this component to leverage Cohesity's built-in CBT technology. Without this component, only full backups can be performed. Installing this component requires a one-time reboot to load the Cohesity CBT driver. Until a reboot, only full backups can be performed.

   File System CBT (Changed Block Tracker): Install this component to perform MS SQL file-based protection.

   > **Note:** Installing the Volume CBT or File System CBT is not necessary for MS SQL VDI-based protection.

5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

- The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
- The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.

6. Wait for the installation to finish.

   In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.
7. The agent starts automatically.
8. Reboot the machine if you have installed the Volume CBT component.
9. Repeat the agent installation process on each node within the MS SQL FCI.

   After you install the agent on each node within the FCI, the MS SQL FCI can be registered as a source with the Cohesity cluster.

Register an MS SQL FCI as a Source

This registration step makes an MS SQL FCI visible to the Cohesity cluster so that its content is eligible to be protected or backed up.

> **Note:**
> - An MS SQL FCI can be registered with only one Cohesity cluster at a time.
> - Each MS SQL FCI must be registered separately.

1. Select **Data Protection > Sources**.
2. Click **Register** > **Databases** > **MS SQL Cluster**.
3. Enter the **Hostname**.

   Cohesity recommends that you specify the fully qualified domain name of the MS SQL FCI to register. Or you can specify the virtual IP address (VIP) of the MS SQL FCI present under SQL Role.

Fully qualified domain name location:

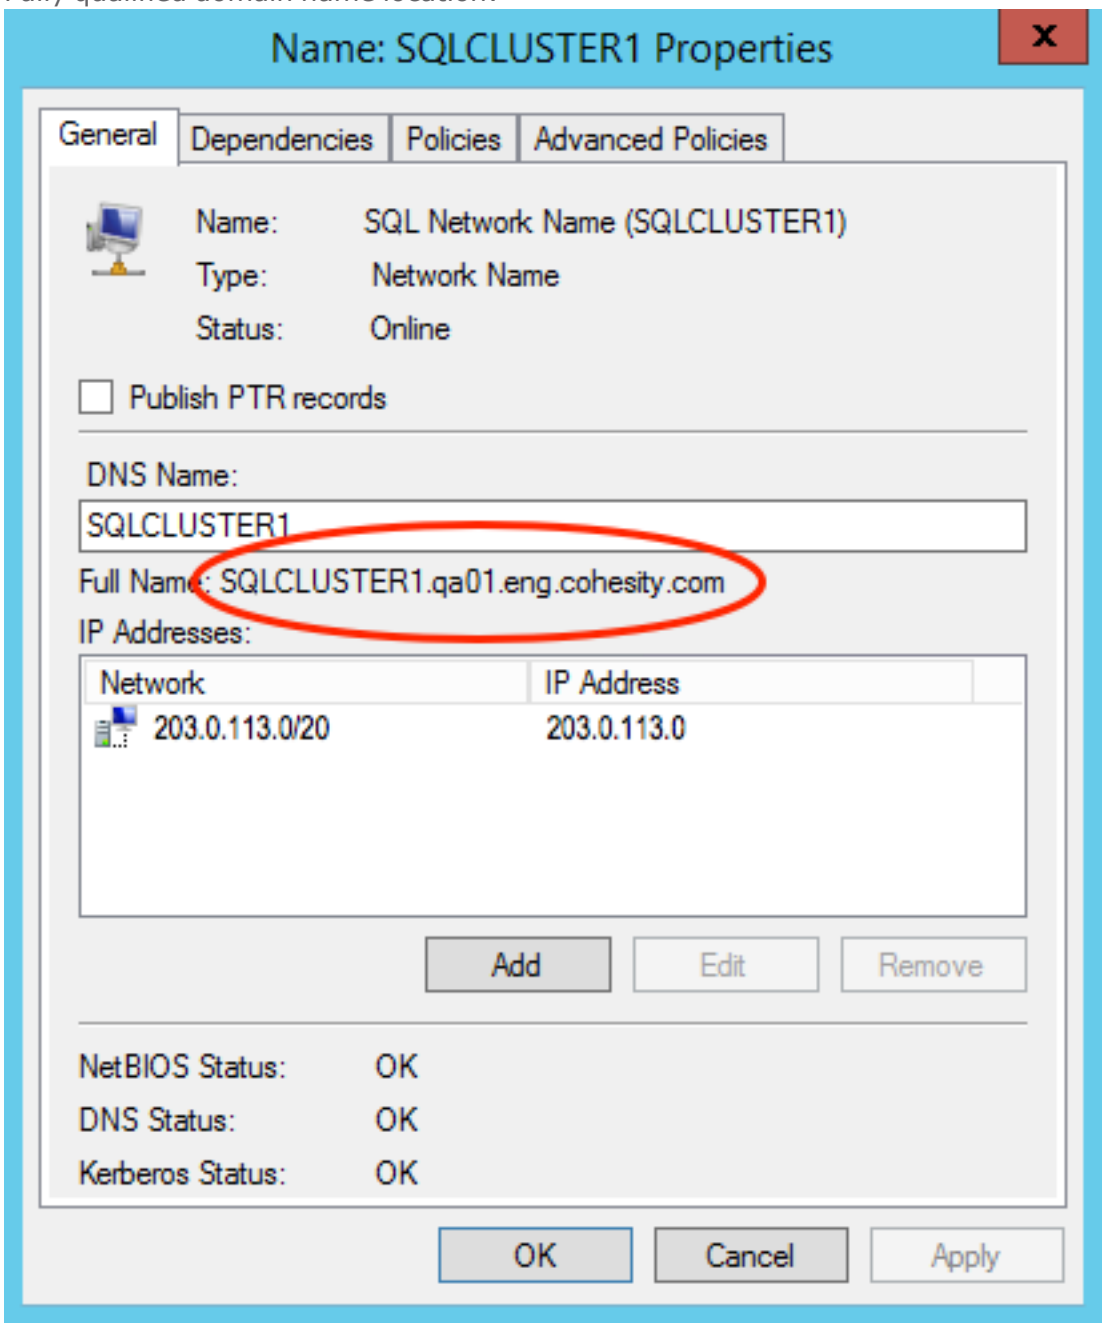

Virtual IP address location:

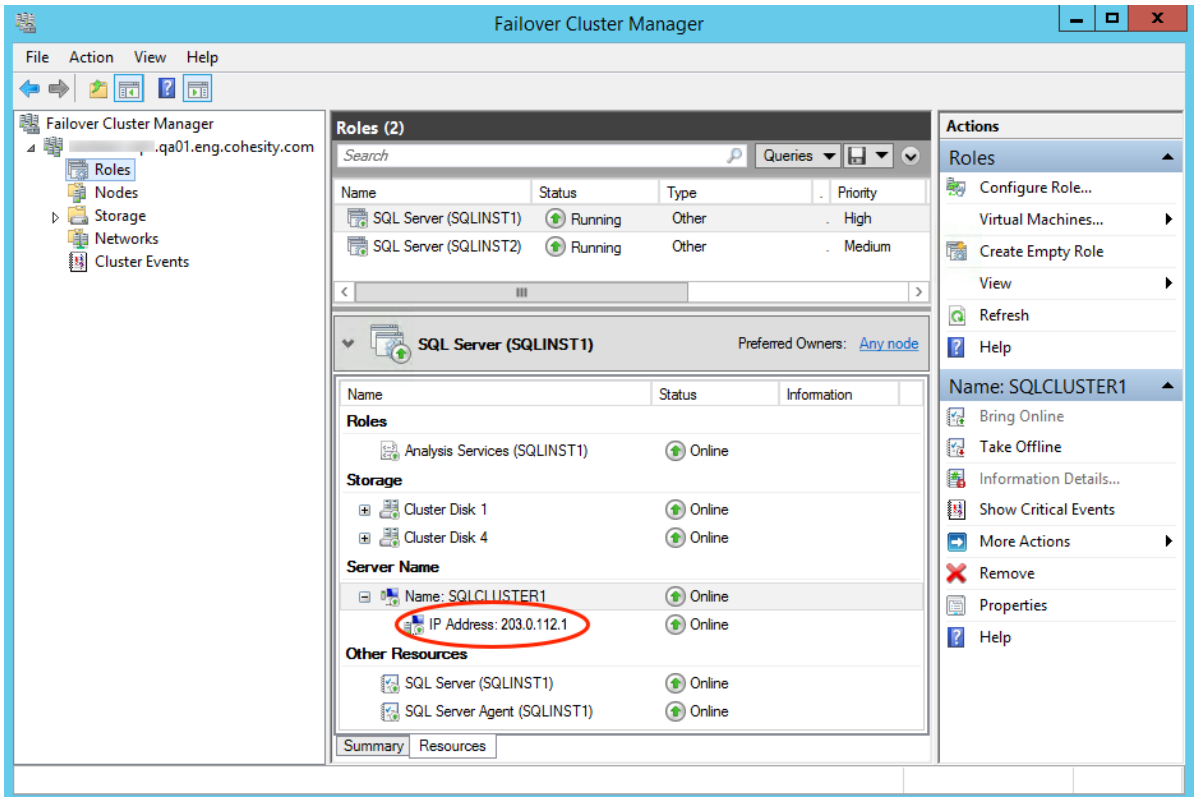

> **Note:** If the MS SQL FCI is not registered using an FQDN and all its nodes are not configured with FQDNs (for example, it is a test environment), changing the IP address of a registered node or adding/removing nodes may not update the FCI's registration configuration information, which could prevent some Cohesity SQL workflows. If you change the configuration information or add/remove nodes under these conditions, Cohesity recommends that you unregister and then reregister the MS SQL FCI. MS SQL FCIs that are registered using an FQDN do not have limitations when updating configuration information such as node IP address changes or adding/removing nodes.

4. Click **Register**.

   When registration finishes, a new entry appears in the physical section on the Sources page.

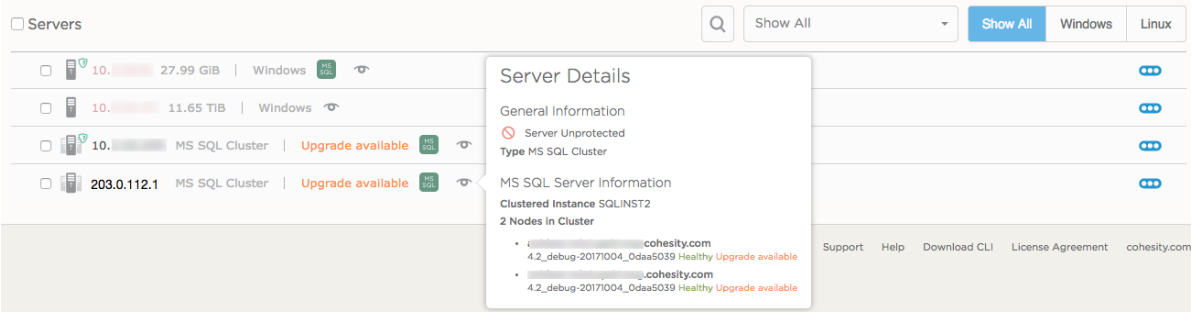

The MS SQL FCI is now registered as a physical server.

Register an MS SQL FCI as MS SQL

This registration step takes an MS SQL FCI source that is already visible (as a physical server) to the Cohesity cluster and registers it as MS SQL Server.

Next to the newly added MS SQL FCI (displayed as a physical server), mouse over the MS SQL FCI's actions menu and click **Register as MS SQL Server**.

When registration finishes successfully, MS SQL Server information displays when you click the icon.

FCI Standalone Multi-Registration

You can register and back up a Windows Server that participates in an MS SQL FCI. You can separately register the FCI clusters and their standalone nodes. This allows you to backup local drives for physical server backups.

Prerequisites

- Ensure that all the SQL agents are upgraded.
- Ensure that the following unsupported protection groups are not present:
  - A protection group that auto-protects a standalone SQL host participating in a WSFC.
  - A protection group that explicitly protects FCI databases on a SQL server participating in an MS SQL FCI.

Error Messages

- The following error message is displayed when unsupported protection groups are present:
  Error
  FCI-node multi-registration cannot proceed, please correct the following protection groups:

- <protectiongroupname1>, <protectiongroupname2> - jobs that autoprotect Windows Cluster node(s).
- <protectiongroupname1>, <protectiongroupname2> - jobs that explicitly back up some FCI databases from Windows Cluster node(s).

Solution

Either delete or fix the protection groups such that they no longer fit the unsupported criteria, and continue the procedure.

- The following error message is displayed when the agent on the Windows cluster node has not been upgraded to a version that supports this feature:
  Error
  Agents on the following sources do not support FCI-node multi-registration: <server1>, <server2>
  Solution
  Upgrade the agent on the Windows cluster node to a version that supports this feature.

Best Practices for Microsoft SQL Server

Cohesity recommends the following best practices when protecting MS SQL:

- Cohesity recommends triggering a full backup whenever a database file (.mdf, .ndf, .ldf) is added or removed from the database.
- Multiple Cohesity clusters concurrently backing up an MS SQL VM is not recommended because the jobs will interfere with the other's log chain.
- For MS SQL point in time restores, the Full Recovery Model is recommended as a best practice. For more information, see Choosing the Recovery Model for a Database.
- Set MS SQL Server to autostart.
- To accurately report the size of transaction log data read/written during log backups, the Windows local or domain Administrator account that is registered with Cohesity for backup must be added as a user to the 'model' system database. Use the following query:

```
-- Configures accurate reporting of the size of transaction log data
read/written during log backup
USE [model]
GO
CREATE USER [HOSTNAME\Administrator] FOR LOGIN [HOSTNAME\Administrator]

GO
```

  Where HOSTNAME is the hostname of the SQL server.
  This assumes that the Administrator account used to register MS SQL is already registered in the SQL server instance's Security Logins and has the "sysadmin" role per documentation. If privileges are not correct, the logical bytes read stat ("Size" column on the Job's details page) for MS SQL log backups will be the same as the

physical bytes read stat ("Read" column on the Job's details page). If privileges are correct, then the size reported will be the sum of the sizes of all the .ldf files for the database.

Considerations for Microsoft SQL Server

General

- For file-based Protection Group, object-level granularity is not supported for Role-based access control (RBAC), such as setting access to objects through Active Directory and roles.
- If using native or third-party backup software to back up an application (such as MS SQL), note that the following behavior breaks the differential chain for the application on that server:
  - Performing an app-consistent backup on a VMware SQL VM.
  - You are allowed to run a file-based physical server Protection Group with a SQL Protection Group or 3rd party SQL protection tools.
    This is a Microsoft limitation. For more information, see Various Causes of Log Chains Breaking Issue in SQL Server.
- The SQL Server VDI log backup and restore operations require Sysadmin privileges which is a Microsoft requirement. For more details, see Microsoft Documentation on SQL Server VDI backup.
- Cohesity cluster uses the Encrypted Network Block Device (NBDSSL) transport method for the vSphere Storage APIs for Data Protection (VADP) workflow. NBDSSL uses SSL to encrypt the data. SAN transport mode operation is not supported.

Protection

- Ensure that all AG databases you want to protect are selected from the respective replica.
- For file-based backups, Cohesity operations rely on default MS SQL file extensions (.mdf, .ndf, .ldf).
- SQL Database and logs residing on SMB/NFS can be protected using the MS SQL Native method.
- SQL Database and logs residing on Clustered Shared Volumes (CSV) can be protected using the MS SQL Native and MS SQL VDI methods.
- CBT information does not persist after a reboot for the file-based backup method.
- As per Microsoft requirements, VDI backups executed on secondary replicas are copy-only full backups. Copy-only backups on secondary replicas do not impact the log

chain. For more information, see Offload supported backups to secondary replicas of an availability group.

- Overwriting Filestream database is not supported; this may prevent the completion of recovery.

Recovery

- Performing a point-in-time restore of an Always On Availability Group replica requires full backups from that replica.

  Log backups are merged across all replicas; full backups are not. During restore, the user sees multiple results for the same database from different replicas. Each replica shows a set of Snapshots for full backups done on that particular replica. For point in time restore to a specific replica, the user must select a full backup Snapshot from that replica only. If a full backup does not exist, point in time restore is not possible.

  Restore the database to the desired MS SQL Server instance on a replica that has a full backup.

- Using the "Restore to Original SQL Server instance" option is not supported if any of the following occurs:
  - The Windows host name changed and you attempt to restore a snapshot that was taken before the host name change.
  - Granular DB restore performed from a higher SQL Server version to a lower version. This is a Microsoft limitation.

  You can toggle off "**Restore to Original SQL Server Instance**" and perform a database restore.

  For more information, see Rename computer hosting instance - SQL Server.

  Restoring a snapshot taken after the name change or upgrade is supported.
- Restoring MS SQL Server from backups of dynamic disks to the same machine is not supported.
- Use File-based or VDI-based backup methods to protect TDE databases. Follow the Microsoft TDE best practices for key and cert management.
- If multiple restore tasks are run on the same target host, the restore tasks will be serialized.
- **For system database recovery**:
  - Recovering the system database to an alternate instance is not supported.
  - When recovering to the original database, the original database is deleted before the recovery. Therefore, a recovery failure will also lead to the loss of the original database.

  - If you are recovering the system databases individually, then the recovery of the system database must be in the following order:
    1. Master
    2. MSDB
    3. Model
  - You cannot recover a system database along with a user database.
  - If you are recovering more than one system database in a recovery task, then the recovery point for the selected system databases must be the same.
- For recovering a database to an Amazon RDS Custom SQL Server, the recovery destination path must be under the **D:\rdsdbdata** folder.

Clone

- Cloning an MS SQL database and renaming the database via SSMS can leave the cloned database in the "Recovery Pending" state during a clone teardown task. This is because the Cohesity cluster database clone operation assumes the cloned database name and files are associated with the original clone operation. To prevent a disconnect between the Cohesity cluster and SQL Server operation, it is recommended to clone a new database with the intended database name rather than rename a cloned database.

Migration

- Cohesity does not support the migration of MS SQL FILESTREAM databases.
- Once a migration task is configured and created, the target SQL Server or database paths cannot be modified for the existing migration task.
- Cohesity does not detect changes/deletions to the target database with a running migration task. Changes made external to the migration task before finalizing can cause the migration task to fail.
- Cohesity does not support the migration of MS SQL databases that are replicated to an alternate Cohesity cluster.

General Recommendations for Microsoft SQL Server Backup Methods

Filestream databases can be protected using the VDI-based approach.

Normal databases can be protected using the following approaches, listed in the order of recommendation:

- File-based
- VDI-based

The following table can help you select the appropriate backup type:

	File-based Backup	VDI-based Backup
Microsoft Approach	Leverages Microsoft VSS approach	Leverages Microsoft SQL Server native backup and restore commands through API calls
Built-in Change Block Tracking (CBT)	Leverages Cohesity's built-in CBT technology	Leverages SQL server's differential backups
Database clone from backup	Supported	Not Supported
Recovery: Cross-domain AD without trust	Supported (GRPC is used to transfer data)	Supported (GRPC is used to transfer data)
Database clone: Cross-domain AD without trust	Not Supported (Clone uses SMB and cross-domain trust is required for SMB access)	-
Implicit Full backup after SQL host restart	The first backup after the SQL server reboot will be a full backup (incremental backup automatically switches to full backup)	-
Memory Optimized Filestream Database	Not Supported	Supported
SQL server-side compression	Not Supported	Supported but not recommended to use, since it may cause low deduplication rate on the Cohesity clusters, which leads to high storage utilization on the Cohesity cluster.
Parallel Log Backup to Full/ Incremental Backup	Supported	Supported
Fast Backup (multi-node)	Supported	Supported
Fast Restore (multi-node)	Supported	Supported
Database-level recovery from archived backup	Supported (Beta till 7.0)	Supported (Beta till 7.0)

	File-based Backup	VDI-based Backup
CAD Policy	Supported	Supported
Transparent Data Encryption (TDE) Databases	Supported	Supported
Database on ReFS volumes	Supported	Supported
Quiesce Support	SQL Database goes through quiesce (freeze/thaw)	Not required. The SQL server performs consistent backup
Free Space on SQL Host (VSS shadow storage requirement)	At least 10% is recommended for each volume on the SQL host by MSFT. Large databases with high IO rates may require additional space.	-
Must-Have	-	Must have Full and incremental (differential). NOTE: Setting extended retention is required for full backup to avoid expiry before the incremental backup expiry.
Parallel run-now backups (same type) for the same job	When the host is already part of a previous run-now submitted, the backup is queued until the previous submitted run is complete.	Backup of the selected database runs in parallel with the existing run-now request.
Known Issues that lead to backup failure	Very High transaction rate on databases may occur resulting in backup failure. This is due to Microsoft's VSS system provider not being able to maintain the VSS snapshot on SQL host with very high transactions. This is very rare. Configuring the shadow storage on volume with no/less IO helps to mitigate the issue.	Agent/SQL host may restart during backup.
Known Issues that lead to incremental backup converting to full backup	First backup after SQL server reboot will be a full backup (incremental backup automatically switches to full backup)	If no explicit policy is available for full backup, incremental backup is converted to full backup after 7 days.

	File-based Backup	VDI-based Backup
Known issues that lead to PIT clone failure	-	-
FortKnox	Supported	Supported

> **Note:** REFS is supported by VDI-based and File-based SQL server.

Backup Amazon RDS Custom for SQL Server

The following topics provide guidance to help you protect MS SQL databases:

For more information on MS SQL deployments, See Protect SQL Server with Cohesity—Deployment Considerations Guide.

About Microsoft SQL Server Backup

Cohesity supports the following MS SQL Server backups:

- File-based Backup
- VDI-based Backup

Cohesity also supports VM-based Backup and SQL Native Backup.

File-based Backup

A Cohesity cluster MS SQL file-based backup protects only the MS SQL databases you choose. It captures only the database files for those selected databases. The Cohesity cluster captures full database server backups and can optionally back up database transaction logs, so you can roll forward to any point in time.

The advantage of a file-based backup is that a smaller amount of data is captured, which allows the Protection Group to run more quickly.

> **Note:** To enable file-based backup, you must install the File System CBT component during Cohesity agent installation.

For information on performing MS SQL file-based backup, see Microsoft SQL Server (File-based).

All configurations are for illustrative purposes only. Cohesity recommends following MS SQL best practices.

An example of MS SQL running on a Windows server that contains the following:

- Multiple volumes, C, D, E and F in the figure below
- Volume D is running two MS SQL databases
- Volume E is running three MS SQL databases
- Volume F contains database logs

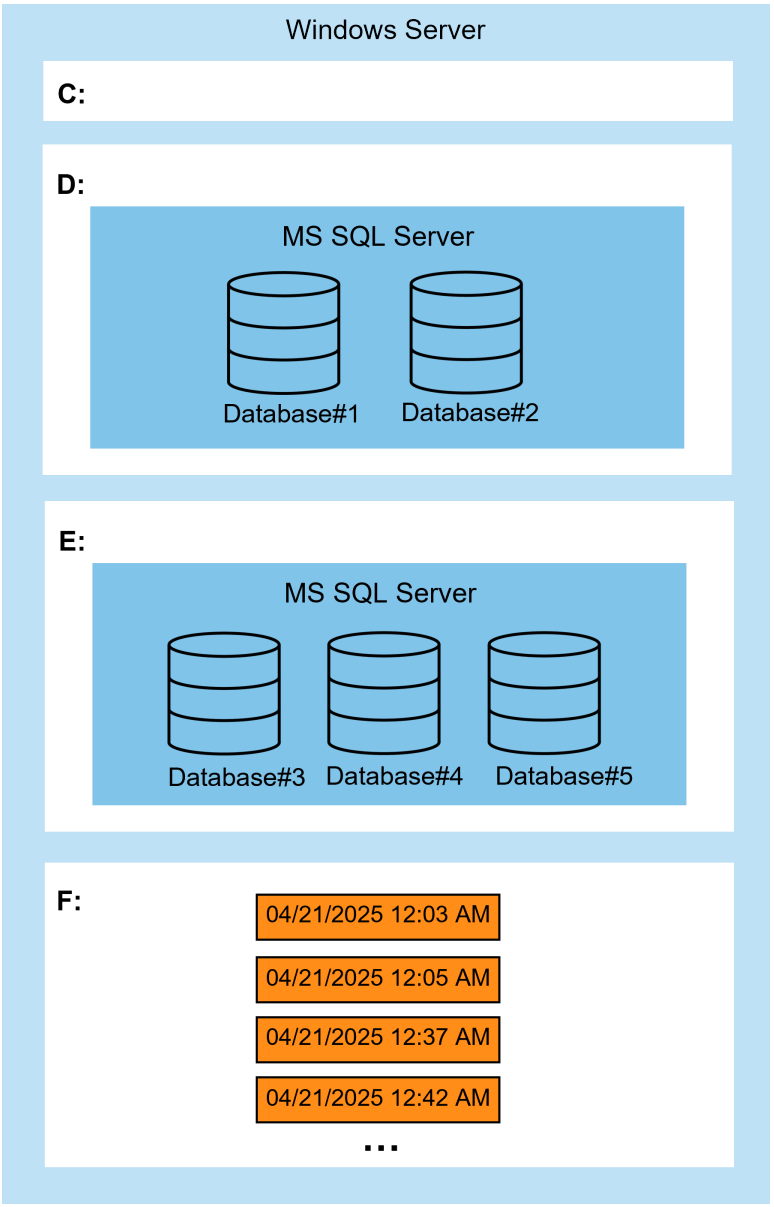

The MS SQL file-based backup task consists of the following actions:

1. During SQL database protection run, a VSS snapshot of the specified MS SQL databases is captured using Volume Shadow Copy Service (VSS). The VSS snapshot resides on the local drive as a special VSS volume. This VSS snapshot is a full database server backup.

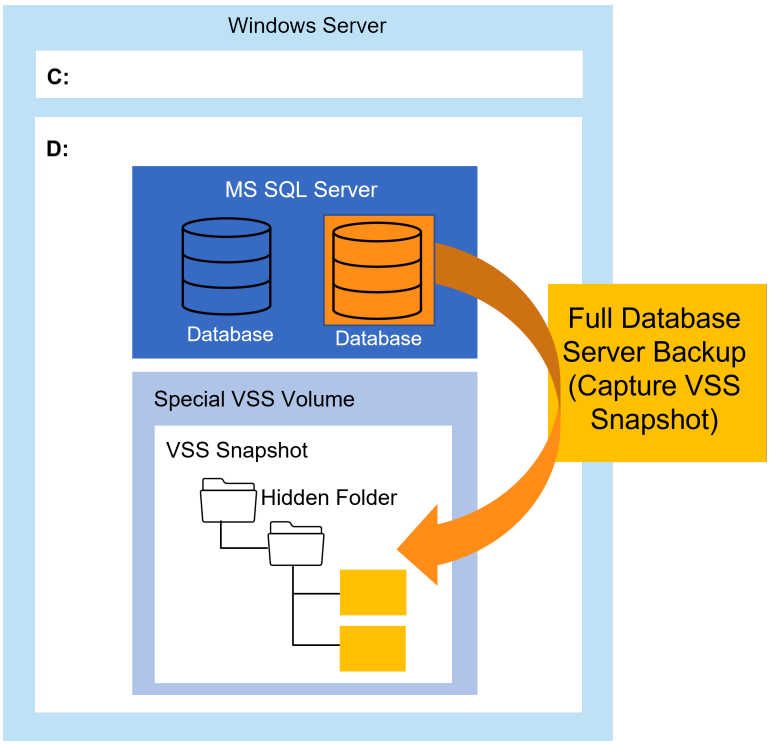

2. The snapshot is captured and saved on the Cohesity cluster with a date and time stamp. You can optionally choose for the Cohesity cluster to utilize Change Block Tracking (CBT) when backing up the snapshot. VMware VMs leverage VMware CBT to back up the snapshot. Physical servers leverage Cohesity CBT to back up the snapshot.

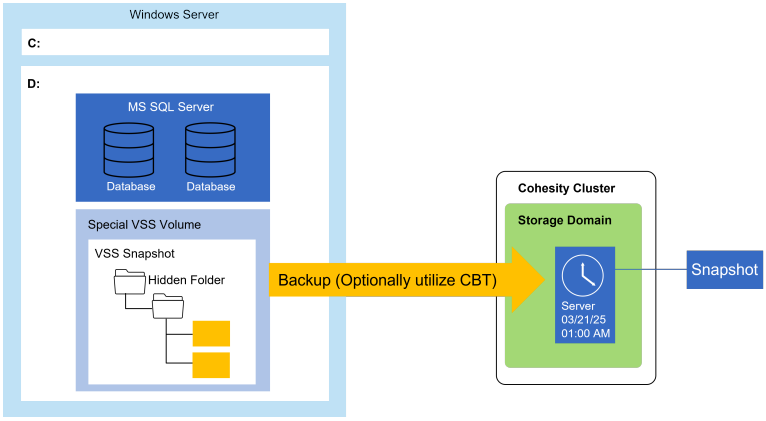

3. Backups of database transaction logs use MS SQL Server Virtual Device Interface (VDI) and are stored on the Cohesity cluster. The logs allow you to roll back to a specific point in time. MS SQL Server automatically truncates the SQL logs after a successful SQL log backup and marks the Virtual Log Files (VLFs) for reuse.

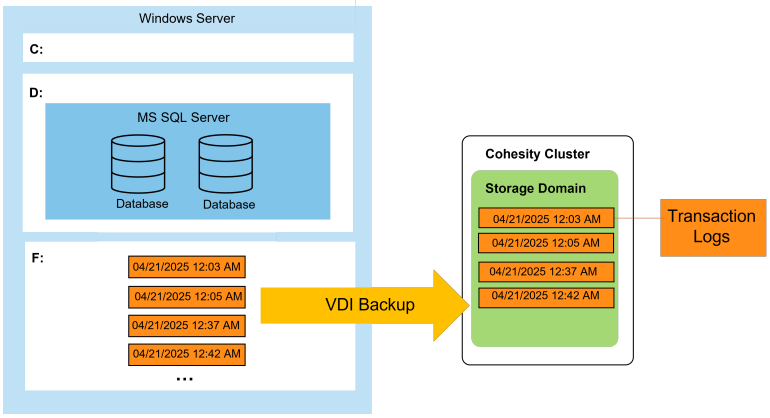

VDI-based Backup

A Cohesity cluster MS SQL VDI-based backup captures the native MS SQL server backups. It captures VDI-based backups of a specific database on a SQL Server instance running on a VMware or physical server.

The Virtual Device Interface (VDI) is a Microsoft interface that allows the Cohesity agent to execute SQL Server Native backup and restore commands. This backup type produces a standard SQL Server Native backup of type (.BAK), transaction logs, (.TRN), and differentials, (.DIFF). These backup files are stored on a Cohesity View and take advantage of View compression, deduplication, encryption, replication, and archiving features.

All configurations are for illustrative purposes only. Cohesity recommends following MS SQL best practices.

An example of MS SQL running on a Windows server that contains the following:

- Multiple volumes, C, D and F in the figure below
- Volume D is running multiple MS SQL databases
- Volume F contains database logs

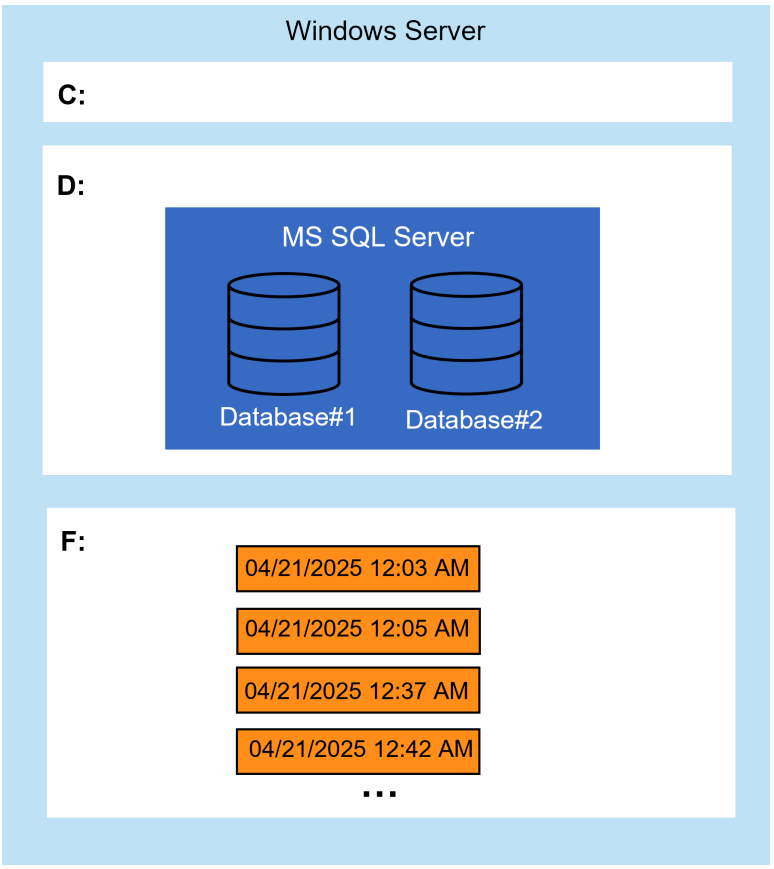

The MS SQL VDI-based backup task consists of the following actions:

1. During an MS SQL protection run, a "BACKUP DATABASE" command for the specified MS SQL databases is initiated using MS SQL Server Virtual Device Interface (VDI). The generated backup file is saved on the Cohesity cluster with a date and time stamp.

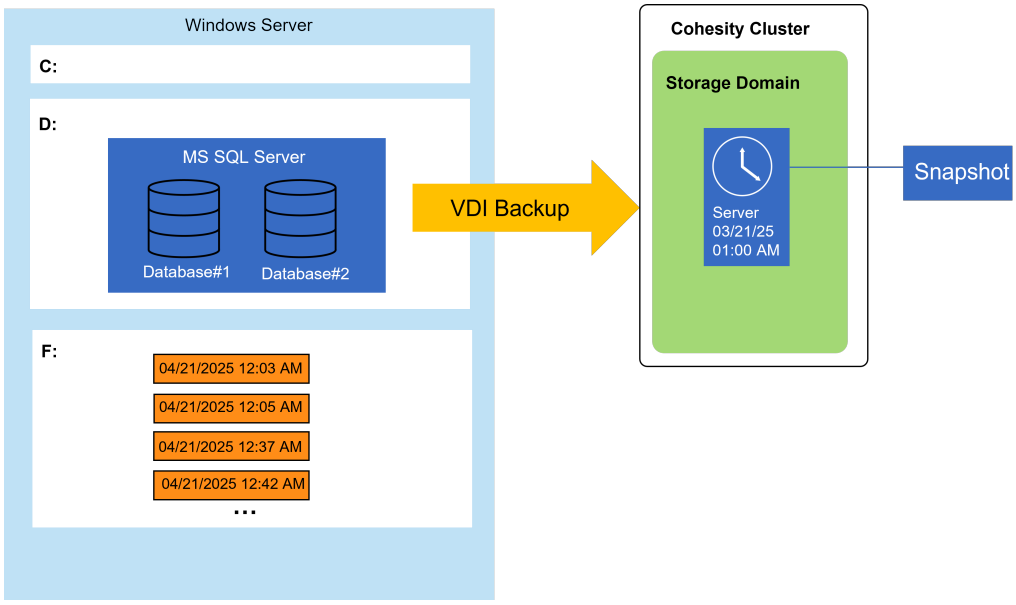

2. If the policy is configured to take incremental database backups, it translates to differential backups in case of VDI-based backup.
3. If transaction log backups are specified in the policy, the database transaction logs are backed up to the Cohesity cluster using MS SQL Server Virtual Device Interface (VDI). The logs allow you to roll back to a specific point in time. MS SQL Server automatically truncates the SQL logs after a successful SQL log backup and marks the Virtual Log Files (VLFs) for reuse.

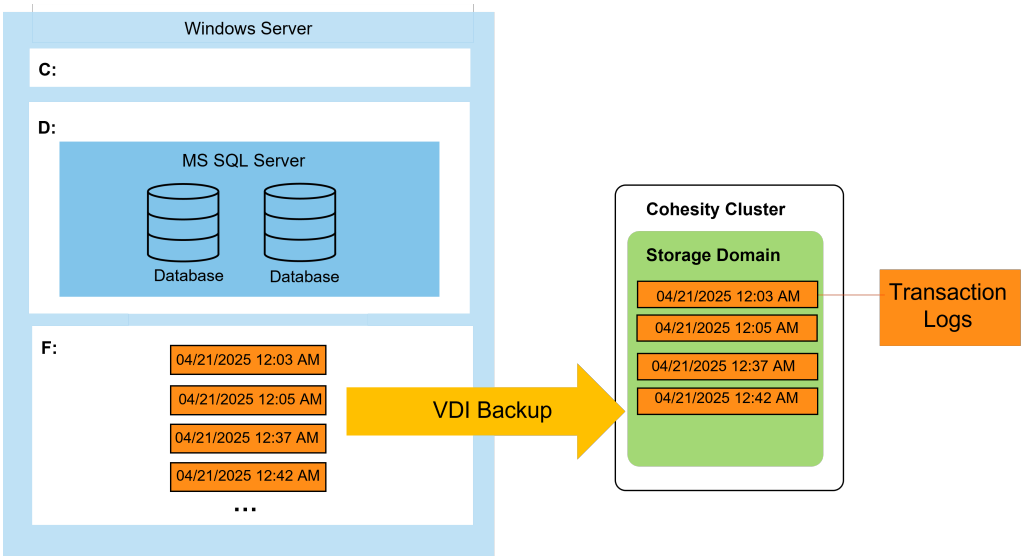

Compare File-Based and VDI-Based Backup Methods

The following table compares the implications of each method.

Cohesity Facet	File-based backup	VDI-based backup
Technology Employed	Cohesity Agent uses VSS and Cohesity technology for file CBT, and VDI for T-logs.	Cohesity agent uses SQL Native backup commands.
What does it capture?	Captures database files.	Captures the database in a SQL Server standard .BAK, .TRN, or .Diff file.
Logs	Yes. Captures a SQL log backup of the database.	Yes. Captures a SQL Server native log backup of the database.
Cohesity Source Registration	Register as a SQL Server application.	Register as a SQL Server application.

Cohesity Facet	File-based backup	VDI-based backup
Compatibility with SQL Server High Availability Features	**Windows Failover Cluster**: **Yes**, backup will succeed if a failover occurs. Must be registered as a Cohesity "SQL Cluster" source. **AG**: **Yes**, backup on the correct replica will succeed after an AG failover occurs. **Mirroring**: **Yes** Each SQL instance must have its own Protection Group. **Log Shipping**: **Yes**. Each SQL instance must have its own Protection Group. **Note:** While Cohesity currently does not include specific features for mirroring and log shipping, it does not conflict with these SQL Server features. **Note:** Log shipping via SQL Server Management Studio supports only full/incremental backup and not log backup since it can break the log chain.	**Windows Failover Cluster**: **Yes**, backup will succeed if a failover occurs. Must be registered as a Cohesity "SQL Cluster" source. **AG**: **Yes**, backup on the correct replica will succeed after an AG failover occurs. Mirroring: No. Log Shipping: No
Method	This is an MS SQL Server database backup at the **file level**. Allows for FULL SQL Server backups with logs, but only the database files are captured. The Cohesity Agent uses VSS, the Cohesity File CBT, and VDI to take an application-consistent snapshot of the volumes on that server. Faster because it's not backing up extraneous files.	This is a native MS SQL Server database backup at the **file level**. Allows for FULL SQL Server backups with logs and Differentials, but only the database is captured. Agent uses SQL Server native backup and restore commands , to take an application-consistent backup of the database.
Recovery/Restore	The database host server should be backed up and restored via a different protection group. Supports Point-in-Time (PIT) restore of the database from across the entire backup history.	

Cohesity Facet	File-based backup	VDI-based backup
Cloning	**Yes**. Can produce a representative database. *Advantage*: Applications can be tested with data that is as close to the production environment as possible: Can reproduce a database. *Advantage*: Applications can be tested with databases that can be destroyed and re-created on demand. No load on Production servers. The backup is NOT modified and therefore always safe. A successful clone is immediate validation of the backup's integrity. Uses storage resources on Cohesity cluster.	**Cloning is scheduled as a future enhancement.**
Source-side deduplication	Supported	Not supported
Archival	**CloudArchive**: Cohesity supports the ability to send backups to public cloud providers like AWS, Azure, and Google Cloud Platform. It also supports search and restore from them, back to the local site when needed.	
Replication	**Managing Copies of backups on another Cohesity cluster**: Backups, files, and volumes can be replicated to a second Cohesity cluster.	

About Log and Full DB Backup Retention Periods

The retention periods of log files and full database server backups are independent of each other. Consider the following examples to help determine the best retention periods for your database server backups. These retention guidelines also apply to full database server and log replication schedules. Retention periods of log files and full database backups must adhere to Microsoft rules for restoring databases.

1. A point-in-time restore using SQL log backups is dependent on a full backup because a full database server backup must also be available to apply the logs.

   When the Cohesity cluster restores to a specific point in time, it must first restore a full database server backup and then apply the appropriate set of logs. If a full database server backup is not available because the retention period of the full database server backup is too short, the Cohesity cluster cannot restore to a point in time within the specified log retention period. For example if the log retention period

is three weeks and the full database server backup captures snapshots weekly with a retention period of two weeks, the Cohesity cluster will not be able to restore to any point in time from three weeks ago. In addition, the logs from three weeks ago cannot be utilized and are using disk space.

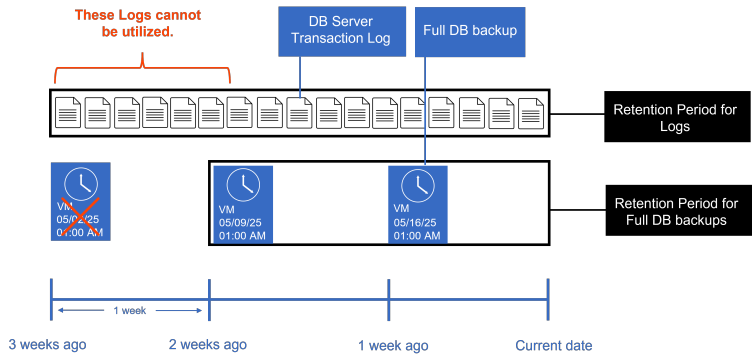

2. If the log retention period is shorter than the full database server backup retention period, the Cohesity cluster cannot restore to a point in time where the logs are not saved. For example, if the log retention period is two weeks and the full database server backup captures snapshots weekly with a retention period of three weeks, the Cohesity cluster will not be able to restore to any point in time between two and three weeks ago. The Cohesity cluster can, however, restore to the snapshot from three weeks ago.

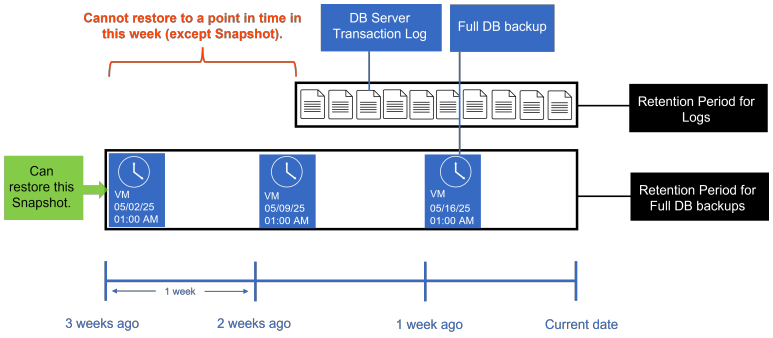

3. Cohesity recommends specifying a log retention period with extra time to allow for failed full database server backups as demonstrated by the following example. A full database server backup is scheduled weekly with a three-week retention period for both full database server and logs. If the full database server backup failed two weeks ago, the Cohesity cluster can still restore to the point in time from the log indicated by the green arrow.

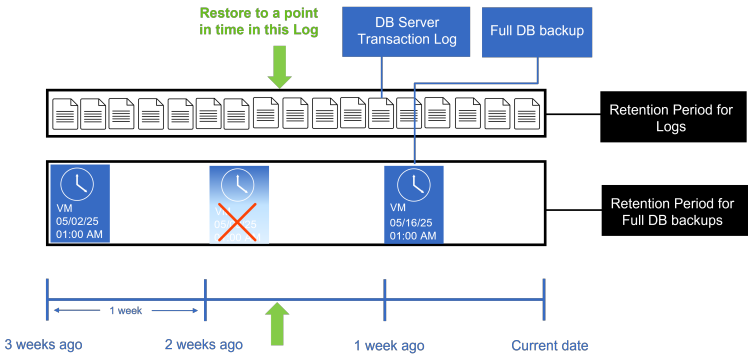

The Cohesity cluster restores full database server backup from three weeks ago and applies the appropriate logs to restore the state to a point.

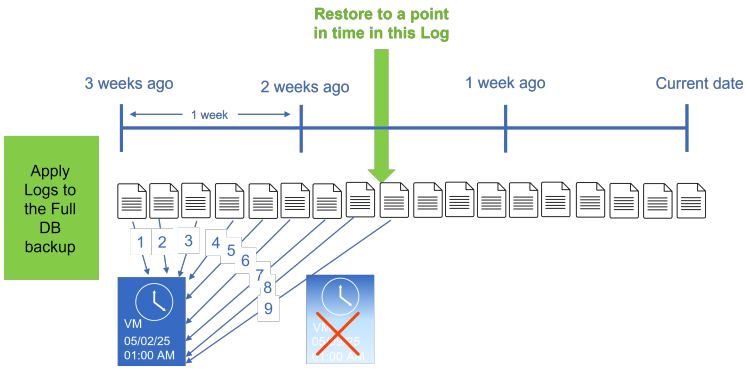

You can restore to a specific transaction in the log.

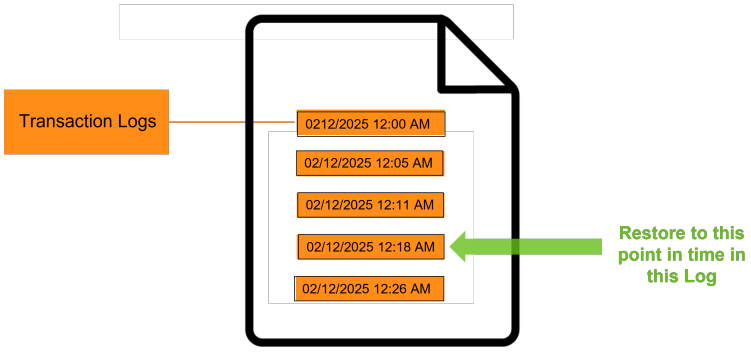

> **Note:** A full database backup is not a transaction log backup. Ensure that your SQL Server protection policy has transaction log backup for full and bulk-logged recovery model requirements.

Microsoft SQL Server (File-based)

This topic describes how to protect MS SQL with a file-based Protection Group. File-based Protection Groups protect only the specific MS SQL databases that you select.

MS SQL Protection Group Support

The following table describes supported job types for MS SQL running on physical servers and VMs.

	File-based Jobs	Volume-based Jobs	VDI-based Jobs	Both Job Types Using the Same Server but Protecting Distinct Databases
Physical MS SQL	Multiple file-based jobs protecting all or a subset of DBs	One volume-based job protecting all or a subset of DBs	Multiple VDI-based jobs protecting all or a subset of databases	Supported
SQL Server Virtual Machine Backup	NA	Volume-base with Hypervisor to provide application-consistent backup	NA	NA

Protect MS SQL Overview

Perform the following steps in the Cohesity cluster to protect SQL Server:

1. Register the SQL Server you want to protect as a source with the Cohesity cluster. See Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs or Set Up Microsoft SQL Server FCI.
2. Select **Dashboard** and from the drop-down, select **MS SQL**.
3. From the **Hosts** tab, you can protect all databases on a host. Select **Protect Databases** from the host's actions menu.

   From the **Databases** tab, you can protect a single database. Select **Protect** from the database's actions menu.

   You can change which objects you want to protect when creating the job.

You can also protect MS SQL by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. Select the objects to protect.

Create a File-based MS SQL Protection Group

Perform the following to create a file-based SQL Server Protection Group:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. In the **New Protection** pop-up, do one of the following:
   - Protect MS SQL by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered MS SQL source or click Register Source to register an MS SQL source. Select the objects to protect.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. **Backup Type**: From the drop-down, select **File-based**.
     vi. Click **Protect**.
   - To protect MS SQL by specifying all the options, click **More Options**.
4. **Source:** Select the objects to protect.

   Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.

   > **Note:** If you have registered a physical server that is a part of the registered MS SQL Failover Cluster, then the MS SQL Failover Cluster Instance will not appear as a source in the object list of the registered physical server.

   If appropriate for your organization, enable Auto Protect.

   Click to turn on **Auto Protect**. By enabling this option:
   - All the SQL databases that are added to the SQL server in the future are automatically added to the Protection Group and are protected from the next

protection run.

- All the SQL databases that are removed from the SQL server in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the user is preserved until the snapshot expires.

For example, you create a Protection Group that backs up a SQL Server. The Protection Group runs for two weeks and then a new database is added to the SQL Server. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new database is also backed up even though the new database has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours.

To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click .

Considerations

Consider the following when using the **Auto Protect** option on the backup job

- If new databases are added to the host, the existing filters are applied to the newly added databases. For example, if the exclude filter matches the newly added databases then these databases will be excluded in the backup job.
- If any system databases are excluded by the filters in an SQL server, it will exclude all the system databases belonging to that SQL server.
- When defining filters to exclude AG databases, ensure that the filters are configured on all the AG hosts. For example, to exclude an AG database and if the AG database has 4 hosts, then ensure that the filters are defined on all the four AG hosts.
- Auto protect functionality is supported only for physical SQL jobs (file-based and native)
- To selectively exclude a SQL database from the protection job, you must upgrade all the cluster nodes to version 6.5.1 and later.

If MS SQL belongs to an AG network, selecting that server displays options for protecting the AG.

- Select all servers that belong to this AG network.

> **Note:** If backing up VMs within an AG, protecting the AG databases requires that all servers within the AG are registered with the Cohesity cluster and are selected for protection. Otherwise AG databases are skipped and only local databases will be backed up. New servers added to AG are not automatically protected.

When using this option when there are multiple AGs, only interlinked AG nodes are automatically selected. For example:

- If there are AGs that use only Node1 and Node2 and other AGs that use Node3 and Node4, if you select Node2 as a source, only Node1 is added automatically.
- If AG1 uses Node1 and Node2, AG2 uses Node2 and Node4, AG3 uses Node3 and Node5. If you select Node2, then these nodes are selected: Node1, Node2 and Node4. Node3 and Node5 are not selected.

When adding nodes to an AG that is protected by Cohesity, ensure you register the node with Cohesity and add it to the Protection Group. Otherwise the AG databases will not be protected.

When removing nodes from an AG that is protected by Cohesity, you can remove it from the Protection Group or keep it; the AG databases remain protected in either case. If the removed node remains in the Protection Group, only its local databases are backed up. If you do not want to back up the removed node, you can safely remove it.

- Select only this server.

Click **Select**.

Click **Add**.

5. If appropriate, set options for specific objects by clicking (**Edit Server Options**) next to the object. A successful SQL log backup automatically truncates the SQL logs.
6. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

> **Note:** This page displays registered MS SQL objects only, so only MS SQL servers are selected regardless of the other servers within the source object hierarchy.

7. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.
8. **Storage Domain:** Select an existing storage domain or create a new one by selecting **New Storage Domain**.

> **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

The selected storage domain stores the snapshots and logs created by this Protection Group.

9. **Backup Method:** Select **File-based**. This method protects only the specific MS SQL databases that you select.
10. **Make Full Backup Copy-only:** Toggle on if you want full backups to be copy-only backups so they do not affect the differential base.

    Copy-only full backups do not take log backups even if they are scheduled by the policy.
11. **Log Chain Break Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the log chain is broken.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.
12. **New Database Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the new database is detected.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.
13. **User Databases:** Does not apply to file-based Protection Groups. Only selected databases are backed up.
14. **AG Backup Preferences:** Select this option if AG databases will be backed up.
    - **Use Server Preferences** uses MS SQL preferences.
    - **Override Preferences** enables you to override MS SQL's preferences with your selection.

    AAG replica preference for MS SQL backups

    Cohesity uses "replica priority" to select the best replica when more than one qualified replica matches the backup preference. For the following AAG Backup Preference settings, Cohesity uses replicas in the indicated order of preference to back up MS SQL databases:

Backup Preference Setting	Replica Used for MS SQL Backup
Prefer Secondary or Any	1. Sync Secondary Replica 2. Async Secondary Replica 3. Primary Replica
Secondary Only*	1. Sync Secondary Replica 2. Async Secondary Replica
Primary Only	1. Primary Replica

* If the AAG uses the **Secondary Only** Backup Preference setting, ensure the AAG replicas are set to "**Readable secondary**=**Yes**". If all replicas are set to "**Readable secondary**=**No**" with **Secondary Only** as the Backup Preference setting, then the backup of AAG databases fails with "*unavailable_for_vss*".

You can set the **Readable Secondary** field to **No** if the Backup Preference Setting is set to **Prefer Secondary or Any** or **Primary Only**.

> **Note:** You must apply the exclusion filter on each AAG replica to ensure the database is excluded when an SQL Server failover occurs.

15. **System Databases:** Does not apply to file-based Protection Groups. Only selected databases are backed up. When you select any system database, however, all system databases are selected automatically.
16. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
17. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
18. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.

19. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
20. **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

    For configuration details, see Configure Pre & Post Scripts.
21. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.

    Indexing is not supported for file-based Protection Groups. Without indexing, file/folder-level recovery and searching are not supported.
22. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.
23. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
24. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

> **Note:**
> The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

25. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
26. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Edit an MS SQL Protection Group

Perform the following to edit a SQL Server Protection Group:

1. Select **Data Protection > Protection**.
2. Locate the Protection Group you want to edit.
3. Mouse over the action icon (⋮) next to the Protection Group and select one of the following:
    1. **Run Now** - Start a run of this Protection Group. You have the option to back up one or more objects.
        1. **Backup all Objects in the Protection Group** - Select this option to backup all the databases in the protection group. Any databases excluded from the backup will not be included in the protection run.
        2. **Backup only Selected Objects in the Protection Group** - you can manually select the databases to be backed up at the Protection Group level.
           If the host (parent of the selected database) is already part of a previous run-now submitted, the backup will be queued until the previous run submitted is complete.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

2. **Pause Future Runs**
    1. Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the Protection Group is currently executing a run, the current run continues to execute and only future runs of this Protection Group are paused.
    2. Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this Protection Group.
3. **Edit** - Edit the Protection Group settings of the protection group and click Protect.
4. **Delete** - Deletes the protection group.

4. Save the Protection Group settings by clicking **Protect**.

Backup Microsoft SQL Server (VDI-based)

This topic describes how to protect MS SQL with a VDI-based Protection Group.

MS SQL Protection Group Support

The following table describes supported job types for MS SQL running on physical servers and VMs.

	File-based Jobs	Volume-based Jobs	VDI-based Jobs	Both Job Types Using the Same Server but Protecting Distinct Databases
Physical MS SQL	Multiple file-based jobs protecting all or a subset of DBs	One volume-based job protecting all or a subset of DBs	Multiple VDI-based jobs protecting all or a subset of databases	Supported
SQL Server Virtual Machine Backup	NA	Volume-base with Hypervisor to provide application-consistent backup	NA	NA

Protect MS SQL Overview

Perform the following steps in the Cohesity cluster to protect SQL Server:

1. Register the SQL Server you want to protect as a source with the Cohesity cluster. See Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs or Set Up Microsoft SQL Server FCI.
2. Select **Dashboard** and from the drop-down, select **MS SQL**.
3. From the **Hosts** tab, you can protect all databases on a host. Select **Protect Databases** from the host's actions menu.

   From the **Databases** tab, you can protect a single database. Select **Protect** from the database's actions menu.

   You can change which objects you want to protect when creating the job.

You can also protect MS SQL by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. Select the objects to protect.

Create a VDI-based MS SQL Protection Group

Perform the following to create a VDI-based SQL Server Protection Group:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. In the **New Protection** pop-up, do one of the following:
   - Protect MS SQL by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered MS SQL source or click Register Source to register an MS SQL source. Select the objects to protect.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. **Backup Type**: From the drop-down, select **VDI-based**.

vi. Click **Protect**.

- To protect MS SQL by specifying all the options, click **More Options**.

4. **Source:** Select the objects to protect.

   Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.

   > **Note:** If you have registered a physical server that is a part of the registered MS SQL Failover Cluster, then the MS SQL Failover Cluster Instance will not appear as a source in the object list of the registered physical server.

   If appropriate for your organization, enable Auto Protect.

   Click to turn on **Auto Protect**. By enabling this option:

   - All the SQL databases that are added to the SQL server in the future are automatically added to the Protection Group and are protected from the next protection run.
   - All the SQL databases that are removed from the SQL server in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the user is preserved until the snapshot expires.

   For example, you create a Protection Group that backs up a SQL Server. The Protection Group runs for two weeks and then a new database is added to the SQL Server. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new database is also backed up even though the new database has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours.

   To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click . To exclude an object from Auto Protect, click .

   - Indicates an Auto Protected object.
   - Use this option to exclude SQL databases from the protection group. Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

To exclude a database from the protection group, specify the absolute path of the database to be excluded from the backup at a protection job level. You can optionally use:

- **Wildcard Exclusion** - You can optionally use the Wildcard * in any location in the search string to exclude all the matching databases from the backup at a protection job level.
- **Regex Exclusion** - You can use the RegEx string name match and manual user selection to exclude DBs from backup at a protection job level.

> **Note:** You can use the exclude SQL database feature only if Auto Protect is enabled for a backup job.

Protection of dismounted databases, recovery database types, and deleted databases are automatically skipped by the Protection Group.

After defining the filter click **Apply** to save the selection.

The following table lists the examples for excluding objects in an MS SQL database:

Sample Configuration		
10.2.157.20 (Host) *** SQL2012** (SQL Server Instance) - UserDB1 - UserDB2 - TestDB1 - TestDB2	*** SQL2014** (SQL Server Instance) - UserDB3 - UserDB4 - TestDB3 - TestDB4	*** MSSQLSERVER** (SQL Server Instance) - ProdDB1 - ProdDB2
Example of Absolute Path		

Sample Configuration
• **10.2.157.20/SQL2012/UserDB1** - Excludes the database 'UserDB1' from the SQL server instance 'SQL2012' on the host 10.2.157.20. • **10.2.157.20/SQL2012/** - Excludes the 'SQL2012' from the SQL server instance and its databases. **Note:** The absolute instance path must end with '/' and the plain text instance path must contain two or more slashes.
Example of Wildcard Exclusion
• **10.2.157.20/SQL2012/User*** - Excludes the databases 'UserDB1' and 'UserDB2' from the SQL server instance 'SQL2012' on the host 10.2.157.20. • **10.2.157.20/SQL201?/*** - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.
Example of Regex Exclusion
10.2.157.20/SQL.*/.* - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.

Considerations

Consider the following when using the **Auto Protect** option on the backup job

- If new databases are added to the host, the existing filters are applied to the newly added databases. For example, if the exclude filter matches the newly added databases then these databases will be excluded in the backup job.
- If any system databases are excluded by the filters in an SQL server, it will exclude all the system databases belonging to that SQL server.
- When defining filters to exclude AG databases, ensure that the filters are configured on all the AG hosts. For example, to exclude an AG database and if the AG database has 4 hosts, then ensure that the filters are defined on all the four AG hosts.
- Auto protect functionality is supported only for physical SQL jobs (file-based and native)
- To selectively exclude a SQL database from the protection job, you must upgrade all the cluster nodes to version 6.5.1 and later.

If MS SQL belongs to an AG network, selecting that server displays options for protecting the AG.

- Select all servers that belong to this AG network.

  > **Note:** If backing up VMs within an AG, protecting the AG databases requires that all servers within the AG are registered with the Cohesity cluster and are selected for protection. Otherwise AG databases are skipped and only local databases will be backed up. New servers added to AG are not automatically protected.

  When using this option when there are multiple AGs, only interlinked AG nodes are automatically selected. For example:

  - If there are AGs that use only Node1 and Node2 and other AGs that use Node3 and Node4, if you select Node2 as a source, only Node1 is added automatically.
  - If AG1 uses Node1 and Node2, AG2 uses Node2 and Node4, AG3 uses Node3 and Node5. If you select Node2, then these nodes are selected: Node1, Node2 and Node4. Node3 and Node5 are not selected.

  When adding nodes to an AG that is protected by Cohesity, ensure you register the node with Cohesity and add it to the Protection Group. Otherwise the AG databases will not be protected.

  When removing nodes from an AG that is protected by Cohesity, you can remove it from the Protection Group or keep it; the AG databases remain protected in either case. If the removed node remains in the Protection Group, only its local databases are backed up. If you do not want to back up the removed node, you can safely remove it.

- Select only this server.

Click **Select**.

Click **Add**.

5. If appropriate, set options for specific objects by clicking (**Edit Server Options**) next to the object. A successful SQL log backup automatically truncates the SQL logs.
6. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

   > **Note:** This page displays registered MS SQL objects only, so only MS SQL servers are selected regardless of the other servers within the source object hierarchy.

7. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**. Cohesity recommends you schedule periodic full backups for MS SQL VDI-based protection. For more information, see Create or Edit a Standard Policy.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

   The selected storage domain stores the snapshots and logs created by this Protection Group.
8. **Backup Method:** Select **VDI-based**. This method protects only the specific MS SQL databases that you select.
9. **WITH Clause:** Define the With clause that you want to use to customize the backup. For more information, see BACKUP (Transact-SQL).
10. **Number of Streams:** Define the number of Cohesity nodes that will be used for the backup process. Using multiple nodes provides more throughput for the backup process.
11. **Log Chain Break Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the log chain is broken.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.
12. **New Database Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the new database is detected.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.
13. **Make Full Backup Copy-only:** Toggle on if you want full backups to be copy-only backups so they do not affect the differential base.

    Copy-only full backups do not take log backups even if they are scheduled by the policy.
14. **User Databases:** Select how to handle AG databases during the backup.
15. **AG Backup Preferences:** Select this option if AG databases will be backed up.
    - **Use Server Preferences** uses MS SQL preferences.
    - **Override Preferences** enables you to override MS SQL's preferences with your selection.

    AAG replica preference for MS SQL backups

Cohesity uses "replica priority" to select the best replica when more than one qualified replica matches the backup preference. For the following AAG Backup Preference settings, Cohesity uses replicas in the indicated order of preference to back up MS SQL databases:

Backup Preference Setting	Replica Used for MS SQL Backup
Prefer Secondary or Any	1. Sync Secondary Replica 2. Async Secondary Replica 3. Primary Replica
Secondary Only*	1. Sync Secondary Replica 2. Async Secondary Replica
Primary Only	1. Primary Replica

* If the AAG uses the **Secondary Only** Backup Preference setting, ensure the AAG replicas are set to "**Readable secondary**=**Yes**". If all replicas are set to "**Readable secondary**=**No**" with **Secondary Only** as the Backup Preference setting, then the backup of AAG databases fails with "*unavailable_for_vss*".

You can set the **Readable Secondary** field to **No** if the Backup Preference Setting is set to **Prefer Secondary or Any** or **Primary Only**.

> **Note:** You must apply the exclusion filter on each AAG replica to ensure the database is excluded when an SQL Server failover occurs.

> **Note:** AG secondary replicas do not support regular full backups, and differential backups. If the policy is configured to take regular full backups, or differential backups, then select override preferences to backup only from the primary replica. All other AG backup preferences are only available if the policy is configured to take copy-only full backups, and log backups. Incremental backups are not supported on the AG secondary nodes. For more details, see Backup Types Supported on Secondary Replicas.

16. **System Databases:** Does not apply to VDI-based Protection Groups. Only selected databases are backed up. When you select any system database, however, all system databases are selected automatically.

17. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.

18. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

19. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.

20. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.

21. **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

    For configuration details, see Configure Pre & Post Scripts.

22. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

For information about customizing indexing, see Customize Indexing.

Indexing is not supported for VDI-based Protection Groups. Without indexing, file/folder-level recovery and searching are not supported.

23. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.
24. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
25. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

    > **Note:**
    > The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

26. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
27. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Edit an MS SQL Protection Group

Perform the following to edit a SQL Server Protection Group:

1. Select **Data Protection > Protection**.
2. Locate the Protection Group you want to edit.
3. Mouse over the action icon (⋮) next to the Protection Group and select one of the following:
    1. **Run Now** - Start a run of this Protection Group. You have the option to back up one or more objects.
        1. **Backup all Objects in the Protection Group** - Select this option to backup all the databases in the protection group. Any databases excluded from the backup will not be included in the protection run.
        2. **Backup only Selected Objects in the Protection Group** - you can manually select the check box for the databases to be backed up at the Protection Group level.
           This backup of the selected database runs in parallel with the existing run-now request. However, if a subset of the selected database includes a run-now/scheduled backup already active, these databases will be skipped from the run-now being submitted.

           > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.
    2. **Pause Future Runs**
        1. Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the Protection Group is currently executing a run, the current run continues to execute and only future runs of this Protection Group are paused.
        2. Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this Protection Group.
    3. **Edit** - Edit the Protection Group settings of the protection group and click Protect.
    4. **Delete** - Deletes the protection group.
       Delete only the protection group and keep the snapshots or delete the snapshots also.
4. Save the Protection Group settings by clicking **Protect**.

Backup Microsoft SQL Server FILESTREAM Databases

This topic describes how to protect MS SQL FILESTREAM databases.

MS SQL Protection Group Support

The following table describes supported job types for MS SQL running on physical servers and VMs.

	File-based Jobs	VDI-based Jobs	Both Job Types Using the Same Server but Protecting Distinct Databases
Physical MS SQL	Not supported	Multiple VDI-based jobs protecting all or a subset of databases	Supported
VM MS SQL	Not supported	Not supported	Not supported

> **Note:** Protecting a FILESTREAM database (in general and not specific to memory-optimized) is performed only using VDI-based backup.

Protect MS SQL Overview

Perform the following steps in the Cohesity cluster to protect SQL Server:

1. Register the SQL Server you want to protect as a source with the Cohesity cluster. See Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs or Set Up Microsoft SQL Server FCI.
2. Select **Dashboard** and from the drop-down, select **MS SQL**.
3. From the **Hosts** tab, you can protect all databases on a host. Select **Protect Databases** from the host's actions menu.

   From the **Databases** tab, you can protect a single database. Select **Protect** from the database's actions menu.

   You can change which objects you want to protect when creating the job.

You can also protect MS SQL by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. Select the objects to protect.

Create an MS SQL Protection Group

1. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

> **Note:** This page displays registered MS SQL objects only, so only SQL servers are selected regardless of the other servers within the source object hierarchy.

2. **Source:** Select the objects to protect.

   Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the ☐ unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box (☑) indicates the object is selected to be backed up.

   If MS SQL belongs to an AG network, selecting that server displays options for protecting the AG.

   - Select all servers that belong to this AG network.

     > **Note:** If backing up VMs within an AG, protecting the AG databases requires that all servers within the AG are registered with the Cohesity cluster and are selected for protection. Otherwise AG databases are skipped and only local databases will be backed up. New servers added to AG are not automatically protected.

     When using this option when there are multiple AGs, only interlinked AG nodes are automatically selected. For example:

     - If there are AGs that use only Node1 and Node2 and other AGs that use Node3 and Node4, if you select Node2 as a source, only Node1 is added automatically.
     - If AG1 uses Node1 and Node2, AG2 uses Node2 and Node4, AG3 uses Node3 and Node5. If you select Node2, then these nodes are selected: Node1, Node2 and Node4. Node3 and Node5 are not selected.

     When adding nodes to an AG that is protected by Cohesity, ensure you register the node with Cohesity and add it to the Protection Group. Otherwise the AG databases will not be protected.

     When removing nodes from an AG that is protected by Cohesity, you can remove it from the Protection Group or keep it in the job; the AG databases remain protected in either case. If the removed node remains in the Protection Group, only its local databases are backed up. If you do not want to back up the removed node, you can safely remove it.

   - Select only this server.

Click **Select**.

Click **Add**.

3. If appropriate, set options for specific objects by clicking (**Edit Server Options**) next to the object. A successful SQL log backup automatically truncates the SQL logs.
4. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

   The selected storage domain stores the snapshots and logs created by this Protection Group.
5. **Backup Method:** Select **VDI-based**. VDI-based jobs can backup FILESTREAM databases.
6. **Make Full Backup Copy-only:** Ensure this is toggled off to ensure log backups are taken for FILESTREAM databases.
7. **User Databases:** Select how to handle AG databases during backup.
8. **System Databases:** Select whether to back up or skip system databases.
9. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

   The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
10. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
11. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.

- **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
- **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.

12. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
13. **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

    For configuration details, see Configure Pre & Post Scripts.
14. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.
15. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.
16. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
17. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

> **Note:**
> The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

18. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
19. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Edit an MS SQL Protection Group

1. Select **Data Protection > Protection**.
2. Locate the Protection Group you want to edit.
3. Mouse over the action icon (⋮) next to the job and select one of the following:
    1. **Run Now** - Start a run of this Protection Group. You have the option to back up one or more objects.
        1. **Backup all Objects in the Protection Group** - Select this option to backup all the databases in the protection group. Any databases excluded from the backup will not be included in the protection job run.
        2. **Backup only Selected Objects in the Protection Group** - you can manually select the databases to be backed up at the protection job level. However, the selection of the database that is excluded at the backup job level is not supported.

    You can run log backups in parallel with existing log backup runs. This helps reduce storage buildup on the primary SQL Server when multiple large backups are in progress.

    This feature is supported only for file-based, VDI-based, and Filestream (backed up using VDI-based method) backups.

> **Note:** When you run a log backup for a database backup job, any preceding backup job for the same database may skip the log backup if the new backup job is completed before the previous backup job for the database is started.

2. **Pause Future Runs**
    1. Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the job is currently executing a run, the current run continues to execute and only future runs of this job are paused.
    2. Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this job.
3. **Edit** - Edit job settings of the protection group and click Protect.
4. **Delete** - Deletes the protection group.

4. Save the job settings by clicking **Protect**.

Configure Microsoft SQL Server T-log Backup

Cohesity enhances the SQL Server T-log backup and restoration performance by splitting the T-log backup into multiple backup files and ingesting them in parallel through the Cohesity nodes.

This feature is applicable for File-based and VDI-based backups.

Follow the below steps to configure the T-log backup settings for a Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect > Databases > MS SQL Server**.
3. Click **Objects** and select the required objects to be protected.
4. In the **Policy** drop-down list, select a policy that includes Log Backup (Database). The **Log Backup Settings** option will be enabled.
5. Optionally configure the following details:
    - **WITH Clause**: Define the WITH clause that you want to use to customize the backup. For more information, see BACKUP (Transact-SQL). The default WITH clause is "*WITH MAXTRANSFERSIZE = 4194304, BUFFERCOUNT = 16*".

> **Important:** The WITH clause text box allows you to invoke the SQL Server general options during backups and restores. The text field accepts any SQL Server general backup or restore option like "ENCRYPTION". Note that specifying an option could add to the overall duration of the backup or restore process.

- **Number of Streams**: Define the number of Cohesity nodes that will be used for the backup process. Using multiple nodes provides more throughput for the backup process. The default number of streams is *1*.

6. Click **Protect**.

> **Note:**
> - Ensure that the agents pre-installed on the SQL VMware source are upgraded to the latest version.
> - If a replication job is configured for the SQL VMware log backup, ensure that the replication cluster is upgraded to the latest version to enable seamless recovery.

Requirements

Ensure the following requirements are met for using SQL log backups:

- The Cohesity cluster's VIPs on **port 11117** must be reachable from the SQL source.
- In a multi-tenant environment, the Hybrid Extender IPs on **port 11117** must be reachable from the SQL source.

> **Note:** If the required port is unreachable, the log backup will fail to start, especially after upgrading from a version lower than 7.2.

Recover Amazon RDS Custom for SQL Server

Cohesity provides the ability to recover the MS SQL databases protected by a Protection Group.

- Recover MS SQL Database to the Original Database.
- Recover MS SQL Database as a New Database in the Original SQL Server Instance.
- Recover MS SQL Database as a New Database in an Alternate SQL Server Instance.

Prerequisites

- Before you can restore MS SQL databases, a full backup must exist.
- Cohesity cluster VIPs must be reachable (DNS) and accessible (Ping) from the target MS SQL server.
- For recovery, the server and the cluster must be on the same VLAN.

The following topics provide guidance to help you restore MS SQL databases:

About Restoring Microsoft SQL Server Databases

A restore task copies a database from the backup copy and optionally applies transaction logs depending on the options you choose. Cohesity provides the ability to restore individual databases to a specific point in time.

The following options are supported when restoring an individual database:

- Restore to the original MS SQL server and original instance and overwrite the database.
- Restore to the original MS SQL server and original instance and rename the database.
- Restore to the original MS SQL server and an alternate instance as a specified database name.
- Restore to an alternate MS SQL server and instance as a specified database name.

Restoring physical server to physical server only is supported.

MS SQL DB Restore Example

The following example depicts restoring an individual MS SQL database running on a physical server to its original location and overwriting the existing database. The example assumes the following about the database:

- Failure has already occurred.
- Its MS SQL server is registered as MS SQL on the cluster.
- It was backed up by a Cohesity Protection Group that includes transaction logs.

The action steps in the following procedure are represented in the figure below with blue numbered time squares:

.

T1) The Cohesity user must select the following task options:
  - Restore to the original SQL server and instance and overwrite the database
  - An MS SQL database and snapshot
  - Point in time (optional)
  - Capture tail logs (optional)

T2) The Cohesity cluster copies the database files from the specified snapshot to the original location.

T3) If applicable, transaction logs are applied to the database for the specified point in time.

T4) The database is then brought online.

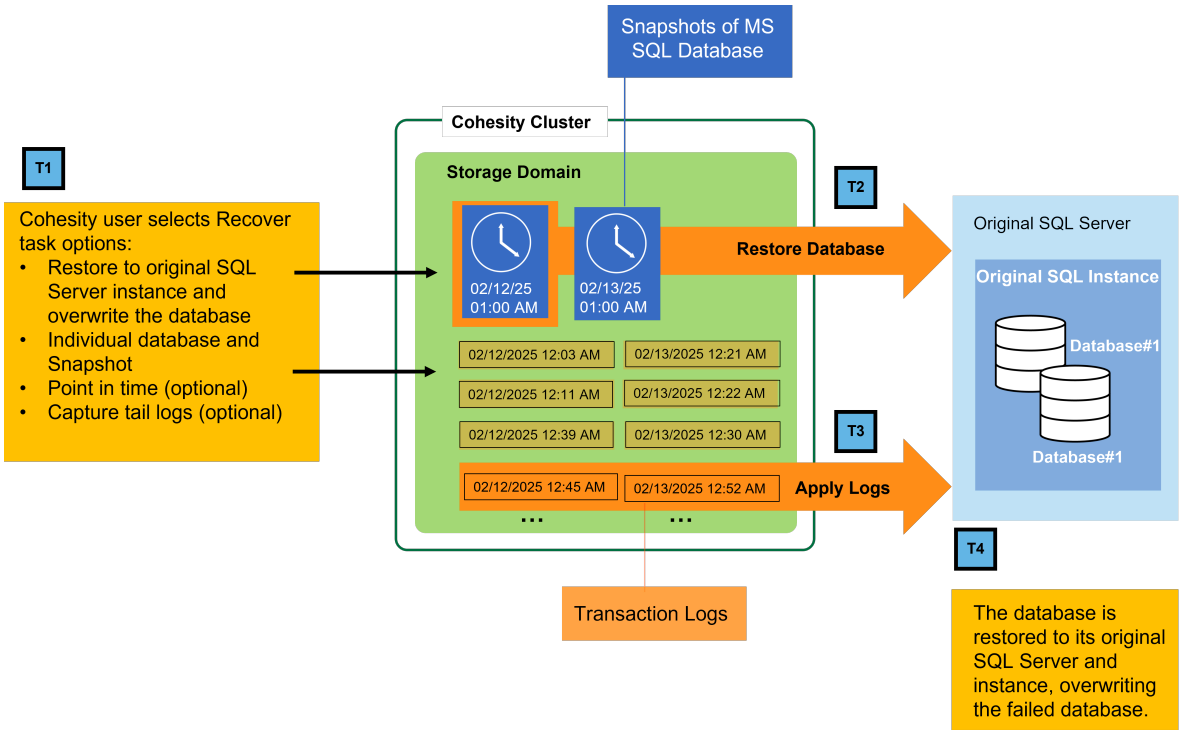

Resume MS SQL Database Recovery After Failure

You can resume a failed MS SQL Database recovery for databases in the following scenarios:

- If the databases are recovered using the WITH NORECOVERY option
- If the previous restore operation failed and the database remains in the restoring state.

This feature improves the RTO for database restore operations and is applicable only for point-in-time (PIT) restore operations.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Recover Microsoft SQL Server System Databases

System databases can be restored in case of system failures. The system databases are restored from the same instance and location where the system databases were backed up. You can also overwrite the existing system database files.

> **Note:**
> Restoring system databases to a different instance is not supported for the VM-based backup method.

The SQL Server system databases are recovered in the following order:

1. Master
2. MSDB
3. Model

Recovery Workflow

You can recover the system databases to the original SQL Server instance.

The following animated image shows how to recover the system databases to the original server instance:

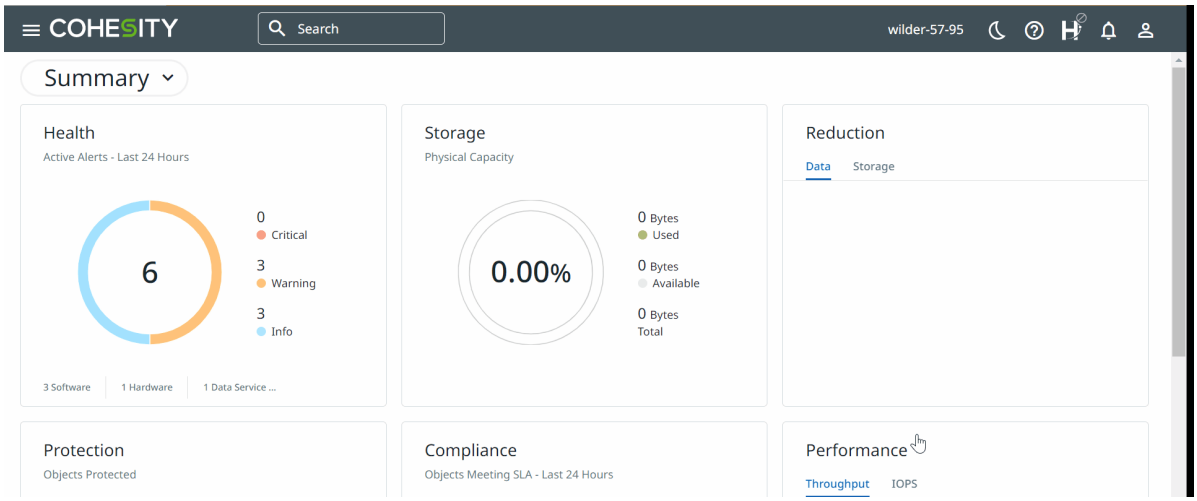

For the detailed steps on recovering system databases, see:

- Recover Microsoft SQL Server Database Over the Original Database

Consideration

- Recovering the system database to an alternate instance is not supported.
- When recovering to the original database, the original database is deleted before the recovery. Therefore, a recovery failure will also lead to the loss of the original database.
- If you are recovering the system databases individually, then the recovery of the system database must be in the following order:

1. Master
2. MSDB
3. Model

- You cannot recover a system database along with a user database.
- If you are recovering more than one system database in a recovery task, then the recovery point for the selected system databases must be the same.

Recover Microsoft SQL Server Database Over the Original Database

Recovering an MS SQL database over the original database overwrites the original database. The recovered database will have the original database name.

> **Note:** The original database is deleted before the recovery. Therefore a recovery failure will also lead to the loss of the original database.

To recover the database to the original database:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Databases** > **MS SQL**.

   The Restore SQL page appears.
3. Search for the databases you want to recover from by entering the characters of the database name, Protection Group name, or by specifying the wildcard character *.

   A list of databases is displayed with details like Cohesity Cluster, Protection Group, Physical Server, and Create Date. The Create Date option indicates the date and timestamp when the database was created.

   > **Note:** Operations like rename may update the creation date at the SQL Server.
4. From the list of search results, select one or more databases you want to recover. You can select multiple MS SQL databases only if they are protected by the same protection group.

   > **Note:** The option to select multiple databases in a recovery task is an Early Access feature. Contact Cohesity Support for enabling this feature.

   By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the edit icon.

   The Edit recovery point for <database_name> page appears.

By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

- To select a recovery point from the Timeline mode:
    a. From the **Choose a date** drop-down, select a date.

       The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.

    b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

    c. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    d. Click **Select Recovery Point** and then go to step 5.
- To select a recovery point from the List mode:
    a. Select the **List** tab from the top-right corner of the Edit recovery point for <database_name> page.

       By default, the page displays the recovery points available for the last 30 days.

    b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

    c. Select the recovery points to be recovered.

    d. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. From the **Targets** option, select **Overwrite Original Database** to recover the database to the original location.

> **Note:** The original database is deleted before the recovery. Therefore a recovery failure will also lead to the loss of the original database.

7. Specify the following **Recovery Options**:

Recovery Options	Description
VDI Restore Settings	Enable this option to specify the WITH clause that you want to use for the restore. **Note:** • This option is available only for restoring the databases that are protected using Protect MS SQL (VDI-based). • This option is not available for system database restore.
Keep CDC	Enable this option to restore a protected database that has the change data capture (CDC) enabled. If the protected database is not CDC enabled and the user tries to restore it with Keep CDC then, the database will be restored without CDC. **Note:** • This option is available only if the database is restored with recovery. • This option is not available for system database restore.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Grou**p drop-down.
Task Name	Change the default name of the recovery task.
Capture Tail Logs	Enable this option to capture records that have not yet been protected. They are captured to ensure all transactions are protected before restoring the database. **Note:** This option is not available for system database restore.

8. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Microsoft SQL Server VDI-based Recovery Recommendations

The Virtual Device Interface (VDI) is a Microsoft interface that allows the Cohesity agent to execute SQL Server Native backup and restore commands. This backup type produces a standard SQL Server Native backup of type (.BAK), transaction logs, (.TRN), and differentials, (.DIFF). These backup files are stored on a Cohesity View and take advantage of View compression, deduplication, encryption, replication, and archiving features.

The aim of the SQL DBA is to have a combination of backups so that the database can be restored to any point in time. A good combination of backups consists of having a full backup, differential (in Cohesity, incremental) backups, and log backups. For example, all SQL database restores must begin with a full database backup, secondly, a differential can then be applied to the full, and finally, log backups can be applied in sequence to complete the database restore.

When the backups are applied during the database restore process, you are sequentially adding the captured changes to the database: FULL+DIFF+Log1+Log2+Log3 = Restored Database.

All Microsoft VDI backups, their differentials, and their logs are dependent on a full backup to perform a database restore. Microsoft requires that in order to restore a SQL database, you must start with a full backup, then its transaction logs can be applied. This means your backup retention policy must keep a full backup along with its log backups in order to successfully restore a database.

Simply put, a SQL VDI-based database restore requires a full backup to seed the database, then differential and/or log backups are applied to the specified point in time.

We recommend retaining two sets of full backups with their differential and log backups.

> **Note:** A good backup plan always includes a combination of full, differential, and log backups.

Recover Microsoft SQL Server Database as a New Database in the Original Microsoft SQL Server Instance

To recover the database as a new database in the original MS SQL server instance:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Databases** > **MS SQL**.

   The Restore SQL page appears.
3. Search for the databases you want to recover from by entering characters of the database name, Protection Group name, or by specifying the wildcard character *. A list of databases is displayed with details like Cohesity Cluster, Protection Group, Physical Server, and Create Date. The Create Date option indicates the date and timestamp when the database was created.

> **Note:** Operations like rename may update the creation date at the SQL Server.

4. From the list of search results, select one or more databases you want to recover. You can select multiple MS SQL databases only if they are protected by the same protection group.

   > **Note:** The option to select multiple databases in a recovery task is an Early Access feature. Contact Cohesity Support for enabling this feature.

   By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the edit icon.

   The Edit recovery point for <database_name> page appears.

   By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

   - To select a recovery point from the Timeline mode:
     a. From the **Choose a date** drop-down, select a date.

        The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.

     b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

     c. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     d. Click **Select Recovery Point** and then go to step 5.
   - To select a recovery point from the List mode:
     a. Select the **List** tab from the top-right corner of the Edit recovery point for <database_name> page.

        By default, the page displays the recovery points available for the last 30 days.

b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

c. Select the recovery points to be recovered.

d. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. From the **Targets** option, select **Recover as a new Database** to recover the database as a new database.
7. Enable the **Restore to Original SQL Server Instance** option to restore the database in the original MS SQL Server instance.
8. In the **Data Files** field, specify the path to an existing directory. The newly created database files will reside in this path.

> **Note:** If the database file names to be restored have a conflict with the existing database files of the database to which files are restored, then Cohesity resolves the conflict by appending a number to the file names.

9. In the **Log Files** field, provide the full path location to which the log files can be restored.
10. (Optional) Click **Add File Path Rules** to add custom rules for the data files.

    You can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

    Accept the file patterns and examples:

    - Wildcard character used for file matching such as *
      - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
      - *m:*.ndf* - Every ndf file on drive m:
      - *c:*1.ndf* - Every ndf file on drive c: that contains 1.ndf
    - Regex pattern
      - *.*\.ndf* - All files with the ndf extension
      - *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf
    - Full path to a source database file

Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

Click **Add File Path Rule** to add multiple custom rules.

11. Specify the following **Recovery Options**:

Recovery Options	Description
Rename	Select one of the following options: • **Bulk Rename**: Select this option if you want to rename the restored databases in bulk. You can specify a string in the **Suffix** field. The restored databases will be renamed by adding this suffix string with the original database names. • **Rename Individual Objects**: Select this option if you want to individually rename the restored databases. You can specify the new name of the databases in the **New Name** field.
WITH RECOVERY	Disable the option to perform a restore WITH NO RECOVERY (the database will be offline when restored). By default, this option is enabled and the database is restored with recovery (the database will be online when restored).
VDI Restore Settings	Enable this option to specify the WITH clause that you want to use for the restore. **Note:** This option is available only for restoring the databases that are protected using Backup Microsoft SQL Server (VDI-based).
Keep CDC	Enable this option to restore a protected database that has the change data capture (CDC) enabled. If the protected database is not CDC enabled and the user tries to restore it with Keep CDC then, the database will be restored without CDC. **Note:** This option is available only if the database is restored with recovery.
Overwrite Alternate Database	Enable this option to recover the database by overwriting the existing database.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Grou**p drop-down.
Task Name	Change the default name of the recovery task.

12. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Microsoft SQL Server Database as a New Database in an Alternate Microsoft SQL Server Instance

To recover the database as a new database in an alternate MS SQL server instance:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Databases** > **MS SQL**.

   The Restore SQL page appears.
3. Search for the databases you want to recover from by entering characters of the database name, Protection Group name, or by specifying the wildcard character *. A list of databases is displayed with details like Cohesity Cluster, Protection Group, Physical Server, and Create Date. The Create Date option indicates the date and timestamp when the database was created.

   > **Note:** Operations like rename may update the creation date at the SQL Server.
4. From the list of search results, select one or more databases you want to recover. You can select multiple MS SQL databases only if they are protected by the same protection group.

   By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the edit icon.

   The Edit recovery point for <database_name> page appears.

   By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

   - To select a recovery point from the Timeline mode:
     a. From the **Choose a date** drop-down, select a date.

        The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.
     b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.
     c. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

d. Click **Select Recovery Point** and then go to step 5.

- To select a recovery point from the List mode:
  a. Select the **List** tab from the top-right corner of the Edit recovery point for <database_name> page.

     By default, the page displays the recovery points available for the last 30 days.

  b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

  c. Select the recovery points to be recovered.

  d. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

  e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. From the **Targets** option, select **Recover as a new Database** to recover the database as a new database.
7. Disable the **Restore to Original SQL Server Instance** option to restore the database to an alternate SQL Server instance.
8. From the **MS SQL Instance** drop-down, select an alternate SQL server instance.
9. In the **Data Files** field, specify the path to an existing directory. The newly created database files will reside in this path.

   > **Note:**
   > - If the database file names to be restored have a conflict with the existing database files of the database to which files are restored, then Cohesity resolves the conflict by appending a number to the file names.

10. In the **Log Files** field, provide the full path location to which the log files can be restored.
11. (Optional) Click **Add File Path Rules** to add custom rules for the data files.

    You can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

This is only used for the destination path in a custom rule and not the source. File path rules are case sensitive. Partial file name or extension (not only directory path) are required for a match to be identified.

> **Note:** For VDI based backup, use two backslashes after the drive. For example, C:\\path.

Accept the file patterns and examples:

- Wildcard character used for file matching such as *
  - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
  - *m:*.ndf* - Every ndf file on drive m:
  - *c:*1.ndf* - Every ndf file on drive c: with a suffix of 1.ndf
- Regex pattern
  - *.*\.ndf* - All files with the ndf extension
  - *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf
- Full path to a source database file

Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

Click **Add File Path Rule** to add multiple custom rules.

12. Specify the following **Recovery Options**:

Recovery Options	Description
Rename	Optional. Enable the Rename option and select one of the following options to rename the restored databases: • **Bulk Rename**: Select this option if you want to rename the restored databases in bulk. You can specify a string in the **Suffix** field. The restored databases will be renamed by adding this suffix string with the original database names. • **Rename Individual Objects**: Select this option if you want to individually rename the restored databases. You can specify the new name of the databases in the **New Name** field.

Recovery Options	Description
WITH RECOVERY	Disable the option to perform a restore WITH NO RECOVERY (the database will be offline when restored). By default, this option is enabled and the database is restored with recovery (the database will be online when restored). **Note:** When restoring a database to an alternate MS SQL server instance with no recovery, a staging directory is created under the database files and logs file directories. To overcome the filename conflict with existing files in the database, the recovered files are saved in this staging directory.
VDI Restore Settings	Enable this option to specify the WITH clause that you want to use for the restore. **Note:** • This option is available only for restoring the databases that are protected using Protect MS SQL (VDI-based).
Keep CDC	Enable this option to restore a protected database that has the change data capture (CDC) enabled. If the protected database is not CDC enabled and the user tries to restore it with Keep CDC then, the database will be restored without CDC. **Note:** • This option is available only if the database is restored with recovery. • This option is not available for system database restore.
Overwrite Alternate Database	Enable this option to recover the database by overwriting the existing database.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Grou**p drop-down.
Task Name	Change the default name of the recovery task.

13. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the recoveries page.

Recover Microsoft SQL Server FILESTREAM Databases

Before you can restore MS SQL FILESTREAM databases, a snapshot must exist.

Restore Database Overview

1. Select **Dashboard** and from the drop-down, select **MS SQL**.
2. Click **Databases**.
3. Select **Recover** from the database's actions menu.
4. Set restore options.

You can also start a restore task by navigating to **Data Protection > Recoveries** > **Recover** > **Databases** > **MS SQL** and selecting a database to restore.

Set Restore Options

1. **Task Name:** Optionally change the default name for the task.
2. **Restore to Original SQL Server Instance:** Toggle on if you want to restore the database to its original instance.

   Overwriting a FILESTREAM database is not supported and prevents the restore from completing.

   a. Select a **SQL Host** and **SQL Instance** if restoring to an alternate MS SQL instance.

   > **Note:** Local MS SQL instances that are running within an MS SQL FCI do not display. You can, however, manually enter the local SQL instance name to restore to.

   b. Provide the full path location to restore all database files.

   You can optionally restore data files and log files to different locations. Click **Add Alternative Log File Locations and Custom Rules** to provide locations for data and log files and any custom rules. For data files, you can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

   This is only used for the destination path in a custom rule and not the source.

   Accept the file patterns and examples:

   - Wildcard character used for file matching such as *
     - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
     - *m:*.ndf* - Every ndf file on drive m:
     - *c:*1.ndf* - Every ndf file on drive c: with a suffix of 1.ndf
   - Regex pattern

- *.*\.ndf* - All files with the ndf extension
- *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf

- Full path to a source database file

Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

Click **Add Another Custom Rule** to add multiple custom rules.

> **Note:** Ensure to follow the standard regular expression rules. Specify a dot '.' before asterisk '*' to match all characters, for example, L:.*.ndf.

c. Provide a **Database Name** for the restored database.

3. By default, the latest snapshot is selected. You can optionally choose to restore from a different snapshot by clicking Edit ( ).

   You can also provide a specific time to restore to if transaction logs are available.

   Click a date on the calendar to select a snapshot from that date. The default point selected will be the latest log transaction or the latest full or incremental, whichever is the latest. The slider also displays available points in time and displays snapshots as circles. Manually select a point on the slider if you want to restore to a specific point in time. Selecting an invalid time from the slider automatically changes to the closest available snapshot. If your manual selection is within 10 minutes of a snapshot, your selection automatically snaps to it. You can also select a specific point in time by entering it in the time fields.

   After selecting the desired recover point, click **Recover**. You can choose to clone the selected recover point by clicking the arrow adjacent to recover and selecting **Clone**.

   In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

4. By default, an MS SQL restore WITH RECOVERY is performed. You can optionally toggle this off to perform a restore WITH NORECOVERY.
5. **Auto Select:** When enabled, the recovery task automatically selects the correct VLAN.

   > **Note:** If using VLANs and the cluster is joined to an AD domain—One of the machine accounts must be a DNS hostname for a virtual IP (VIP) of a VLAN configured on the cluster in order to recover to servers within that VLAN.

6. **Capture tail logs:** This option is available only when overwriting a database in its original location. This option is unavailable if the database's Protection Group has been paused or deleted.

   Tail logs capture records that have not yet been backed up. They are captured to ensure all transactions are backed up before restoring the database.

7. Click **Recover** to start the task.

SQL Server Maintenance Mode

You can enable maintenance mode for SQL Server instances and databases. Enabling maintenance mode skips any scheduled or manually initiated protection runs.

To place a SQL Server instance or database in maintenance mode:

1. Navigate to **Data Protection > Sources**.
2. Click on the action menu (⋮) of the source and select **Maintenance**.
3. You can initiate the maintenance mode immediately or later.

   To initiate the maintenance mode immediately:

   1. Select **Start Now**. The **Start Maintenance Now** dialog displays the list of Protection Groups available in the source.
   2. Review the list of Protection Groups and enter **Yes** to confirm.
   3. Optionally, you can enter a note.
   4. Click **Confirm**.

   To schedule the maintenance mode:

   1. Select **Schedule**.
   2. On the **Schedule Maintenance** dialog, enter the start date and time and the end date and time.
   3. Optionally, you can enter a note.
   4. Click **Save**.

The SQL Server instance or database will be set in maintenance mode accordingly.

AG in Maintenance Mode

When maintenance mode is enabled for a SQL instance hosting the AG database, the replica hosting the instance is not considered for the AG backup plan. This behavior is common to both legacy (non-first-class citizen) and first-class citizen features.

For legacy AG backup, only the replica for which maintenance mode is enabled will be skipped. You must enable maintenance mode on all the replicas to skip the entire backup.

Migrate Microsoft SQL Server

The following topics provide guidance to help you migrate MS SQL databases:

About Migrating Microsoft SQL Server Databases

You can use MS SQL database migration for the following use cases:

- Back up and migrate a database from a single pane of glass.

  Leverage database backup to achieve database migration. No more shipping logs across multiple sites or multiple data centers.

  Upgrade the MS SQL Server version while minimizing application downtime and maintaining the ability to restore to any point in time.
- Staging a database on the target server for faster recovery with Cohesity's database migration.

  Stage a database on the target SQL server and perform only incremental updates at the file level to improve RPO.

Database migration for file-based MS SQL backups consists of the following actions, which must be initiated manually:

1. Perform the initial copy of database files to the target MS SQL server.
2. Sync forward to the latest incremental/full snapshot available on the cluster. You can trigger multiple syncs.

   For each sync, Cohesity performs incremental file updates to database files on the target MS SQL server. Syncs ensure files are up to date. The files are not attached to the database until you complete the finalized workflow. This ensures that the database is up to date when finalization occurs.
3. Finalize the migration. Finalizing is available only if the target database is synced to the latest snapshot.

   This brings the database online. After the database is online and the finalize step succeeds, the recovery task is completed successfully. You cannot perform syncs after this point because the migration workflow is complete.

Migrate Microsoft SQL Server Databases

Before you migrate an MS SQL database, ensure that a file-based snapshot exists on the cluster.

Migrate MS SQL Database

This section lists the steps to migrate an MS SQL database.

1. Select **Dashboard** and from the drop-down, select **MS SQL**.
2. Select **Migrate** from the database's actions menu.

   You can also start a restore task by navigating to **Data Protection > Recoveries> Recover > Migrate Database (MS SQL only)**.

3. Search for an MS SQL server or database. Enter characters of the name and a list of items displays. You can optionally specify the wildcard character *. For more information, see Use Wildcard * in Search for Recovery. You can also optionally narrow the search results by specifying filter criteria, for example, you can filter the search results by a specific Protection Group. Click the **Add Filters** icon, specify the filter and click **Add**.
4. Select a database and set the migration options for the MS SQL database you want to migrate.
    1. **Task Name:** Optionally change the default name for the task.
    2. **Objects Being Recovered**: This displays the database to be migrated.
    3. **Sync Options**: Specify the sync options for the database to be migrated.

        > **Note:** Transaction log backup is not synced. Only the incremental and full backup is synced.

        - **Auto Sync Databases**: This option automatically keeps the database in sync with the latest available snapshot until you finalize the database and bring it online.
        - **Manually Sync Databases**: Once the initial database migration is completed, you can sync the migrated database files with newer snapshots available until you finalize the database and bring it online.
    4. **Recover Point:** This displays the selected recover point for the migration.
    5. **Settings:** Specify the **SQL Host** and **SQL Instance** to migrate the database to.

        Provide the full path location to restore all database files.

        You can optionally restore data files and log files to different locations. Click **Add Alternative Log File Locations and Custom Rules** to provide locations for data and log files and any custom rules. For data files, you can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

        This is only used for the destination path in a custom rule and not the source.

        Accept the file patterns and examples:

        - Wildcard character used for file matching such as *
            - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
            - *m:*.ndf* - Every ndf file on drive m:
            - *c:*1.ndf* - Every ndf file on drive c: with a suffix of 1.ndf

- Regex pattern
  - *.*\.ndf* - All files with the ndf extension
  - *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf
- Full path to a source database file

Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

Click **Add Another Custom Rule** to add multiple custom rules.

> **Note:** Ensure to follow the standard regular expression rules. Specify a dot '.' before asterisk '*' to match all characters, for example, L:.*.ndf.

5. By default, an MS SQL restore WITH RECOVERY is performed. You can optionally toggle this off to perform a restore WITH NO RECOVERY.
6. Click **Start Migration**.

   The initial file copy task starts. This is the first step in the migration process.
7. After the file copy task is completed, the options available for the next migration step are displayed.
   - **Finalize** - Available if the target database is synchronized to the latest snapshot.
   - **Sync** - Available if there is a snapshot that is more recent than the target database.

   If you select the **Auto Sync Databases** option in Step 4, the **Auto Sync to latest** toggle is enabled by default. You can manually disable this option if required.

> **Note:** Database migration is not supported for AG databases.

Cancel Migration

You can cancel an in-progress migration task. Canceling the migration prevents syncing or finalizing the database and may leave it in an unusable state. Delete the database files manually and then start another migration task as needed.

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Find the migration task in the list and click the task name to display its detail page.
3. Click the **Cancel** button located at the top right of the page.
4. Respond to the confirmation prompt.

Amazon Relational Database Service for MariaDB

Amazon Relational Database Service (RDS) for MariaDB is a fully managed, cloud-based relational database service offered by Amazon Web Services (AWS). Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Amazon Relational Database Service for MariaDB instances.

> **Important:**
> Cohesity supports the backup and recovery of Amazon RDS for MariaDB only at the instance level.

This topic provides information on how you can protect Amazon RDS for MariaDB on Cohesity Cloud Edition on AWS. For information on protecting Amazon RDS for MariaDB on Physical and Virtual Edition clusters, see Amazon RDS and Amazon Aurora.

Plan and Prepare for Amazon RDS for MariaDB Protection

To protect Amazon RDS for MariaDB with Cohesity, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. Ensure you meet the pre-requisites and understand the considerations before registering the AWS source with the Cohesity cluster.

Ports Used for Communication

For information on port numbers required, see Amazon Elastic Block Store.

Considerations for AWS RDS for MariaDB Data Protection

Review and understand the following considerations before you protect Amazon RDS for MariaDB:

- Cohesity supports only snapshot-based protection of RDS for MariaDB. Ingest-based protection is not supported.
- Cross-account restore of an Amazon RDS for MariaDB protected using the AWS Snapshot protection type is not supported.
- When recovering encrypted RDS instance to the same AWS account but a different AWS region, you must create a KMS encryption key in the target region with the same alias name as the KMS key used to encrypt the source RDS instance.
- When performing cross-region cross-account recovery of an Amazon RDS instance encrypted with a customer-managed KMS key, ensure the following:
    - The source account's KMS key alias is available in both the source and destination regions of the destination account.

  Example:

If A1 is the source account with region R1, and A2 is the destination account with region R2, ensure the KMS key alias from A1 is available in both R1 and R2 of A2.

- The destination account has permission to access the source account's KMS key alias.
- The option group from the source account is available in both the source and destination regions of the destination account.

Backup Amazon RDS for MariaDB

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads including Amazon RDS for MariaDB associated with the registered AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **RDS MariaDB** tab to view all discovered instances.

The **RDS MariaDB** tab displays the entity hierarchy for the Amazon RDS for MariaDB environment, allowing navigation from the organization or account level down to individual RDS for MariaDB instances.

RDS MariaDB Entity Hierarchy

The **RDS MariaDB** tab displays the entity hierarchy for the RDS instance, allowing you to navigate from the organization or account level down to individual RDS for MariaDB instances.

The **RDS MariaDB** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with Cohesity for protecting Amazon RDS for MariaDB. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The organizational units within the AWS organization that contain member accounts.

Structure	Description
Member Account	The AWS member accounts within the organization that contain Amazon RDS for MariaDB.
AWS Account	The AWS account registered with Cohesity for RDS protection. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the RDS resources reside.
Availability Zone	The availability zone within the region where the RDS instance is deployed.
RDS Instance	The specific RDS MariaDB instance within the selected availability zone.

Backup Scope by RDS MariaDB Hierarchy

The Cohesity cluster allows you to back up Amazon RDS for MariaDB at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level.

- **AWS Organization** – Protect all RDS for MariaDB instances across all member accounts in the registered AWS organization.
- **Organizational Unit (OU)** – Protect all RDS for MariaDB instances within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all RDS for MariaDB instances within a specific member account in the registered AWS organization.

- **AWS Account** – Protect all RDS for MariaDB instances within the registered standalone AWS account.
- **AWS Region** – Protect all RDS for MariaDB instances within a specific AWS region.
- **Availability Zone** – Protect RDS for MariaDB instances within a specific Availability Zone.
- **RDS Instance** – Protect a specific RDS for MariaDB instance.

RDS for MariaDB Backup Methods

AWS Snapshot method: Cohesity uses AWS Snapshot method to protect Amazon RDS for MariaDB instances. In this method, Cohesity leverages the AWS native snapshot and stores them in the same AWS account and region as the source Amazon RDS for MariaDB instance.

Supported Backup Types

Cohesity performs only full backup of Amazon RDS for MariaDB.

Protect Amazon RDS for MariaDB Using AWS Snapshot Method

With the AWS Snapshot protection method, Cohesity cluster protects the Amazon RDS MariaDB of your AWS account at the instance level.

To protect your Amazon RDS for MariaDB instance using AWS snapshot method:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **RDS MariaDB**.
3. In the New Protection page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Amazon RDS for MariaDB instances for protection using one of the following views:
    - **List View**: Use this option to select the RDS for MariaDB instances for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select instances for protection at various hierarchy levels.

      Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new RDS for MariaDB instances added to your AWS source.
        - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
            - **AWS Organization**
            - **Organizational Unit (OU)**

    - **Member Account**
    - **AWS Region**
    - **Availability Zone**
  - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
    - **AWS Account**
    - **AWS Region**
    - **Availability Zone**

  You can also exclude the Amazon RDS for MariaDB instances from auto-protection within an account, region, or availability zone by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the instances for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all instances with that tag for protection.

  Optionally,

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Amazon RDS for MariaDB instances with this tag.
  - You can exclude the Amazon RDS for MariaDB instances from auto-protection within a tag by clicking the auto-protect icon ( ) next to the instance. The instances excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude instances with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. Create or select an existing policy from the **Policy** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon RDS for MariaDB Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon RDS for MariaDB instances to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Amazon RDS for MariaDB.

3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MariaDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MariaDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon RDS for MariaDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.

- Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.

5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon RDS for MariaDB Instance

You can recover the Amazon RDS for MariaDB instance with to the same location or to a different location:

> **Note:**
> - Cohesity supports only instance-level recovery of Amazon RDS for MariaDB instance. Database-level recovery is not supported.
> - We recommend that you review the Plan and Prepare for Amazon RDS for MariaDB Protection.

To recover Amazon RDS for MariaDB:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **RDS MariaDB**.

1. Select **Instance** as the recovery task.
2. Search for the instance by entering characters of the instance name, or by specifying the wildcard character *.
3. From the list of search results, select the instance to recover. By default, the latest snapshot of the instance is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for <*instance_name*>** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:

     a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*instance_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. You can also select any point on the green solid line on the slider to restore to a specific point in time. Selecting an invalid time from the slider

automatically selects the closest available snapshot.

b. Select the snapshots to be recovered.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

- To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

4. Click **Next: Recover Options**.
5. In the **Recover To** section, select any one of the following options based on where you want to restore the instance and then provide the required information:
   - **Original Location**: Restore the instance to the same region.

     > **Note:** If an RDS instance with the same name already exists, the restore will fail. Ensure the original instance is not present or restore with a different name.

     Type **Yes** to confirm.
   - **New Location**: Restore the instance to a different region. You can restore the instance to the same or a different region within the same account as the original RDS instance.

     Based on the registration level of the AWS source, select the recovery destination hierarchy:
     - If the AWS source is registered at the organization level:
       - **Source**: Select the AWS organization that contains the destination instance. You can recover the instance within the same organization as the source instance.
       - **AWS Account**: Select the member account to which you want to restore the instance.
       - **Region**: Select a region to restore the instance.

- **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
- **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
- **Availability Zone**: Select an availability zone in AWS to recover the RDS.

- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account to which you want to restore the instance.
  - **Region**: Select a region to restore the instance.
  - **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered instance.
  - **Subnet**: Select a subnet in the Amazon VPC to store the recovered instance.
  - **Availability Zone**: Select an availability zone in AWS to recover the RDS.

To overwrite any existing instance with the same name, enable **Overwrite existing instance(s)** and type **Yes** to confirm.

> **Note:**
> - Cross-account restore of RDS instance protected using the AWS Snapshot protection type is not supported.
> - If an RDS instance with the same name already exists, the restore will fail.

6. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
7. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
8. **DB Parameter Group**: Select the database (DB) parameter group to associate with the DB instance.
9. **DB Port:** Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port. The default port is 3306.

10. Configure the following **Recovery Options**:
    a. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
    b. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
    c. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
    d. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
    e. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.
11. **Task name**: Optionally, change the default name of the recovery task.
12. Click **Recover**.

Amazon DocumentDB

Amazon DocumentDB is a fully managed, MongoDB API–compatible document database service. Cohesity natively integrates with AWS to provide simple, fast, and cost-effective backup, recovery, and data management for both Instance-based DocumentDB Clusters and Elastic DocumentDB Clusters in your AWS source.

> **Important:** The backup and recovery of Amazon DocumentDB is supported only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster).

Plan and Prepare for Amazon DocumentDB Protection

Before protecting Amazon DocumentDB with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon DocumentDB with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Supported Amazon DocumentDB Cluster Types

Cohesity supports database-level protection for the following DocumentDB cluster types in your AWS account:

- Instance-based DocumentDB clusters
- DocumentDB Elastic Clusters

Requirements for Amazon DocumentDB Instance-based Cluster Protection

To protect Amazon DocumentDB Instance-based Cluster, you can register:

- AWS at account level using both **IAM User** and **IAM Role** methods.
- AWS at organization level using both **Preconfigured IAM Roles** and **CFT** methods.

Supported versions

The following table lists the DocumentDB Instance-based Cluster versions that are supported based on the registration method you choose.

Registration Method	Supported DocumentDB Versions
IAM User registration	Version 5.x and below
IAM Role registration	Version 5.x

Authentication Requirements

When creating a Protection Group to protect Amazon DocumentDB at the database level, the Cohesity cluster must authenticate to the DocumentDB primary instance using one of the following methods:

- **Username and Password Authentication**: Username and password of the Amazon DocumentDB Cluster instances. Provide the DocumentDB Cluster's database username and password.
- **IAM-authenticated DocumentDB User**: This option is available only if the AWS account is registered using the IAM Role method.

If you choose the **IAM-authenticated DocumentDB** user method, ensure the following:

- DocumentDB supports the MongoDB AWS authentication mechanism (MONGODB-AWS).
- The database user must be created inside the Amazon DocumentDB cluster.
- The database username must be the ARN of the IAM Role attached to the Cohesity cluster nodes (when registering the AWS account).
- The IAM Role must allow the `sts:AssumeRole` flow for MONGODB-AWS authentication.

To create the IAM-Authenticated DocumentDB User:

1. Connect to the DocumentDB cluster using **mongosh**.
2. Switch to the external authentication database:

```
use $external;
```

3. Create the user with appropriate roles:

```
db.createUser({
    user: "<ARN of IAM Role attached to EC2/NGCE>",
    mechanisms: ["MONGODB-AWS"],
    roles: [
        { db: "readWriteDB", role: "readWrite"},
        { db: "admin", role: "readWriteAnyDatabase"},
        { db: "admin", role: "userAdmin"},
        { db: "admin", role: "dbAdminAnyDatabase"}
    ]
});
```

For information on the steps to authenticate the Cohesity cluster to the Amazon DocumentDB in the Cohesity cluster UI, see Protect Amazon DocumentDB.

Requirements for Amazon DocumentDB Elastic Cluster Protection

To protect Amazon DocumentDB Elastic Cluster, you must register AWS at account level using the **IAM User** method.

Authentication Requirements

To protect Amazon DocumentDB Elastic Clusters at the database level, authentication must be performed using the username and password of the DocumentDB Elastic cluster.

Download the AWS CA Certificate Bundle for Amazon DocumentDB

For secure TLS communication between the Cohesity cluster and Amazon DocumentDB, you must download the AWS CA certificate bundle (PEM file) to the Cohesity cluster.

To download the PEM file:

1. SSH to the Cohesity cluster:

```
ssh -p 2222 support@<node_IP_address>21`1
```

   Ensure that the support user is enabled and a password is configured. For more information, see Manage Support User.

2. Switch to Cohesity user:

```
sudo su - cohesity
```

3. Navigate to the following directory:

```
cd /home/cohesity/bin/dataprotect_connectors/AWSDocumentDB
```

4. Download the AWS CA certificate bundle:

```
wget https://truststore.pki.rds.amazonaws.com/global/global-bundle.pem
```

The **global-bundle.pem** file is downloaded to the following location on the Cohesity cluster:

```
/home/cohesity/bin/dataprotect_connectors/AWSDocumentDB
```

> **Important:** Repeat this procedure on every node in the Cohesity cluster, and after each cluster upgrade. The PEM file must be present on all cluster nodes.

Ports Used for Communication

For information on port numbers required, seeAWS Cloud Database Workloads.

Considerations for Amazon DocumentDB Protection

Review and understand the following limitations and known issues before you protect Amazon DocumentDB database:

- Cohesity performs only full backup of Amazon DocumentDB.
- DocumentDB Version 8 for instance-based clusters is not supported for protection.
- IAM-authentication using DocumentDB User is not supported for DocumentDB Elastic Cluster.
- Restoring DocumentDB to an alternate AWS account and region is not supported.
- Restoring DocumentDB Elastic Cluster databases to DocumentDB instance-based cluster and vice-versa are not supported.
- Restoring a backup taken from a higher DocumentDB database version to a lower version is not supported.

Backup Amazon DocumentDB

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads in the specified AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **DocumentDB** tab to view all discovered databases. For information on registering AWS, see Register AWS Source.

The DocumentDB tab displays the entity hierarchy for Amazon DocumentDB, allowing navigation from the organization or account level down to individual databases.

Amazon DocumentDB Entity Hierarchy

The Amazon DocumentDB entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS Organization registered with Cohesity for Amazon DocumentDB protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The Organizational Units within the AWS Organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain Amazon DocumentDB.
AWS Account	The AWS account registered with Cohesity for DocumentDB protection. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the AWS resources reside.
DocumentDB Cluster	The specific Amazon DocumentDB cluster within the selected region.
DocumentDB Database	The individual databases hosted within the DocumentDB cluster. The database will be listed in the entity hierarchy only if you have provided credentials to discover the databases of the Amazon DocumentDB cluster.

Backup Scope by Amazon DocumentDB Hierarchy

The Cohesity cluster allows you to back up Amazon DocumentDB databases at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level. You can choose to back up:

- **AWS Organization** – Protect all Amazon DocumentDB clusters and databases across all member accounts in the registered AWS organization.
- **Organizational Unit (OU)** – Protect all Amazon DocumentDB clusters and databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all Amazon DocumentDB clusters and databases within a specific member account in the registered AWS organization.
- **AWS Account** – Protect all the Amazon DocumentDB clusters and databases within the registered AWS account.
- **AWS Region** – Protect all the Amazon DocumentDB clusters and databases within a specific region of the registered account.

- **DocumentDB Cluster** – Protect all Amazon DocumentDB databases within a specific Amazon DocumentDB cluster and all databases hosted on it.
- **DocumentDB Databases** – Protect individual Amazon DocumentDB database by providing the database credentials.

Supported Backup Types

Cohesity supports only full backups for Amazon DocumentDB databases. Incremental backups are not supported.

Protect Amazon DocumentDB

In Cohesity, Amazon DocumentDB databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

Prerequisites for Amazon DocumentDB Protection

Authenticate for DocumentDB Protection

To protect the Amazon DocumentDB at database level, you must authenticate the DocumentDB primary instance. Once authentication is successful, Cohesity discovers the databases of the DocumentDB for protection.

Perform the following steps for authentication:

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AWS source and click into it.
3. Click the **DocumentDB** tab.
4. Select the checkbox for the hierarchy level (AWS Organization, OU, member account, AWS account, region, or cluster level), and then click **Database Credentials** above the checkboxes.

   The **Database Credentials** for AWS DocumentDB dialog appears.
5. Select one of the following authentication methods for Amazon DocumentDB database-level discovery:
   1. **Credentials**: Enter the Username, and Password for the Amazon DocumentDB cluster.

> **Note:** This option is supported for both Instance-based DocumentDB clusters and Elastic DocumentDB clusters.

2. **IAM Authentication**: Enter the Username for IAM authentication.

   This option is supported only if:

   - The DocumentDB cluster is an instance-based cluster.
   - The cluster runs DocumentDB version 5.
   - The AWS source was registered in Cohesity using the IAM role method.

   The databases corresponding to the selected DocumentDB cluster or hierarchy level are discovered. For more information, see Authentication Requirements.

6. Click **Save**.
7. Click the **Refresh** icon to discover the databases.

Protect Amazon DocumentDB Database

To protect Amazon DocumentDB databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **AWS** > **DocumentDB**.
2. In the **New Protection** page, click **Add Objects** and select your registered source, or click **Register Source** to add a new one.
3. Select the DocumentDB databases using one of the following views or option:
   - **List View**: Use this option to select the DocumentDB databases for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select DocumentDB databases for protection at various hierarchy levels.

     Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon RDS for MS SQL Server instances.

     - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
       - **AWS Organization**
       - **Organizational Unit (OU)**
       - **Member Account**

        - **AWS Region**
        - **DocumentDB Cluster**
    - If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
        - **AWS Account**
        - **AWS Region**
        - **DocumentDB Cluster**

    You can exclude the Amazon DocumentDB databases from auto-protection within an account or region by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the database for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all DocumentDB databases with that tag for protection.

  Optionally:

    - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added DocumentDB databases with this tag.
    - You can exclude the Amazon DocumentDB databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The DocumentDB databases excluded from auto-protection are displayed with the exclude ( ) icon.
    - You can exclude DocumentDB databases with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

4. Click **Continue.**
5. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
6. From the **Policy** drop-down, create or select an existing policy.
7. Select **More Options** to change or configure any of the Advanced Settings.
8. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Section	Setting	Description
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Section	Setting	Description
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon DocumentDB Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon DocumentDB databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for DocumentDB databases.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon DocumentDB databases.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon DocumentDB databases.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for DocumentDB databases.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon DocumentDB

You can recover the Amazon DocumentDB database either to its original location or to a different location.

You can also perform PITR when restoring to the original AWS account where the backups were created. Restoring to an alternate AWS account using PITR is not supported.

> **Note:** We recommend that you review the Plan and Prepare for Amazon DocumentDB Protection topic.

Recover Amazon DocumentDB Database

To recover Amazon DocumentDB database:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **DocumentDB**.
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
   - By default, the latest snapshot of the database is selected for recovery.
   - To recover from a different snapshot, hover over the snapshot and click the **Edit** icon. The **Edit Recovery Point for <*database*_name>** page appears.
   - By default, snapshots are displayed in the **List** mode. You can switch to the Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   To select a snapshot from the **List** mode:

   a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*database_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

   b. Select the snapshots to be recovered.

   c. By default, the snapshot is recovered from the local Cohesity cluster.

   d. Click **Select Recovery Point** and then go to step 5.

   To select a snapshot from the **Timeline** mode:

   a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

   b. Select the snapshot to be recovered by clicking the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

   c. By default, the snapshot is recovered from the local Cohesity cluster.

   d. Click **Select Recovery Point** and then go to step 5.
5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:
   - **Original Location**: Restore the database to the original Amazon DocumentDB Cluster.

> **Note:** Restoring to the original location overwrites the existing database.

Type **Yes** to confirm.

- **New Location:** Restore the database to an Amazon DocumentDB Cluster within the same region and AWS account. Restoring to another region in the same AWS account and restoring to a different AWS account are not supported.

  Based on the registration level of the AWS source, select the recovery destination hierarchy:

  - If the AWS source is registered at the organization level:
    - **Source**: Select the AWS organization that contains the destination Amazon DocumentDB cluster. You can recover data only within the same AWS Organization as the source.
    - **Account**: Select the AWS member account to which you want to restore the database.
    - **Cluster**: Select the Amazon DocumentDB cluster to which you want to restore the database.
  - If the AWS source is registered at the account level:
    - **Source**: Select the AWS account to which you want to restore the database.
    - **Cluster**: Select the Amazon DocumentDB cluster to which you want to restore the database.

  To overwrite any existing database with the same name, enable Overwrite existing database(s) and type **Yes** to confirm.

> **Note:**
> Restoring to a new location also replaces any existing database with the same name in the selected instance.

7. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Click the Edit icon and enter a New Name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

8. Click **Recover**.

Amazon DynamoDB

Amazon DynamoDB is a fully managed NoSQL database service provided by AWS, designed for high performance at scale. In DynamoDB, data is stored in tables. Cohesity natively integrates with AWS to provide a simple, fast, and cost-effective solution for the backup, recovery, and data management of both the DynamoDB Standard table class and the DynamoDB Standard-Infrequent Access (DynamoDB Standard-IA) table class in your AWS source.

> **Important:** The backup and recovery of Amazon DynamoDB is supported only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster). This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Plan and Prepare for Amazon DynamoDB Protection

Before protecting Amazon DynamoDB with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Amazon DynamoDB with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Ports Used for Communication

For information on port numbers required, seeAWS Cloud Database Workloads.

Considerations for Amazon DynamoDB Protection

Review and understand the following limitations and known issues before protecting Amazon DynamoDB tables:

- **General Protection Limitations**
  - Tables in Amazon DynamoDB can be protected across all regions in your AWS source.
  - During recovery, you can restore Amazon DynamoDB tables to their original region. However, the restore operation fails if another table with the same name already exists in that region.
  - To restore a table using its original name, you must first manually delete the

existing table. Then, perform the recovery without specifying any prefix or suffix for the table name.

- **Restore Method Behavior and Priority**

  When a restore is initiated, Cohesity attempts recovery using the AWS Point-in-Time Recovery (PITR) method first. If PITR is unavailable, Cohesity falls back to the snapshot-based restore method.

  Capabilities and limitations differ between the two methods:

  - **PITR Restore**
    - By default, restores both Local Secondary Indexes (LSIs) and Global Secondary Indexes (GSIs) within the 35-day PITR window.
    - Restore is supported only to the original AWS account where the backups were created.
    - Cross-region and cross-account restores are not supported using PITR.
  - **Snapshot-Based Restore**
    - Supports restoring Local Secondary Indexes (LSIs) only when the Restore Local Secondary Indexes option is enabled during recovery. Otherwise, Cohesity uses the AWS ImportTable-based restore workflow, which does not restore LSIs.
    - Cross-region and cross-account restores are supported.
    - Used automatically when PITR is unavailable.
- **LSI Restore Considerations**
  - If Restore Local Secondary Indexes is enabled for snapshot-based restore, but the selected DynamoDB tables do not contain any Local Secondary Index (LSIs), Cohesity automatically switches to the ImportTable-based recovery method to reduce restore cost.
  - If you select multiple Amazon DynamoDB tables in a single recovery task, the configured restore settings are applied to all selected tables. Cohesity recommends creating separate recovery tasks for each table when the table sizes differ significantly. This allows you to optimize throughput allocation, restore performance, and cost per hour for individual tables.
  - If an LSI-enabled snapshot-based restore fails or is cancelled, Cohesity may not be able to automatically delete the partially created DynamoDB table. In such cases, the table might remain in the AWS account. If the table is provisioned with high write capacity units, AWS billing continues until the table is deleted.

    Recommendation: To avoid additional AWS charges, verify and manually delete any partially restored tables after a failed or cancelled recovery operation.
- **Global Tables**

    - Global tables are recovered as local tables.
    - After restore, you must manually reconfigure the table as a global table.
    - This behavior applies only to EVENTUAL CONSISTENCY Global Tables.
    - **Backup strategy**: Perform the backup from any one of the regions participating in the global table.
    - **Restore strategy**: Restore the table to the desired region, then manually configure the global table.
- **Index Restore Considerations**
    - Tables with Global Secondary Indexes (GSIs) may take longer to restore because indexes are recreated during recovery.
    - Restore of Local Secondary Indexes (LSIs) is not supported when using snapshot-based protection. This limitation applies only to snapshot restores and not to PITR restores. This is an AWS consideration for importing DynamoDB tables from Amazon S3.
- **Large Table Backup and Restore Considerations**

  For DynamoDB tables larger than 2.5 TB:
    - The default Max task DPUs per account must be increased in the AWS region where the AWS Glue job runs.
    - The increase is required on:
        - The AWS side (AWS Glue DPU quota)
        - The Cohesity side (cluster G-flag configuration)

  Contact Cohesity Support for assistance with increasing the DPU limits on both the AWS and Cohesity sides.

  For Large Restores (15 TB or Larger)
    - AWS does not natively support restores of 15 TB or larger.
    - Before attempting a restore of this scale, you must contact AWS Support to request an increase to the import size limit.
- **AWS Charges**

  For protected DynamoDB tables, the following AWS charges may apply:
    - **AWS Glue job costs**: For data transformation and movement during backup or recovery.
    - **Export table costs**: For exporting table data to S3.
    - **Egress costs**: For data transferred out of AWS.
- **Point-in-Time Recovery (PITR) Considerations**

    - Cohesity enables PITR for Amazon DynamoDB tables, allowing restoration to any specific point within the past 35 days.
    - PITR is available only when restoring to the original AWS account.
    - After unprotecting a DynamoDB table, you must manually disable PITR to avoid unnecessary retention and associated AWS costs.
    - Manually turning off PITR after a full backup breaks the backup chain. Subsequent incremental backups cannot run until a new full backup is completed.
- **DynamoDB Version Limitation**
    - Restoring a backup taken from a higher DynamoDB version to a lower version is not supported.
- **Recovery Cleanup on Failure or Cancellation**

  If the recovery workflow fails or is cancelled, Cohesity may be unable to automatically delete or roll back the DynamoDB table created during the recovery process, because the configured IAM role has insufficient deletion permissions. In such cases, you must verify the recovered table in AWS and manually delete any incomplete or unwanted tables.
- **Cross-Region and Cross-Account Restore Summary**

Restore Method	Same Region	Cross Region	Cross Account
PITR	Supported	Supported	Not Supported
Snapshot	Supported	Supported	Supported

Backup Amazon DynamoDB

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads in the specified AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **DynamoDB** tab to view all discovered databases.

For information on registering AWS, see Register AWS Source.

The **DynamoDB** tab displays the entity hierarchy DynamoDB, allowing navigation from the organization or account level down to individual tables.

Amazon DynamoDB Entity Hierarchy

The **DynamoDB** tab displays the entity hierarchy for DynamoDB, allowing you to navigate from the organization or account level down to individual tables. The **DynamoDB** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS Organization registered with Cohesity for Amazon DynamoDB protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The Organizational units within the AWS Organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain Amazon DynamoDB.
AWS Account	The AWS account registered with Cohesity for DynamoDB protection.
AWS Region	The AWS region where the DynamoDB resources reside.
DynamoDB Tables	The DynamoDB tables that can be selected for backup and recovery.

Backup Scope by Amazon DynamoDB Hierarchy

The Cohesity cluster allows you to back up Amazon DynamoDB at different levels of the entity hierarchy. You can choose to back up:

- **AWS Organization** – Backup all Amazon DynamoDB tables across all member accounts in the registered AWS organization.
- **Organizational Unit (OU)** – Backup all Amazon DynamoDB tables within the member accounts that belong to a specific organizational unit.
- **Member Account** – Backup all Amazon DynamoDB tables within a specific member account in the registered AWS organization.
- **AWS Account** – Backup all the AWS DynamoDB tables within the registered AWS account.
- **AWS Region** – Backup all AWS DynamoDB tables within a specific region of the registered account.
- **DynamoDB Tables** – Backup individual DynamoDB tables.

Backup Process

When you initiate a protection run for DynamoDB tables, Cohesity performs a DynamoDB API call to export the table data to a temporary AWS S3 staging bucket . This S3 bucket is created by Cohesity in the same AWS account and region where the DynamoDB table to be backed up resides.

Once the data is successfully exported to the S3 bucket, Cohesity ingests the exported data from the S3 bucket to the Cohesity cluster. Once the data is fully backed up to the Cohesity cluster, the data on the staging bucket is deleted to free up storage space.

During the first full backup, Cohesity also enables AWS Point-In-Time Recovery (PITR) for the DynamoDB table. This allows you to perform PITR restores, enabling recovery to a specific point in time within the past 35 days by default.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Supported Backup Types

Cohesity supports both full and Incremental backups for Amazon DynamoDB.

Although Cohesity does not support log backups for DynamoDB, it enables Amazon DynamoDB Point-In-Time Recovery (PITR) to provide point-in-time restore capability. This allows Cohesity to perform point-in-time recovery of DynamoDB tables back to the original location using AWS-managed PITR snapshots.

Protect Amazon DynamoDB Tables

In Cohesity, Amazon DynamoDB tables are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

To protect Amazon DynamoDB tables:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **AWS** > **DynamoDB**.
2. In the **New Protection** page, click **Add Objects** and select your registered source, or click **Register Source** to add a new one.
3. Select the DynamoDB tables using one of the following views or option:
    - **List View**: Use this option to select the DynamoDB tables for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select DynamoDB tables for protection at various hierarchy levels.

      Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new DynamoDB tables added to the AWS source.

        - If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-

protect icon:

- **AWS Organization**
- **Organizational Unit (OU)**
- **Member Account**
- **AWS Region**

- If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
  - **AWS Account**
  - **AWS Region**

You can exclude the Amazon DynamoDB tables from auto-protection within an account orregion by clicking the auto-protect icon ( ) next to the table. The tables excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the tables for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all tables with that tag for protection.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added DynamoDB tables with this tag.
  - You can exclude the Amazon DynamoDB tables from auto-protection within a tag by clicking the auto-protect icon ( ) next to the table. The DynamoDB tables excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude DynamoDB tables with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

4. Click **Continue.**
5. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
6. Since the data is ingested to the Cohesity cluster from the AWS S3 staging bucket, Ingest based policy is enabled by default. From the **Ingest Method** of Policy, create or select and existing policy
7. Select **More Options** to change or configure any of the Advanced Settings.
8. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon DynamoDB Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon DynamoDB tables to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for DynamoDB tables.

3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon DynamoDB tables.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon DynamoDB tables.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon DynamoDB tables.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.

- Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.

5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon DynamoDB

After protecting your Amazon DynamoDB tables in the Cohesity cluster, you can recover them to either their original region or a different region. Cohesity recovers the table as a new table within the location you choose. When you restore a table, all the tags that were present on the original table at the time of backup are automatically added to the restored table. Additionally, any new tags specified in the custom recovery options are also applied to the restored table.

You can also perform PITR when restoring to the original AWS account where the backups were created. Restoring to an alternate AWS account using PITR is not supported.

> **Note:** We recommend that you review the Plan and Prepare for Amazon DynamoDB Protection topic.

Recovery Process

Depending on the target location for recovery, Cohesity supports the following recovery methods for Amazon DynamoDB:

- Point-in-Time Recovery (PITR)
- Snapshot-based Recovery

Point-in-Time Recovery (PITR)

Cohesity uses Point-in-Time Recovery (PITR) when restoring a DynamoDB table within the same AWS account and region as the source table.

- By default, Cohesity attempts PITR method for recovering to the same region and account as the original DynamoDB table.
- PITR uses AWS-native APIs to restore the table directly from DynamoDB-managed backups.
- When PITR is successful, the table is restored with all configurations, including Local Secondary Indexes (LSIs) and Global Secondary Indexes (GSIs).
- PITR is available only when:
    - The restore point is within the configured PITR retention window (up to 35 days).
    - The source table exists and PITR was enabled.

- If PITR is not available or fails, Cohesity automatically falls back to snapshot-based recovery.

Snapshot-based Recovery

Cohesity uses this method when restoring the DynamoDB table to an alternate location (different AWS account or region).

Depending on whether Local Secondary Index (LSI) restore is required, Cohesity uses one of the following approaches:

- Snapshot Restore using ImportTable
- Snapshot Restore with LSI Support

Snapshot Restore using ImportTable

This is the default snapshot-based recovery method when LSI restore is not required.

In this method:

1. Cohesity exports the selected table data from the Cohesity cluster to an Amazon S3 staging bucket. This bucket is created by Cohesity in the same AWS region as the target DynamoDB table.
2. After the data export is complete, Amazon DynamoDB imports the data directly from the S3 staging bucket into the target table. Cohesity invokes the **AWS ImportTable** API to import the table from the S3 bucket.
3. After the import is complete, Cohesity deletes the exported data from the staging bucket to reduce storage costs.

Snapshot Restore with LSI Support

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

If you have enabled the Restore Local Secondary Indexes option on the Cohesity restore workflow, Cohesity uses an alternative restore method to support LSI restoration.

In this method:

1. Cohesity creates a new DynamoDB table using the original table schema, including LSI definitions.
2. Data is read from the Cohesity snapshot and written to DynamoDB using the BatchWriteItem API.
3. Data ingestion is performed in parallel using multiple workers to optimize throughput.
4. After data ingestion completes:

- Global Secondary Indexes (GSIs) are created and backfilled sequentially using AWS APIs.
- Table settings are reverted to the original configuration.

Recover Amazon DynamoDB Tables

To recover Amazon DynamoDB:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **DynamoDB**.
3. Search for the table by entering characters of the table name, or by specifying the wildcard character *.
4. From the list of search results, select the table to recover.
   - By default, the latest snapshot of the table is selected for recovery.
   - To recover from a different snapshot, hover over the snapshot and click the **Edit** icon. The **Edit Recovery Point for <*table*_name>** page appears.
   - By default, snapshots are displayed in the **List** mode. You can switch to the Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   To select a snapshot from the **List** mode:

   a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*table_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

   b. Select the snapshots to be recovered.

   c. By default, the snapshot is recovered from the local Cohesity cluster.

   d. Click **Select Recovery Point** and then go to step 5.

   To select a snapshot from the **Timeline** mode:

   a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

   b. Select the snapshot to be recovered by clicking the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

   c. Click **Select Recovery Point** and then go to step 5.
5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on where you want to restore the table and then provide the required information:

- **Original Location**: Restore the DynamoDB table within the same region and account as the original DynamoDB table.
- **New Location:** Restore the DynamoDB table to a new region within the same AWS account or to a different AWS account.

  Based on the registration level of the AWS source, select the recovery destination hierarchy:

  - If the AWS source is registered at the organization level:
    - **Source**: Select the AWS organization that contains the destination Amazon DynamoDB. You can recover data only within the same AWS Organization as the source.
    - **AWS Account**: Select the member account that contains the destination Amazon DynamoDB.
    - **Region**: Select the destination AWS Region.
  - If the AWS source is registered at the account level:
    - **Source**: Select the AWS account that contains the destination Amazon DynamoDB.
    - Region: Select the destination AWS Region.

> **Note:**
> You cannot restore a table if another table with the same name already exists in that region—the restore operation will fail.
>
> To restore a table using its original name, you must first manually delete the existing table. Then, perform the recovery without specifying any prefix or suffix for the table name.

7. In the **Recovery Options** section, configure the following options:

Options	Description

Encryption	Select the encryption type for the table to be recovered:
	• **DynamoDB**: The AWS KMS key is owned and managed by DynamoDB. This option is selected by default.
	• **AWS Managed**: The AWS KMS key is stored in your account and managed by AWS Key Management Service.
	• **Self Managed**: The AWS KMS key is stored in your account and managed by you. Provide the following information:
	  • **KMS Key Location**: From the drop-down, select whether you want to use the KMS key from the Same AWS Account where you are recovering to or from a Different AWS Account.
	  • **KMS Key**: If you are using the KMS key from the Same AWS Account, then from the drop-down, select the KMS Key. If you are using the KMS key from a Different AWS Account, then enter the KMS Key ARN in the following format: `arn:<partition>:kms:<region>:<accountid>:key/<key-id>`.

Options	Description
Restore Local Secondary Indexes **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.	Enable this option to restore all associated Local Secondary Index (LSIs) along with the table. When enabled: • LSIs are restored along with the base table. • Additional write capacity is consumed during the restore process. When disabled: • LSIs are not restored. • Only the base table data is recovered. **Note:** • For point-in-time recovery, Local Secondary Index (LSIs) are automatically restored as part of the recovery process. No additional configuration is required. • The Restore Local Secondary Indexes option applies only to snapshot-based recovery, where you must explicitly enable it to restore LSIs. • If Restore Local Secondary Indexes is enabled, but the selected DynamoDB tables do not contain any Local Secondary Index (LSIs), Cohesity automatically switches to the ImportTable-based recovery method to reduce restore cost. **Choose Restore Speed and Cost** Select the restore performance mode: • **Balanced (Recommended)**: (Default) Automatically configures **Provisioned Write Capacity Units** and **GSI Write Capacity Multiplier** to provide balanced throughput for moderate recovery speed and controlled costs per hour or unit time. • **Fastest Restore**: Automatically configures **Provisioned Write Capacity Units** and **GSI Write Capacity Multiplier** to provide maximum throughput for fastest restore at the highest cost per hour. • **Lowest Cost**: Automatically configures **Provisioned Write Capacity Units** and **GSI Write Capacity Multiplier** to provide minimal throughput for the longest restore at the lowest cost per hour. • **Custom**: Manually configure **Provisioned Write Capacity Units** and **GSI Write Capacity Multiplier**. **View Per-Table Allocation** Click **View per-table allocation** to review how write capacity is distributed.

Options	Description
	The **Per-table Write Capacity Allocation** window displays: • **Table name** • **Table size** • **Base WCU** allocated to the table • **GSI WCU** (write capacity allocated for Global Secondary Indexes) This helps you understand how the selected restore option distributes throughput across tables and indexes. **Note:** • If you select multiple Amazon DynamoDB tables in a single recovery task, the configured restore settings are applied to all selected tables. Cohesity recommends creating separate recovery tasks for each table when the table sizes differ significantly. This allows you to optimize throughput allocation, restore performance, and cost per hour for individual tables. • The restore operation uses **batch writes**, which consume provisioned write capacity units. • The configured values for **Provisioned Write Capacity Units (WCU)** and **GSI Write Capacity Multiplier** are applied to: The base table and all associated indexes, including LSIs and GSIs. • Increasing write capacity speeds up the restore process and increases the per-hour AWS cost. • Lower write capacity reduces cost per hour and increases restore duration.
Custom AWS Resource Tags	Provide tags (key-value pairs) to attach to the recovered table. You can add multiple key-value pairs by clicking the **Add i**con.
Rename	Provide a prefix and/or suffix to the name of the table being recovered.
Task Name	Update the default name of the recovery task.

8. Click **Recover**.

> **Note:** If the recovery workflow fails or is cancelled, Cohesity may be unable to automatically delete or roll back the DynamoDB table created during the recovery process, because the configured IAM role has insufficient deletion permissions. In such cases, you must verify the recovered table in AWS and manually delete any incomplete or unwanted tables.

Amazon Redshift Cluster

Amazon Redshift cluster is a fully managed, petabyte-scale cloud data warehouse service offered by AWS to efficiently analyze large datasets using standard SQL and existing business intelligence (BI) tools. Cohesity natively integrates with AWS to provide solution for the backup, recovery, and data management of Amazon Redshift databases in your AWS source.

> **Important:** The backup and recovery of Amazon Redshift cluster is supported only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster). This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Plan and Prepare for Amazon Redshift Cluster

Before protecting Amazon Redshift Cluster with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect AmazonRedshift Cluster with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

AWS S3 Staging Bucket Requirement

Cohesity uses an Amazon S3 staging bucket as temporary storage during the backup and recovery of Amazon Redshift databases. You can either allow Cohesity to automatically create the staging bucket when you initiate a backup or recovery task, or you can manually create a custom bucket in your AWS account before starting the process.

You must use the same bucket for both backup and recovery. If you choose to manually create the bucket, ensure to create the AWS S3 bucket of type, S3 Standard in the same AWS region as the Redshift database selected for backup and restore.

You can specify only one custom bucket name per Protection Group. Because an S3 bucket belongs to a single AWS region, a custom bucket is supported only when all Redshift

databases in the Protection Group are in the same region. If your Protection Group spans multiple regions, use the **Auto Create** option instead.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Ports Used for Communication

For information on port numbers required, seeAWS Cloud Database Workloads.

Considerations for Amazon Redshift Data Protection

Review and understand the following limitations and known issues before you protect Amazon Redshift databases:

- **Redshift Support Scope**: Cohesity supports protection only for provisioned Amazon Redshift clusters. Amazon Redshift Serverless is not supported.
- **Restore Scope**: does not support restoring Amazon Redshift databases to a different region or account. Restores can only be performed to the original account within the same region.
- **FIPS Certification**: Backup and recovery operations for Redshift databases using Cohesity are not FIPS-certified.
- Database Coverage: Only user databases are supported for protection. System databases within Redshift are not supported for backup or recovery.
- **Cluster State Requirement**: The Redshift cluster must be in an available state for Cohesity to initiate a backup. Backups cannot be initiated if the cluster is offline or unavailable.
- **Empty Tables**: It is not recommended to back up databases that contain empty tables.
- **Table Format**: Tables with names that use double quotes or are case-sensitive will not be restored.
- **Crash-Consistent Backup**: Cohesity provides crash-consistent protection for Amazon Redshift databases. Before initiating a backup, review your application's activity and requirements to determine if it is necessary to perform a checkpoint or quiesce the application.
- **Tag Restore**: Cohesity does not support restoring tags for Amazon Redshift clusters since the protection is at database level.
- **Point-in-Time Restore (PITR)**: Cohesity does not support PITR of Amazon Redshift cluster.

- **AWS S3 Staging Bucket**: You must use the same AWS S3 staging bucket for both backup and recovery. Recovery may fail if you do not use the same staging bucket that was used during the backup operation.
- **Version Downgrade Not Supported**: Restoring a backup taken from a higher Redshift database version to a lower version is not supported.
- **External Database Support Limitation**: External databases in Amazon Redshift created through AWS Glue Data Catalog registration may appear in the Cohesity cluster entity hierarchy; however, they are not currently supported for backup and recovery.
- **IAM Auto-Created User Ownership**: Recovery of Redshift cluster database to an alternate cluster can fail if a table is owned by an IAM auto-created database user (for example a user created through Redshift Query Editor v2 using temporary IAM credentials). During recovery, the workflow attempts to restore table ownership using an `ALTER TABLE ... OWNER TO` statement. This statement fails if the referenced IAM user does not already exist on the target cluster, because IAM auto-created users are not restored as part of the recovery process.

  To resolve this issue, log in to the target cluster using the same IAM identity (example, using Redshift Query Editor). This automatically creates the corresponding database user on the target cluster. After the user is created, rerun the recovery task. The subsequent recovery task succeeds.
- **Custom Staging Bucket and Multi-Region Considerations**: When you use the **Specify Bucket to Use** option, you can specify only one bucket name per Protection Group. Because an S3 bucket resides in a single AWS region, this option is not supported in scenarios where Redshift databases from multiple regions or accounts might be protected together, such as:
  - Auto-protect at the AWS Account, OU, or Organization level.
  - Tag-based auto-protection.
  - A Protection Group that includes Redshift databases from multiple AWS regions.
  - AWS Organization registration spanning multiple accounts with Redshift databases.
  - For these scenarios, use the **Auto Create** option to allow Cohesity to automatically create the appropriate staging bucket per region.

Attach CohesitySourceRegistrationRole to Amazon Redshift Cluster

After completing source registration by running the CloudFormation Template (CFT), the following IAM roles and profiles are created:

- **CohesityInstanceRole**
- **CohesityInstanceProfile**
- **CohesitySourceRegistrationRole**

The **CohesitySourceRegistrationRole** includes the permissions defined in the CFT that allow Cohesity to discover AWS resources within the account hierarchy in the product UI.

To perform backup and recovery operations for Amazon Redshift, you must attach the **CohesitySourceRegistrationRole** to the Redshift cluster you want to protect and set it as the default IAM role for the cluster.

Attach CohesitySourceRegistrationRole

Perform the following steps to attach the permissions to **CohesitySourceRegistrationRole** to the Redshift cluster:

1. In the AWS Management Console, go to the Amazon Redshift console.
2. In the navigation pane, choose **Clusters**.
3. Select the Redshift cluster you want to protect.
4. Go to the **Properties** tab.
5. Under **Permissions,** locate the IAM roles section.
6. Click **Associate IAM roles**.
7. Select **CohesitySourceRegistrationRole** from the list.
8. Click **Associate IAM roles**.

Set as Default Role

After attaching the role:

1. In the same IAM roles section, locate **CohesitySourceRegistrationRole**.
2. Click **Set as default**.

Backup Amazon Redshift Cluster

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads in the specified AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS account source, each workload type is displayed on its own tab. Select the **Redshift Cluster** tab to view all discovered Redshift clusters and databases. For information on registering AWS, see Register AWS Source.

The **Redshift Cluster** tab displays the entity hierarchy Redshift, allowing navigation from the organization or account level down to individual databases.

Amazon Redshift Cluster Entity Hierarchy

The Amazon Redshift entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with Cohesity for Redshift cluster protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The organizational units within the AWS organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain Redshift cluster.
AWS Account	The AWS account registered with Cohesity for protecting Amazon Redshift databases. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the Redshift resources reside.
Redshift Cluster	A Redshift cluster hosts one or more Redshift databases.

Structure	Description
Redshift Database	The individual Redshift databases hosted within the Redshift cluster.

Backup Scope by Amazon Redshift Cluster Hierarchy

The Cohesity cluster allows you to back up Amazon Redshift at different levels of the entity hierarchy. You can choose to back up:

- **AWS Organization** – Protect all Amazon Redshift cluster databases across all member accounts in the registered AWS Organization.
- **Organizational Unit (OU)** – Backup all Amazon Redshift cluster databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Backup all Amazon Redshift cluster databases within a specific member account in the registered AWS organization.
- **AWS Account** – Backup all the AWS Redshift cluster databases within the registered AWS account.
- **AWS Region** – Backup all the AWS Redshift cluster databases within a specific region of the registered account.
- **Redshift Cluster** – Backup all databases within a specific Redshift cluster.
- **Redshift Database** – Backup individual Redshift cluster database.

Backup Process

When you initiate a protection run in the Cohesity cluster to back up Amazon Redshift databases, Cohesity first exports the selected Redshift data to an Amazon S3 staging bucket.

Once the data is successfully exported to the S3 bucket, Cohesity ingests the staged data from the S3 bucket into the Cohesity cluster. Once the ingestion is successful and the data is fully backed up to the Cohesity cluster, Cohesity automatically deletes the exported data from the staging bucket to free up storage and minimize S3 costs.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Supported Backup Types

Cohesity supports only full backups for Redshift cluster databases. Incremental backups are not supported.

Protect Amazon Redshift Cluster

In Cohesity, Amazon Redshift databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

Prerequisites

Before initiating a backup of Amazon Redshift cluster databases, an Amazon S3 staging bucket must be available in the AWS account where the Redshift cluster resources reside. This bucket is used to temporarily store exported data during the backup process. You can manually create or select an existing S3 bucket in your AWS account, or allow Cohesityto automatically create one when the backup process begins. When creating a Protection Group, you can choose either option based on your preference.

AWS Source Registered at the Account Level

When creating a Protection Group, you can manually specify an S3 bucket only if all selected databases belong to the same AWS account and region.

If all selected databases are in the same AWS account and region:

You can manually specify a newly created or existing S3 bucket, or allow Cohesity to automatically create one. If you are manually specifying an S3 bucket, ensure the bucket is in the same account and region as the selected database

If the selected databases span multiple regions:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

AWS Source Registered at the Organization Level

When protecting databases across member accounts, Cohesity supports manual bucket selection only when all selected databases for protection belong to the same member account and region.

If all selected databases are in the same AWS member account and region:

You can manually specify a newly created or existing S3 bucket, or allow Cohesity to automatically create one. If you are manually specifying an S3 bucket, ensure the bucket is in the same account and region as the selected database

If the selected databases span multiple AWS member accounts and region:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Protect Amazon Redshift Database

To protect Amazon Redshift databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **AWS** > **Redshift Cluster**.
2. In the **New Protection** page, click **Add Objects** and select your registered source, or click **Register Source** to add a new one.
3. Select the Redshift databases using one of the following views or option:
   - **List View**: Use this option to select the Redshift cluster database databases for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select Redshift cluster database databases for protection at various hierarchy levels.

     Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon Redshift cluster databases.

     If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:
     - **AWS Organization**
     - **Organizational Unit (OU)**
     - **Member Account level**
     - **AWS Region**
     - **Availability Zone**
     - **Namespace**

     If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:

- **AWS Account**
- **AWS Region**
- **Availability Zone**
- **Namespace**

You can exclude the Redshift databases from auto-protection within a hierarchy level by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the databases for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all databases that match the selected tag for protection and backup operations.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added databases with this tag.
  - You can exclude the databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude databases with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.

> **Note:** If you enable auto-protect, the Protection Group may include Redshift databases from multiple AWS regions. In this case, you must use the **Auto Create** option for the staging bucket. Using **Specify Bucket to Use** with a multi-region selection will cause the backup to fail.

4. Click **Continue.**
5. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
6. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary AWS S3 staging bucket in the same region as your Redshift databases selected for backup. The bucket follows the naming format: `cohe-<source-id>-<region-name>`.

Where:

- `source-id` is the AWS account ID.
- `region-name` is the AWS region where the Redshift database reside.

- **Specify Bucket to Use**: Enter the name of an existing AWS S3 bucket in the same region and account as the Redshift databases being protected. If the Protection Group includes databases from multiple AWS regions, do not use this option, the backup will fail. Use **Auto Create** instead.

After each protection run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations.

> **Note:** The **Specify Bucket to Use** option is not supported when the Protection Group includes Redshift databases from multiple AWS regions. This includes Protection Groups using auto-protect at the AWS Account, OU, or Organization level, tag-based auto-protection, or manual selection of databases across regions. In these cases, select **Auto Create**. If you select **Specify Bucket to Use** in these scenarios, the backup will fail.

7. Create or select an existing policy from the **Policy** drop-down.
8. Select **More Options** to change or configure any of the Advanced Settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon Redshift Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon Redshift databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Redshift database.

3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Redshift database.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Redshift database.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Redshift database.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.

- Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.

5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon Redshift Cluster

Cohesity does not support restoring Amazon Redshift databases to a different region or account. Restores can only be performed to the original account within the same region.

> **Note:** We recommend that you review the Considerations for Amazon Redshift Data Protection.

Recovery Process

When you initiate a recovery, Cohesity exports the selected Redshift data from the Cohesity cluster back into the Amazon S3 staging bucket that was used during the original backup. After the data is successfully exported to the staging bucket, Amazon Redshift reads the exported files from the bucket and imports them into the specified Redshift target location. Once the import completes, Cohesity automatically deletes the exported files from the staging bucket to free up storage and minimize S3 costs.

Recover Redshift Database

To recover Amazon Redshift database:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **Redshift Cluster.**
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
   - By default, the latest snapshot of the database is selected for recovery.
   - To recover from a different snapshot, hover over the snapshot and click the **Edit** icon. The **Edit Recovery Point for <*database_name*>** page appears.
   - By default, snapshots are displayed in the List mode. You can switch to the Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   To select a snapshot from the **List** mode:

   a. Select the **List** tab from the top-right corner of the **Edit recovery point for <*database_name*>** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page

displays the available snapshots in the specified range.

b. Select the snapshots to be recovered.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.

    1. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:

        - **Original Location**: Restore the database to the original Amazon Redshift cluster.

            > **Note:** Restoring to the original location will overwrite any existing database with the same name.

            Type **Yes** to confirm.

        - **New Location**: Restore the database to an Amazon Redshift cluster within the same region of the same AWS account. Restoring to different AWS account or to a different region within the same account is not supported.

            Based on the registration level of the AWS source, select the recovery destination hierarchy:

            - If the AWS source is registered at the organization level:
                - **Source**: Select the AWS organization that contains the destination Amazon Redshift cluster. You can recover data only within the same AWS Organization as the source.
                - **AWS Account**: Select the AWS member account that contains the destination Amazon Redshift serverless namespace.

- **Cluster**: Select the Amazon Redshift cluster to which you want to restore the database.

- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account that contains the destination Amazon Redshift cluster.
  - **Cluster**: Select the Amazon Redshift cluster to which you want to restore the database.

To overwrite any existing database with the same name, enable **Overwrite existing database(s)** and type **Yes** to confirm.

> **Note:** Restoring to a new location also replaces any existing database with the same name in the selected cluster.

6. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**:

     Select this option if you chose **Auto Create** for the staging bucket during the backup process.
   - **Specify Bucket to Use**:

     Select this option if you created a staging bucket during the backup process. Ensure to specify the **same S3 bucket** that was used for the backup.

     > **Important:** Recovery may fail if you do not use the same staging bucket that was used during the backup operation.

   Cohesity deletes only the temporary data stored in the staging bucket during each recovery operation. The bucket itself is retained and reused for future recovery operations.
7. Click the **Edit** icon of the **Rename** field and enter a **New Name** for the database if you want to recover with a new name.
8. In the **Task Name** field, update the default name of the recovery task.
9. Click **Recover**.

Amazon Redshift Serverless

Amazon Redshift Serverless is an on-demand, auto-scaling configuration of the Amazon Redshift data warehouse that allows you to run analytics without managing data warehouse infrastructure. Cohesity integrates with AWS to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of databases within the Amazon Redshift Serverless namespace of your AWS source.

> **Important:** The backup and recovery of Amazon Redshift cluster is supported only on Cohesity Cloud Edition on AWS (formerly known as AWS NGCE cluster). This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Plan and Prepare for Amazon Redshift Serverless

Before protecting Amazon Redshift Serverless with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect AmazonRedshift Serverless with the Cohesity cluster, you must register your AWS as a source in the Cohesity cluster and perform backup and recovery operations. For information on registering AWS, see Register AWS Source.

Required Permissions for Amazon Redshift Serverless Data Protection

For information on the permissions required to protect Redshift Serverless namespace database, see Minimum Permissions for Individual Workloads and Workflows.

AWS S3 Staging Bucket Requirement

Cohesity uses an Amazon S3 staging bucket as temporary storage during the backup and recovery of databases within a Redshift Serverless namespace. You can either allow Cohesity to automatically create the staging bucket when you initiate a backup or recovery task, or you can manually create the staging bucket in your AWS account before starting the process.

You must use the same bucket for both backup and recovery. If you choose to manually create the bucket, ensure to create the AWS S3 bucket of type, S3 Standard in the same AWS region as the Redshift Serverless namespace database selected for backup and restore.

You can specify only one custom bucket name per Protection Group. Because an S3 bucket belongs to a single AWS region, a custom bucket is supported only when all Redshift databases in the Protection Group are in the same region. If your Protection Group spans multiple regions, use the **Auto Create** option instead.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Ports Used for Communication

For information on port numbers required, seeAWS Cloud Database Workloads.

Admin Password Secret ARN Requirement

For enabling backup and restore operations in Amazon Redshift Serverless, you must configure an admin password secret ARN on the namespace.

Backup and recovery workflows require secure access to administrative credentials. Configuring an admin password secret ARN ensures that these credentials are securely stored and retrieved using AWS Secrets Manager, enabling successful data protection operations.

This requirement applies to both:

- Newly created namespaces
- Existing namespaces that do not already have an associated secret ARN

You can update existing namespaces using the AWS CLI. For more information, see AWS CLI update-namespace command.

> **Note:** Use the `--manage-admin-password` option to enable management of the admin password via Secrets Manager.

Enabling the admin password secret ARN rotates the existing admin password. As a result, any applications or services using the previous password will fail to connect until they are updated.

Ensure that all dependent applications retrieve credentials dynamically from AWS Secrets Manager.

Considerations for Amazon Redshift Serverless Data Protection

Review and understand the following limitations and known issues before you protect Amazon Redshift databases:

- **Restore Scope**: Cohesity does not support restoring Amazon Redshift databases to a different region or account. Restores can only be performed to the original account within the same region.
- **FIPS Certification**: Backup and recovery operations for Redshift databases using Cohesity are FIPS-certified.
- **Database Coverage**: Only user databases are supported for protection. System databases within Redshift are not supported for backup or recovery.
- **Namespace and workgroup State Requirement**: The Amazon Redshift Serverless Namespace and workgroup must be in an available state for Cohesity to initiate a backup. Backups cannot be initiated if the Namespace and workgroup are offline or unavailable.
- **Empty Tables**: It is not recommended to back up databases that contain empty tables.

- **Table Format**: Tables with names that use double quotes or are case-sensitive will not be restored.
- **Crash-Consistent Backup**: Cohesity provides crash-consistent protection for databases. Before initiating a backup, review your application's activity and requirements to determine if it is necessary to perform a checkpoint or quiesce the application.
- **Tag Restore**: Cohesity does not support restoring tags for namespace since the protection is at database level.
- **AWS S3 Staging Bucket**: You must use the same AWS S3 staging bucket for both backup and recovery. Recovery may fail if you do not use the same staging bucket that was used during the backup operation.
- **Custom Staging Bucket and Multi-Region Considerations**: When you use the **Specify Bucket to Use** option, you can specify only one bucket name per Protection Group. Because an S3 bucket resides in a single AWS region, this option is not supported in scenarios where Redshift Serverless namespace databases from multiple regions or accounts might be protected together, such as:
  - Auto-protect at the AWS Account, OU, or Organization level.
  - Tag-based auto-protection.
  - A Protection Group that includes Redshift databases from multiple AWS regions.
  - AWS Organization registration spanning multiple accounts with Redshift Serverless namespace databases.
  - For these scenarios, use the **Auto Create** option to allow Cohesity to automatically create the appropriate staging bucket per region.

Attach CohesitySourceRegistrationRole to Amazon Redshift Serverless

After completing source registration by running the CloudFormation Template (CFT), the following IAM roles and profiles are created:

- **CohesityInstanceRole**
- **CohesityInstanceProfile**
- **CohesitySourceRegistrationRole**
- **CohesityDelegatedAdminRole** (if AWS is registered at organization-level)

The **CohesitySourceRegistrationRole** includes the permissions defined in the CFT that allow Cohesity to discover AWS resources within the account hierarchy in the product UI.

To perform backup and recovery operations for Amazon Redshift Serverless, you must attach the **CohesitySourceRegistrationRole** to the Amazon Redshift Serverless you want to protect and set it as the default IAM role for the Amazon Redshift Serverless.

Attach CohesitySourceRegistrationRole

Perform the following steps to attach the permissions to **CohesitySourceRegistrationRole** to the Amazon Redshift Serverless:

1. In the AWS Management Console, go to the Amazon Redshift console.
2. In the navigation pane, choose **Redshift**.
3. Select the Redshift Serverless namespace you want to protect.
4. Go to the **Properties** tab.
5. Under **Permissions,** locate the IAM roles section.
6. Click **Associate IAM roles**.
7. Select **CohesitySourceRegistrationRole** from the list.
8. Click **Associate IAM roles**.

Set as Default Role

After attaching the role:

1. In the same IAM roles section, locate **CohesitySourceRegistrationRole**.
2. Click **Set as default**.

Backup Amazon Redshift Serverless

Depending on the services you selected while registering the source, the Cohesity cluster automatically discovers all supported workloads associated with the registered AWS source. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered AWS source, each workload type is displayed on its own tab. Select the **Redshift Serverless** tab to view all discovered databases within the Redshift Serverless namespaces of your AWS source. For information on registering AWS, see Register AWS Source.

Based on the registration level of your AWS source, the **Redshift Serverless** tab displays the entity hierarchy for Redshift Serverless, allowing navigation from the organization or account level down to individual databases.

Amazon Redshift Serverless Entity Hierarchy

The Amazon Redshift Serverless entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
AWS Organization	The AWS organization registered with Cohesity for Redshift Serverless protection. Applicable when the AWS source is registered at the organization level.
Organizational Unit (OU)	The organizational units within the AWS organization that contain member accounts.
Member Account	The AWS member accounts within the organization that contain Redshift Serverless namespaces.
AWS Account	The AWS account registered with Cohesity for protecting Amazon Redshift databases. Applicable when the AWS source is registered at the account level.
AWS Region	The AWS region where the Redshift Serverless resources reside.
Amazon Redshift Serverless Namespace	An Amazon Redshift Serverless namespace hosts one or more databases. Each namespace appears as an entity within Cohesity for backup and recovery operations.
Amazon Redshift Serverless Namespace Database	The individual databases hosted within the Amazon Redshift Serverless namespace.

Backup Scope by Amazon Redshift Serverless Hierarchy

The Cohesity cluster allows you to back up databases within the Amazon Redshift Serverless namespace at different levels of the entity hierarchy. The available hierarchy levels depend on whether the AWS source is registered at the organization level or account level.

You can choose to back up:

- **AWS Organization** – Protect all Amazon Redshift Serverless namespace databases across all member accounts in the registered AWS Organization.
- **Organizational Unit (OU)** – Protect all Amazon Redshift Serverless namespace databases within the member accounts that belong to a specific organizational unit.
- **Member Account** – Protect all Amazon Redshift Serverless namespace databases within a specific member account in the registered AWS organization.
- **AWS Account** – Protect all the AWS Redshift Serverless namespace databases within the registered AWS account.
- **AWS Region** – Protect all the AWS Redshift Serverless namespace databases within a specific region of the registered account.
- **Amazon Redshift Serverless Namespace** – Protect all databases within a specific Amazon Redshift Serverless namespace.
- **Redshift Serverless Databases** – Protect individual database within an Amazon Redshift Serverless namespace.

Backup Process

When you initiate a protection run in the Cohesity cluster to back up Amazon Redshift Serverless namespace databases, Cohesity first exports the selected Redshift data to an Amazon S3 staging bucket.

Once the data is successfully exported to the S3 bucket, Cohesity ingests the staged data from the S3 bucket into the Cohesity cluster. Once the ingestion is successful and the data is fully backed up to the Cohesity cluster, Cohesity automatically deletes the exported data from the staging bucket to free up storage and minimize S3 costs.

Supported Backup Types

Cohesity supports only full backups for Redshift Serverless namespace databases. Incremental backups are not supported.

Protect Amazon Redshift Serverless

In Cohesity, Amazon Redshift Serverless namespace databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

Prerequisites

Before initiating a backup of Amazon Redshift Serverless namespace databases, an Amazon S3 staging bucket must be available in the AWS account where the Redshift Serverless resources reside. This bucket is used to temporarily store exported data during the backup process. You can manually create or select an existing S3 bucket in your AWS account, or

allow Cohesityto automatically create one when the backup process begins. When creating a Protection Group, you can choose either option based on your preference.

AWS Source Registered at the Account Level

When creating a Protection Group, you can manually specify an S3 bucket only if all selected databases belong to the same AWS account and region.

If all selected databases are in the same AWS account and region:

You can manually specify a new or existing S3 bucket, or allow Cohesity to automatically create one. If you are manually specifying an S3 bucket, ensure the bucket is in the same AWS account and region as the selected databases.

If the selected databases span multiple regions:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

AWS Source Registered at the Organization Level

When protecting databases across member accounts, Cohesity supports manual bucket selection only when all selected databases for protection belong to the same member account and region.

If all selected databases are in the same AWS member account and region:

You can manually specify a new or existing S3 bucket, or allow Cohesity to automatically create one. If you are manually specifying an S3 bucket, ensure the bucket is in the same AWS member account and region as the selected databases.

If the selected databases span multiple AWS member accounts and region:

You must allow allow Cohesity to automatically create the staging bucket. If you manually specify a newly created or existing S3 bucket, the backup will fail.

> **Note:** Cohesity supports adding custom tags to S3 staging buckets when the auto-create staging bucket option is selected. Tags are applied Cohesity cluster-wide across all AWS registrations on that Cohesity cluster. To configure tags for S3 staging bucket, contact Cohesity Support.

Protect Amazon Redshift Serverless Namespace Database

To protect Amazon Redshift databases:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **AWS** > **Redshift Serverless**.
3. In the **New Protection** page, click **Add Objects** and select your registered source, or click **Register Source** to add a new one.
4. Select the Redshift Serverless namespace databases using one of the following views or option:
    - **List View**: Use this option to select the Redshift Serverless namespace databases for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select Redshift Serverless namespace databases for protection at various hierarchy levels.

      Optionally, you can configure auto-protection at different hierarchy levels based on the registration level of your AWS source to automatically protect new databases added to the Amazon Redshift Serverless namespaces.

      If the AWS source is registered at the organization level, you can configure auto-protection at the following levels by clicking the auto-protect icon:

        - **AWS Organization**
        - **Organizational Unit (OU)**
        - **Member Account level**
        - **AWS Region**
        - **Availability Zone**
        - **Namespace**

      If the AWS source is registered at the account level, you can configure auto-protection at the following levels by clicking the auto-protect icon:

        - **AWS Account**
        - **AWS Region**
        - **Availability Zone**
        - **Namespace**

      You can exclude the Redshift serverless namespace databases from auto-protection within a hierarchy level by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
    - **Tags**: Use this option to select the databases for protection based on namespace-level tags. Select the checkbox of a tag, and then select **Select All**

Child Objects to include all databases within namespaces that match the selected tag for protection and backup operations.

Optionally:

- You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added databases with this tag.
- You can exclude the databases from auto-protection within a tag by clicking the auto-protect icon ( ) next to the database. The databases excluded from auto-protection are displayed with the exclude ( ) icon.
- You can exclude databases with a specific tag from auto-protection by clicking the exclude icon ( ). The tag excluded from auto-protection are displayed with the exclude icon.

> **Note:** If you enable auto-protect, the Protection Group may include Redshift Serverless namespace databases from multiple AWS regions. In this case, you must use the **Auto Create** option for the staging bucket. Using **Specify Bucket to Use** with a multi-region selection will cause the backup to fail.

5. Click **Continue.**
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary AWS S3 staging bucket in the same region as your Redshift databases selected for backup. The bucket follows the naming format: `cohe-<source-id>-<region-name>`.

     Where:

     - `source-id` is the AWS account ID.
     - `region-name` is the AWS region where the Redshift database reside.
   - **Specify Bucket to Use**: Enter the name of an existing AWS S3 bucket in the same region and account as the Redshift Serverless namespace databases being protected. If the Protection Group includes databases from multiple AWS regions, do not use this option, the backup will fail. Use **Auto Create** instead.

> **Note:** The **Specify Bucket to Use** option is not supported when the Protection Group includes Redshift Serverless namespace databases from multiple AWS regions. This includes Protection Groups using auto-protect at the AWS Account, OU, or Organization level, tag-based auto-protection, or manual selection of databases across regions. In these cases, select **Auto Create**. If you select **Specify Bucket to Use** in these scenarios, the backup will fail.

If the selected databases span multiple AWS accounts or regions, the option to manually specify an S3 bucket isn't available. You can only select the **Auto Create** option for Cohesity to automatically create the staging bucket. For more information, see Prerequisites.

After each protection run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations.

8. Create or select an existing policy from the **Policy** drop-down.
9. Select **More Options** to change or configure any of the Advanced Settings.
10. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: This option is not applicable. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Amazon Redshift Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Amazon Redshift Serverless namespace databases to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Redshift Serverless namespace database.

3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Redshift Serverless namespace database.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Redshift Serverless namespace database.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Amazon Redshift Serverless namespace database.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.

- Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.

5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Amazon Redshift Serverless

Cohesity does not support restoring Amazon Redshift Serverless namespace databases to a different region or account. Restores can only be performed to the original account within the same region.

> **Note:** We recommend that you review the Considerations for Amazon Redshift Serverless Data Protection.

Recovery Process

When you initiate a recovery, Cohesity exports the selected Redshift Serverless data from the Cohesity cluster back into the Amazon S3 staging bucket that was used during the original backup. After the data is successfully exported to the staging bucket, Amazon Redshift Serverless reads the exported files from the bucket and imports them into the specified Redshift Serverless target location. Once the import completes, Cohesity automatically deletes the exported files from the staging bucket to free up storage and minimize S3 costs.

Recover Redshift Serverless Namespace Database

To recover Amazon Redshift Serverless namespace database:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **AWS** > **Redshift Serverless**.
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
   - By default, the latest snapshot of the database is selected for recovery.
   - To recover from a different snapshot, hover over the snapshot and click the **Edit** icon. The **Edit Recovery Point for <*database_name*>** page appears.
   - By default, snapshots are displayed in the List mode. You can switch to the Timeline mode by clicking the **Timeline** tab in the top-right corner of the page.

   To select a snapshot from the **List** mode:

a. Select the **List** tab from the top-right corner of the **Edit recovery point for *<database_name>*** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

b. Select the snapshots to be recovered.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on where you want to restore the database and then provide the required information:

- **Original Location**: Restore the database to the original Amazon Redshift Serverless.

  > **Note:** Restoring to the original location will overwrite any existing database with the same name.

  Type **Yes** to confirm.

- **New Location**: Restore the database to an Amazon Redshift Serverless within the same region of the same AWS account. Restoring to different AWS account or to a different region within the same account is not supported.

  Based on the registration level of the AWS source, select the recovery destination hierarchy:

  - If the AWS source is registered at the organization level:
    - **Source**: Select the AWS organization that contains the destination Amazon Redshift serverless namespace. You can recover data only within the same AWS Organization as the source.
    - **AWS Account**: Select the AWS member account that contains the destination Amazon Redshift serverless namespace.

- **Namespace**: Select the Amazon Redshift Serverless namespace to which you want to restore the database.

- If the AWS source is registered at the account level:
  - **Source**: Select the registered AWS account that contains the destination Amazon Redshift serverless namespace.
  - **Namespace**: Select the Amazon Redshift Serverless namespace to which you want to restore the database.

To overwrite any existing database with the same name, enable **Overwrite existing database(s)** and type **Yes** to confirm.

> **Note:** Restoring to a new location also replaces any existing database with the same name in the selected cluster.

7. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**:

     Select this option if you chose **Auto Create** for the staging bucket during the backup process.
   - **Specify Bucket to Use**:

     Select this option if you created a staging bucket during the backup process. Ensure to specify the **same S3 bucket** that was used for the backup.

     > **Important:** Recovery may fail if you do not use the same staging bucket that was used during the backup operation.

   Cohesity deletes only the temporary data stored in the staging bucket during each recovery operation. The bucket itself is retained and reused for future recovery operations.
8. Click the **Edit** icon of the **Rename** field and enter a **New Name** for the database if you want to recover with a new name.
9. In the **Task Name** field, update the default name of the recovery task.
10. Click **Recover**.

Azure

Cohesity natively integrates with Microsoft Azure to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Azure Services, such as virtual machines, databases and so on.

This chapter includes the following topics:

Required Permissions for Azure Workloads

You must assign custom roles to the service account to enable Azure data protection. Before you Register an Azure Source at Tenant Levelon Cohesity DataProtect, ensure you set the required permissions.

The following tables list the permissions and privileges for the Azure workloads that must be included in a custom role.

Cloud Native Backup and Recovery and File/Folder Recovery

> **Note:** To create custom roles using the Azure Portal, use the format provided here.
>
> To create custom roles using the Azure PowerShell, use the format provided here.

Action	Resource Provider	Operation Name
Cloud Native Backup and Recovery and File/Folder Recovery	Microsoft.Resources	Microsoft.Resources/subscriptions/resourceGroups/read Microsoft.Resources/subscriptions/resourceGroups/write
	Microsoft.Storage	Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/listkeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write
	Microsoft.Network	Microsoft.Network/dnszones/A/read Microsoft.Network/dnszones/A/write Microsoft.Network/dnszones/A/delete Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/write Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/delete Microsoft.Network/networkInterfaces/ipconfigurations/read Microsoft.Network/networkSecurityGroups/read Microsoft.Network/networkSecurityGroups/join/action Microsoft.Network/networkSecurityGroups/securityRules/read Microsoft.Network/privateEndpoints/read Microsoft.Network/privateEndpoints/write Microsoft.Network/privateEndpoints/delete Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/read Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action

Action	Resource Provider	Operation Name
	Microsoft.Compute	Microsoft.Compute/disks/beginGetAccess/action Microsoft.Compute/disks/delete Microsoft.Compute/disks/endGetAccess/action Microsoft.Compute/disks/read Microsoft.Compute/disks/write Microsoft.Compute/diskAccesses/read Microsoft.Compute/diskAccesses/write Microsoft.Compute/diskAccesses/delete Microsoft.Compute/diskAccesses/privateEndpointConnectionsApproval/action Microsoft.Compute/snapshots/read Microsoft.Compute/snapshots/write Microsoft.Compute/snapshots/delete Microsoft.Compute/snapshots/beginGetAccess/action Microsoft.Compute/snapshots/endGetAccess/action Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/powerOff/action Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/runCommand/action Microsoft.Compute/availabilitySets/read
	Microsoft.KeyVault	Microsoft.KeyVault/vaults/deploy/action
Disk-level Recovery	Microsoft.Resources	Microsoft.Resources/subscriptions/resourceGroups/write Microsoft.Resources/subscriptions/resourceGroups/delete

Cloud Snapshot Manager (CSM) Backup and Recovery

> **Note:** To create custom roles using the Azure Portal, use the format provided here.
> To create custom roles using the Azure PowerShell, use the format provided here.

Action	Resource Provider	Operation Name
Cloud Snapshot Manager (CSM) Backup and Recovery	Microsoft.Resources	Microsoft.Resources/subscriptions/resourceGroups/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/listkeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write
	Microsoft.Network	Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/write Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/delete Microsoft.Network/networkSecurityGroups/join/action Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/read Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action
	Microsoft.Compute	Microsoft.Compute/disks/beginGetAccess/action Microsoft.Compute/disks/delete Microsoft.Compute/disks/endGetAccess/action Microsoft.Compute/disks/read Microsoft.Compute/disks/write Microsoft.Compute/snapshots/read Microsoft.Compute/snapshots/write Microsoft.Compute/snapshots/delete Microsoft.Compute/snapshots/beginGetAccess/action Microsoft.Compute/snapshots/endGetAccess/action Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/delete Microsoft.Compute/availabilitySets/read

CloudSpin Policy Based

> **Note:** To create custom roles using the Azure Portal, use the format provided here.
> To create custom roles using the Azure PowerShell, use the format provided here.

Action	Resource Provider	Operation Name
Policy based CloudSpin	Microsoft.Resources	Microsoft.Resources/subscriptions/resourceGroups/read Microsoft.Resources/deployments/read Microsoft.Resources/deployments/write
	Microsoft.Storage	Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/listkeys/action Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/delete
	Microsoft.Compute	Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/delete Microsoft.Compute/availabilitySets/read

CloudSpin On Demand

> **Note:** To create custom roles using the Azure Portal, use the format provided here.
> To create custom roles using the Azure PowerShell, use the format provided here.

Action	Resource Provider	Operation Name
CloudSpin On Demand	Microsoft.Resources	Microsoft.Resources/subscriptions/resourceGroups/read Microsoft.Resources/deployments/read Microsoft.Resources/deployments/write
	Microsoft.Storage	Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/listkeys/action Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/read Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action Microsoft.Network/networkInterfaces/write Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/delete
	Microsoft.Compute	Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/delete Microsoft.Compute/virtualMachines/runcommand/action Microsoft.Compute/availabilitySets/read

Required Permissions for Azure MySQL Protection

You must assign custom roles to the service account to enable Azure MySQL data protection. The following table lists the permissions and privileges that must be included in

a custom role.

Permission	Description
Microsoft.DBforMySQL/flexibleServers/read	List and read the properties of the server.
Microsoft.DBforMySQL/flexibleServers/databases/read	List and read the properties of the database under a server.
Microsoft.DBforMySQL/flexibleServers/databases/write	Create or update the properties of the database.
Microsoft.DBforMySQL/flexibleServers/databases/delete	Delete the database.
Microsoft.Authorization/*/read	List the permission associated with resource for given credentials.
Microsoft.Subscription/*/read	List the permission associated with resource for given credentials.
Microsoft.Resources/*/read	Get resource group and location information.
Microsoft.ManagedIdentity/userAssignedIdentities/read	Get the list of UMI in the given subscription.
Microsoft.ManagedIdentity/operations/read	View available operations, providers, and metadata related to managed identities.

Required Roles for Azure MySQL Protection

To create custom roles using the Azure Portal, use the following format:

Cohesity Azure MySQL Discovery

```
{
"Name": "Cohesity Azure MySQL DBPaaS Discovery",
"IsCustom": true,
"Description": "Allows Cohesity DBPaaS to discover subscriptions, resources, managed identities, and Azure Database for MySQL Flexible Servers using read-only permissions.",
"_comment": "Discovery-only role. No write or delete permissions.",
"Actions": [
"Microsoft.Authorization/*/read",
"Microsoft.Subscription/*/read",
"Microsoft.Resources/*/read",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.ManagedIdentity/operations/read",
"Microsoft.DBforMySQL/flexibleServers/read",
"Microsoft.DBforMySQL/flexibleServers/databases/read"
```

```
],
"NotActions": [],
"DataActions": [],
"NotDataActions": [],
"AssignableScopes": [
"/subscriptions/<SUBSCRIPTION_ID>"
]
}
```

Cohesity Azure MySQL Backup

```
{
"Name": "Cohesity Azure MySQL DBPaaS Backup",
"IsCustom": true,
"Description": "Allows Cohesity DBPaaS to perform backup operations on Azure
Database for MySQL Flexible Servers.",
"_comment": "Used for backup workflows where database-level write access is
required.",
"Actions": [
"Microsoft.DBforMySQL/flexibleServers/read",
"Microsoft.DBforMySQL/flexibleServers/databases/read",
"Microsoft.DBforMySQL/flexibleServers/databases/write",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.ManagedIdentity/operations/read"
],
"NotActions": [],
"DataActions": [],
"NotDataActions": [],
"AssignableScopes": [
"/subscriptions/<SUBSCRIPTION_ID>"
]
}
```

Cohesity Azure MySQL Recovery

```
{
"Name": "Cohesity Azure MySQL DBPaaS Restore",
"IsCustom": true,
"Description": "Allows Cohesity DBPaaS to restore Azure Database for MySQL
Flexible Server databases, including overwrite and cleanup operations.",
"_comment": "Highest-privilege role. Should be assigned only for restore
operations.",
"Actions": [
"Microsoft.DBforMySQL/flexibleServers/read",
"Microsoft.DBforMySQL/flexibleServers/databases/read",
"Microsoft.DBforMySQL/flexibleServers/databases/write",
"Microsoft.DBforMySQL/flexibleServers/databases/delete",
"Microsoft.ManagedIdentity/userAssignedIdentities/read",
"Microsoft.ManagedIdentity/operations/read"
],
```

```
"NotActions": [],
"DataActions": [],
"NotDataActions": [],
"AssignableScopes": [
"/subscriptions/<SUBSCRIPTION_ID>"
]
}
```

Required Permissions for Azure Cosmos DB – Cassandra Protection

You must assign custom roles to the service account to enable Azure Cosmos DB – Cassandra protection.

The following table lists the permissions and privileges that must be included in a custom role.

Permission	Description
Microsoft.Resources/subscriptions/resourceGroups/read	Gets or lists resource groups.
Microsoft.DocumentDB/databaseAccounts/read	Gets or lists Cosmos DB Accounts.
Microsoft.DocumentDB/databaseAccounts/cassandraKeyspaces/throughputSettings/read	Reads Cassandra keyspace throughput.
Microsoft.DocumentDB/databaseAccounts/cassandraKeyspaces/tables/throughputSettings/read	Reads Cassandra table throughput.
Microsoft.DocumentDB/databaseAccounts/cassandraKeyspaces/views/throughputSettings/read	Reads Cassandra Materialized View throughput.
Microsoft.DocumentDB/databaseAccounts/cassandraKeyspaces/read	Gets or lists Cassandra Keyspaces.

Permission	Description
Microsoft.Insights/Metrics/Read	Retrieves operational metrics from Azure Monitor for Cosmos DB Cassandra table/keyspace. This permission enables the backup/recovery process to query Azure Monitor metrics for Request Unit (RU) consumption during data ingestion. The agent uses these metrics to dynamically calculate the optimal backup/recovery speed that maximizes throughput while staying within provisioned RU limits, preventing throttling errors and minimizing retime.
Microsoft.DocumentDB/databaseAccounts/listKeys/action	Lists primary and secondary keys for data transfer.

Required Roles for Azure Cosmos DB – Cassandra Protection

To create custom roles using the Azure Portal, use the following format:

```
{
"properties": {
"roleName": "AzureCosmosDBCassandraBackupRecovery",
"description": "Permission for Azure Cosmos DB Cassandra Keyspaces Backup
```

```
and Recovery",
"assignableScopes": [
"/subscriptions/subscriptionID"
],
"permissions": [
{
"actions": [
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.DocumentDB/databaseAccounts /read",
"Microsoft.DocumentDB/databaseAccounts /cas-
sandraKeyspaces/throughputSettings/read",
"Microsoft.Docu-
mentDB/databaseAccounts/cassandraKeyspaces/tables/throughputSettings/read"
"Microsoft.Docu-
mentDB/databaseAccounts/cassandraKeyspaces/views/throughputSettings/read",
“Microsoft.DocumentDB/databaseAccounts/cassandraKeyspaces/read”,
“Microsoft.Insights/Metrics/Read”,
],
"notActions": [],
"dataActions": [],
"notDataActions": []
}
]
```

For information on creating a custom role, see Create or update Azure custom roles using the Azure portal.

Required Permissions for Azure Cosmos DB - MongoDB Protection

You must assign custom roles to the service account to enable Azure Azure Cosmos DB - MongoDB data protection. The following table lists the permissions and privileges that must be included in a custom role.

Permission	Description
Microsoft.DocumentDB/databaseAccounts/read	Read a database account.
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/read	Read a MongoDB database or list all the MongoDB databases.

Permission	Description
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/write	Create a MongoDB database.
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/collections /read	Read a MongoDB collection or list all the MongoDB collections.
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/collections /write	Create or update a MongoDB collection.
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/collections/throughputSettings/read	Read a MongoDB collection throughput.
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/collections/throughputSettings/write	Update MongoDB collection throughput.
Microsoft.Resources/subscriptions/resourceGroups/read	Discover resource groups containing Cosmos DB accounts.
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/throughputSettings/read	Read database-level throughput settings for capacity metadata.
Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/throughputSettings/write	Set throughput on newly created databases during recovery.

Permission	Description
Microsoft.DocumentDB/locations/read	Read Cosmos DB location metadata for precheck validation.
Microsoft.Insights/Metrics/Read	Read Azure Monitor metrics for database size estimation.
Microsoft.DocumentDB/databaseAccounts/listConnectionStrings/action	Get MongoDB connection strings for data plane access (alternative to listKeys).
Microsoft.DocumentDB/databaseAccounts/mongodbRoleDefinitions/read	Read existing MongoDB role definitions.
Microsoft.DocumentDB/databaseAccounts/mongodbRoleDefinitions/write	Create/update custom MongoDB role definitions.
Microsoft.DocumentDB/databaseAccounts/mongodbUserDefinitions/read	Read existing MongoDB user definitions.
Microsoft.DocumentDB/databaseAccounts/mongodbUserDefinitions/write	Create/update MongoDB user definitions for the Managed Identity.

Setup Azure Cosmos DB — MongoDB

The following is the procedure to setup Azure Cosmos DB — MongoDB.

Prerequisites

- User with the Global Administrator role
- Service Principal Object ID (from Enterprise Applications > your app > Object ID). Do not use the Application (Client) ID.
- Cosmos DB account name, Subscription ID, Resource group name

Variables

Replace the values with your actual values below:

```
SUBSCRIPTION_ID="<your-subscription-id>"
RESOURCE_GROUP="<your-resource-group>"
COSMOS_ACCOUNT="<your-cosmosdb-mongodb-account-name>"
PRINCIPAL_ID="<service-principal-object-id>"
ACCOUNT_SCOPE="/subscriptions/${SUBSCRIPTION_ID}/resourceGroups/${RESOURCE_
GROUP}/providers/Microsoft.DocumentDB/databaseAccounts/${COSMOS_ACCOUNT}"
```

Step 1: Assign Control Plane Role (Account Discovery)

```
az role assignment create \
--assignee-object-id "${PRINCIPAL_ID}" \
--assignee-principal-type "ServicePrincipal" \
--role "Cosmos DB Account Reader Role" \
--scope "/subscriptions/${SUBSCRIPTION_ID}"
```

> **Note:** If the Service Principal already has the Contributor role at the resource group level, this step can be skipped.

Step 2: Create Custom ARM Role for MongoDB CRUD + Throughput (Control Plane)

```
cat > /tmp/mongodb-custom-role.json << 'ROLEEOF'
{
"roleName": "Cohesity Cosmos MongoDB Role",
"assignableScopes": [
"/subscriptions/<your-subscription-id>"
],
"permissions": [{
"actions": [
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.DocumentDB/databaseAccounts/read",
"Microsoft.DocumentDB/databaseAccounts/listKeys/action",
"Microsoft.DocumentDB/databaseAccounts/listConnectionStrings/action",
"Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/read",
"Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/write",
"Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/delete",
"Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/collections/read",
"Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/collections/write",
"Microsoft.DocumentDB/databaseAccounts/mongodbDatabases/collections/delete",
"Microsoft.Docu-
mentDB/databaseAccounts/mongodbDatabases/throughputSettings/read",
```

```
"Microsoft.Docu-
mentDB/databaseAccounts/mongodbDatabases/throughputSettings/write",
"Microsoft.Docu-
mentDB/data-
baseAccounts/mongodbDatabases/collections/throughputSettings/read",
"Microsoft.Docu-
mentDB/data-
baseAccounts/mongodbDatabases/collections/throughputSettings/write",
"Microsoft.DocumentDB/locations/read",
"Microsoft.Insights/Metrics/Read"
]
}]
}
ROLEEOF

az role definition create --role-definition @/tmp/mongodb-custom-role.json
```

Step 3: Assign the Custom Role to the Service Principal

```
az role assignment create \
--assignee-object-id "${PRINCIPAL_ID}" \
--assignee-principal-type "ServicePrincipal" \
--role "Cohesity Cosmos MongoDB Role" \
--scope "${ACCOUNT_SCOPE}"
```

Step 4: (Optional) For UMI/SMI Database-Level RBAC — Create MongoDB User

If using Managed Identity for data plane access instead of connection strings:

```
# Enable Cosmos DB MongoDB RBAC on the account
az cosmosdb update \
--name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--capabilities "EnableMongoRoleBasedAccessControl"

# Create a MongoDB user with dbOwner role (for backup + restore)
az cosmosdb mongodb user definition create \
--account-name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--body '{
"Id": "'${COSMOS_ACCOUNT}'.cohesity-backup-user",
"UserName": "cohesity-backup-user",
"Password": "",
"DatabaseName": "admin",
"CustomData": "Cohesity backup and restore user",
"Roles": [
{"Role": "dbOwner", "Db": "admin"}
],
"Mechanisms": "SCRAM-SHA-256"
}'
```

Step 5: Verify All Assignments

```
echo "=== Azure RBAC (Control Plane) Roles ==="
az role assignment list \
--assignee "${PRINCIPAL_ID}" \
--scope "${ACCOUNT_SCOPE}" \
--output table

echo "=== Cosmos DB MongoDB Users (if RBAC enabled) ==="
az cosmosdb mongodb user definition list \
--account-name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--output table
```

Required Permissions for Azure Cosmos DB – NoSQL Protection

You must assign custom roles to the service account to enable Azure Cosmos DB – NoSQL protection.

The following table lists the permissions and privileges that must be included in a custom role.

Permission	Description
Microsoft.DocumentDB/databaseAccounts/read	Read Cosmos DB account properties and metadata
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/read	Read SQL database properties and metadata
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/write	Create or update SQL databases
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/read	Read container properties and metadata
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/write	Create or update containers

Permission	Description
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/throughputSettings/read	Read container throughput settings (RU/s)
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/throughputSettings/write	Update container throughput settings (RU/s)
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/storedProcedures/read	Read stored procedures
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/storedProcedures/write	Create or update stored procedures
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/triggers/read	Read triggers
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/triggers/write	Create or update triggers
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/userDefinedFunctions/read	Read user-defined functions
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/userDefinedFunctions/write	Create or update user-defined functions
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/throughputSettings/read	Read database throughput settings (RU/s)
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/throughputSettings/write	Update database throughput settings (RU/s)
Microsoft.Insights/Metrics/Read	Used by Go adapter monitor client for RU consumption monitoring.

Permission	Description
Microsoft.DocumentDB/locations/read	Used during precheck validation.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/replace	Required for upsert_item() replace semantics during recovery.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/delete	Required for database overwrite or cleanup on recovery failure.

Required Roles for Azure Cosmos DB – NoSQL Protection

To create custom roles using the Azure Portal, use the following format:

```
{
"RoleName": "CosmosDB-Backup-Restore-Role",
"Type": "CustomRole",
"AssignableScopes": ["/subscriptions/subscriptionID"],
"Permissions": [{
"DataActions": [
"Microsoft.DocumentDB/databaseAccounts/readMetadata",
"Microsoft.DocumentDB/databaseAccounts/sqlDatabases/read",
"Microsoft.DocumentDB/databaseAccounts/sqlDatabases/write",
"Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/read",
"Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/write",
"Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/read",
"Microsoft.Docu-
mentDB/databaseAccounts/sqlDatabases/containers/items/create",
"Microsoft.Docu-
mentDB/databaseAccounts/sqlDatabases/containers/items/upsert",
"Microsoft.Docu-
mentDB/databaseAccounts/sqlDatabases/containers/executeQuery",
"Microsoft.Docu-
mentDB/databaseAccounts/sqlDatabases/containers/readChangeFeed"
]
}]
}
```

Required Permissions for Azure Cosmos DB - Table Protection

The following are the permissions required for the Azure Cosmos DB - Table protection operations such as registration, backup, and recovery.

Azure Cosmos DB - Table Registration

The following are the permissions required for the Azure Cosmos DB - Table source registration, discovery, and DBSize fetch:

Permission	Type	Description
Microsoft.Resources/subscriptions/resourceGroups/read	Action	Discover resource groups containing Cosmos DB accounts.
Microsoft.DocumentDB/databaseAccounts/read	Action	Discover and read Cosmos DB account properties.
Microsoft.DocumentDB/databaseAccounts/tables/read	Action	List tables in the account and fetch table metadata/size.
Microsoft.DocumentDB/databaseAccounts/readMetadata	DataAction	Read account-level metadata for data-plane operations.

Azure Cosmos DB - Table Backup

Permission	Type	Description
Microsoft.DocumentDB/databaseAccounts/tables/read	Action	List tables to determine objects to back up.
Microsoft.DocumentDB/databaseAccounts/readMetadata	DataAction	Read account-level metadata for data-plane operations.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/read	DataAction	Read table entities/items for full backup.

Permission	Type	Description
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/readChangeFeed	DataAction	Track item/entity changes for incremental backup.
Microsoft.Insights/Metrics/Read	Action	Read Azure Monitor metrics for RU consumption during backup.

Azure Cosmos DB - Table Recovery

Permission	Type	Description
Microsoft.DocumentDB/databaseAccounts/tables/write	Action	Create tables at the recovery destination.
Microsoft.DocumentDB/databaseAccounts/tables/read	Action	Check if the target table already exists.
Microsoft.DocumentDB/databaseAccounts/tables/delete	Action	Delete existing table before overwriting recovery.
Microsoft.DocumentDB/databaseAccounts/readMetadata	DataAction	Read account-level metadata for data-plane operations.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/create	DataAction	Create entities in the recovered table.

Permission	Type	Description
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/upsert	DataAction	Create or replace entities during recovery.
Microsoft.DocumentDB/databaseAccounts/throughputSettings/*	DataAction	Set throughput on newly created tables during recovery.
Microsoft.Insights/Metrics/Read	Action	Read Azure Monitor metrics for RU consumption during recovery.

Setup Azure Cosmos DB - Table Permissions

Prerequisites

Ensure the following are available:

- User with Global Administrator role.
- Service Principal Object ID (from Enterprise Applications > your app > Object ID). Do not use the Application (Client) ID.
- Cosmos DB Account name, Subscription ID, Resource Group name.

Variables

Ensure you replace the values with your actual values.

```
SUBSCRIPTION_ID="<your-subscription-id>"
RESOURCE_GROUP="<your-resource-group>"
COSMOS_ACCOUNT="<your-cosmosdb-account-name>"
PRINCIPAL_ID="<service-principal-object-id>"
ACCOUNT_SCOPE="/subscriptions/${SUBSCRIPTION_ID}/resourceGroups/${RESOURCE_
GROUP}/providers/Microsoft.DocumentDB/databaseAccounts/${COSMOS_ACCOUNT}"
```

Procedure

Following are the steps to setup the permissions:

1. Assign Control Plane Role (for Account Discovery)

```
az role assignment create \
--assignee-object-id "${PRINCIPAL_ID}" \
--assignee-principal-type "ServicePrincipal" \
```

```
--role "Cosmos DB Account Reader Role" \
--scope "/subscriptions/${SUBSCRIPTION_ID}"
```

> **Note:** You can skip this step if the Service Principal already has the Contributor role at the resource group level.

2. Assign Cosmos DB Built-in Data Contributor Role (Data Plane)

```
az cosmosdb sql role assignment create \
--account-name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--scope "/" \
--principal-id "${PRINCIPAL_ID}" \
--role-definition-id "00000000-0000-0000-0000-000000000002"
```

Where:
--principal-id: The Object ID of the Service Principal (not the Application/Client ID)
--role-definition-id: Built-in Data Contributor role ID (00000000-0000-0000-0000-000000000002)
--scope: Use / for the entire account

3. Create Custom Role for Throughput Settings (Data Plane)
The Built-in Data Contributor role does not include the throughputSettings permissions. A custom role is required for table creation during recovery.

```
cat > /tmp/throughput-role.json << 'ROLEEOF'
{
"RoleName": "Cohesity Table Throughput Access",
"Type": "CustomRole",
"AssignableScopes": ["/"],
"Permissions": [{
"DataActions": [
"Microsoft.DocumentDB/databaseAccounts/readMetadata",
"Microsoft.DocumentDB/databaseAccounts/throughputSettings/*"
]
}]
}
ROLEEOF
az cosmosdb sql role definition create \
--account-name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--body @/tmp/throughput-role.json
```

4. Assign the Custom Throughput Role

```
ROLE_ID=$(az cosmosdb sql role definition list \
--account-name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--query "[?roleName=='Cohesity Table Throughput Access'].name" -o tsv)

echo "Custom role GUID: ${ROLE_ID}"
az cosmosdb sql role assignment create \
--account-name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--scope "/" \
--principal-id "${PRINCIPAL_ID}" \
--role-definition-id "${ROLE_ID}"
```

5. Verify All Assignments

```
echo "=== Cosmos DB Data Plane Roles ==="
az cosmosdb sql role assignment list \
--account-name "${COSMOS_ACCOUNT}" \
--resource-group "${RESOURCE_GROUP}" \
--output table

echo "=== Azure RBAC (Control Plane) Roles ==="
az role assignment list \
--assignee "${PRINCIPAL_ID}" \
--scope "${ACCOUNT_SCOPE}" \
--output table
```

For Azure Cosmos DB - Table, DataProtect intentionally uses two Azure built-in roles and a custom role (for the throughputSettings/*) instead of a single all-in-one custom role. These permissions are Azure actions required by the workflow, and not a role definition.

This is due to the following reasons:

- **Control-plane vs data-plane separation.** Cosmos DB splits permissions into two different RBAC systems.
  - Action entries (control plane) - managed by Azure RBAC (az role)
  - DataAction entries (data plane) - managed by the Cosmos DB data-plane RBAC (az cosmosdb sql role).

  A single Azure RBAC custom role cannot include Cosmos DB DataAction permissions, and a single Cosmos DB data-plane custom role cannot include control-plane Action permissions. A minimum of two role assignments is required.
- **Data-plane roles cannot be assigned via Azure portal UI.** Cosmos DB data-plane roles must be assigned via Azure CLI. Relying on the existing built-in Data Contributor role keeps the procedure minimal and consistent with the Microsoft pattern.

- **Microsoft-maintained built-in roles are preferred where they fit.** The Cosmos DB Account Reader role and Cosmos DB Built-in Data Contributor role already contain most of the permissions. Using these ensures lesser maintenance, and any further permissions added by Microsoft to these roles will be picked automatically.
- **Custom roles only fix the gap. Microsoft.** DocumentDB/databaseAccounts/throughputSettings/* is the only permission not supported by the Built-in Data Contributor role. A custom role (step 3) grants only that permission, preserving least-privilege and avoiding a sprawling custom role that duplicates what the built-ins already provide.
- **Auditability.** The tables remain the source of truth for what and why permissions are used. The setup steps show how those permissions are granted (which built-in role or custom role covers each). Reviewers can map operation > permission > role without re-deriving the role contents.

Points To Remember

- After assigning roles, wait 5-10 minutes for role propagation before performing backup or recovery.
- Always use the Service Principal Object ID (from Enterprise Applications) and not the Application (Client) ID.
- Cosmos DB data plane roles cannot be assigned via the Azure portal UI; Azure CLI is required.
- The Cosmos DB Built-in Data Contributor role does not include throughputSettings permissions; the custom role is mandatory for recovery.
- User Managed Identity (UMI) and System Managed Identity (SMI) require the same permissions, except that you must use the Object ID of the Managed Identity resource instead of the Service Principal, to obtain the Principal ID.

Required Permissions for Azure SQL Protection

The following permissions are required for the custom role used by Managed identity:

Permission	Description
Microsoft.Authorization/permissions/read	List the permission associated with resource for given credentials
Microsoft.Subscription/subscriptions/read	Get resource group and location information
Microsoft.Resources/subscriptions/resourceGroups/read Microsoft.Resources/subscriptions/resources/read	Get SQL database and corresponding provider protected recovery (PIT,LTR)
Microsoft.ManagedIdentity/userAssignedIdentities/read	Get the list of UMI in the given subscription

Permission	Description
Microsoft.Sql/servers/read	Get the list the SQL server
Microsoft.Sql/servers/databases/read	Get the list of Databases in the SQL server
Microsoft.Sql/servers/databases/write	Create temporary database from PIT in the SQL server
Microsoft.Sql/servers/databases/delete	Delete temporary database from SQL server

Required Permissions for Azure SQL Managed Instance Protection

The following set of permissions are required to discover the SQL Managed Instance to the Cohesity cluster.

Permissions	Description
Microsoft.Authorization/permissions/read Microsoft.Subscription/subscriptions/read Microsoft.Resources/subscriptions/resourceGroups/read Microsoft.Resources/subscriptions/resources/read Microsoft.Sql/managedInstances/read Microsoft.Sql/managedInstances/databases/read	To list the permissions associated with a resource for a given credential.
Microsoft.Resources/*/read	To get resource group and location information.
Microsoft.Sql/*/read	To get SQL Managed Instance.
Microsoft.ManagedIdentity/userAssignedIdentities/read	To get the list of user-assigned managed identities (UMI) in the given subscription.
Microsoft.Sql/managedInstances/databases/write	To create a new database during recovery.
Microsoft.Sql/managedInstances/databases/delete	To delete a database during an overwrite recovery.

Non-CDC-based protection of SQL MI requires the following storage account permissions:

Microsoft.Resources/subscriptions/resources/read	To locate which resource group the storage account provided by the user belongs to.
Microsoft.Storage/storageAccounts/read	To read properties of storage account to decide whether to use SAS token or Managed Identity when creating T-SQL credentials on SQL MI.
Microsoft.Storage/storageAccounts/listKeys/action	To generate SAS token and connection strings.
Microsoft.Storage/storageAccounts/blobServices/containers/read	To check container existence before further operations.
Microsoft.Storage/storageAccounts/blobServices/containers/write	To create a staging container.
Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read	To download the database dump files during backup.
Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write	To write database dump files on staging container during recovery.
Microsoft.Storage/storageAccounts/blobServices/containers/blobs/delete	To delete the data inside container after operation has completed.

Required Permissions for Azure Table Storage Protection

The following set of permissions are required to discover the Azure Table Storage Accounts to Cohesity cluster.

Permissions	Type	Description
Microsoft.Storage/storageAccounts/tableServices/tables/read	DataAction	Read/list tables and get table properties in a storage account.
Microsoft.Storage/storageAccounts/tableServices/tables/write	DataAction	Create or update tables in a storage account.

Permissions	Type	Description
Microsoft.Storage/storageAccounts/tableServices/tables/delete	DataAction	Delete tables in a storage account.
Microsoft.Resources/subscriptions/resourceGroups/read	Action	Read/list resource groups and view their properties.
Microsoft.Storage/storageAccounts/read	Action	Read/list storage accounts and view their properties.
Microsoft.Storage/storageAccounts/tableServices/tables/entities/read	Action	Read/query table entities (retrieve entities from tables).
Microsoft.Storage/storageAccounts/tableServices/tables/entities/write	Action	Insert, update, merge, or upsert table entities.
Microsoft.DocumentDB/databaseAccounts/read	Action	Discover and read Cosmos DB account properties.
Microsoft.DocumentDB/databaseAccounts/tables/write	Action	Create tables at the recovery destination.
Microsoft.DocumentDB/databaseAccounts/tables/read	DataAction	Check if the target table already exists.

Permissions	Type	Description
Microsoft.DocumentDB/databaseAccounts/tables/delete	DataAction	Delete existing table before overwriting recovery.
Microsoft.DocumentDB/databaseAccounts/readMetadata	DataAction	Read account-level metadata for data-plane operations.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/read	DataAction	Read table entities/items for full backup.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/create	DataAction	Create entities in the recovered table.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/items/upsert	DataAction	Create or replace entities during recovery.
Microsoft.DocumentDB/databaseAccounts/sqlDatabases/containers/readChangeFeed	DataAction	Track item/entity changes for incremental backup.
Microsoft.DocumentDB/databaseAccounts/throughputSettings/*	DataAction	Set throughput on newly created tables during recovery.

Setup Azure Table Storage

The following is the procedure to setup Azure Table Storage:

Prerequisites

- User with the Global Administrator role.
- Service Principal Object ID (from Enterprise Applications > your app > Object ID). Do not use the Application (Client) ID.
- Subscription ID, Resource group name.

Variables

Replace the values with your actual values below:

```
SUBSCRIPTION_ID="<your-subscription-id>"
RESOURCE_GROUP="<your-resource-group>"
PRINCIPAL_ID="<service-principal-object-id>"
```

Step 1: Create a Custom Role for Table Storage

```
cat > /tmp/table-storage-role.json << 'ROLEEOF'
{
"roleName": "Cohesity Table Storage Role",
"assignableScopes": [
"/subscriptions/<your-subscription-id>"
],
"permissions": [{
"actions": [
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.Storage/storageAccounts/read"
],
"dataActions": [
"Microsoft.Storage/storageAccounts/tableServices/tables/read",
"Microsoft.Storage/storageAccounts/tableServices/tables/write",
"Microsoft.Storage/storageAccounts/tableServices/tables/delete",
"Microsoft.Storage/storageAccounts/tableServices/tables/entities/read",
"Microsoft.Storage/storageAccounts/tableServices/tables/entities/write",
"Microsoft.Storage/storageAccounts/tableServices/tables/entities/delete"
]
}]
}
ROLEEOF

az role definition create --role-definition @/tmp/table-storage-role.json
```

Step 2: Assign the Custom Role to the Service Principal

```
az role assignment create \
--assignee-object-id "${PRINCIPAL_ID}" \
--assignee-principal-type "ServicePrincipal" \
--role "Cohesity Table Storage Role" \
--scope "/subscriptions/${SUBSCRIPTION_ID}"
```

Step 3: Verify the Role Assignment

```
az role assignment list \
--assignee "${PRINCIPAL_ID}" \
```

```
--scope "/subscriptions/${SUBSCRIPTION_ID}" \
--output table
```

Points to Remember

- After assigning roles, wait a few minutes for role propagation before performing backup or recovery operations.
- Always use the Service Principal Object ID (from Enterprise Applications), not the Application (Client) ID.
- Cosmos DB data plane roles cannot be assigned via Azure portal UI; Azure CLI is required.
- The Cosmos DB Built-in Data Contributor role does not include throughputSettings permissions; the custom role is mandatory for recovery.
- The same permissions apply when using User Managed Identity or System Managed Identity. The only difference is that the PRINCIPAL_ID is obtained using the Object ID of the Managed Identity resource instead of the Service Principal.

Required Permissions for Azure PostgreSQL Protection

The following set of permissions are required to discover PostgreSQL to Cohesity cluster.

Permissions	Description
Microsoft.Authorization/*/read Microsoft.Subscription/*/read	To list the permission associated with the resource for the given credential.
Microsoft.Resources/*/read	To get resource group and location information.
Microsoft.DBforAzurePostgreSQL/servers/databases/write	To create a database in Azure PostgreSQL server.
Microsoft.DBforAzurePostgreSQL/servers/databases/delete	To delete a database in Azure PostgreSQL server.
Microsoft.DBforAzurePostgreSQL/flexibleServers/read	To list and get details of Azure PostgreSQL Flexible server.
Microsoft.DBforAzurePostgreSQL/flexibleServers/databases/read	To list and get details of databases in Azure PostgreSQL Flexible server.
Microsoft.DBforAzurePostgreSQL/flexibleServers/databases/write	To create a database in Azure PostgreSQL Flexible server.

Permissions	Description
Microsoft.ManagedIdentity/*/read	Allows read-only access to managed identity resources (system-assigned and user-assigned identities), including viewing identity properties. Does not allow assigning roles or changing identity configuration.
Microsoft.DBforPostgreSQL/servers/read	Allows read-only access to Azure Database for PostgreSQL (Single Server) server resources, including server configuration and properties.
Microsoft.DBforPostgreSQL/servers/databases/read	Allows read-only access to database resources under an Azure Database for PostgreSQL Single Server. Used to list or inspect databases at the ARM level.
Microsoft.DBforPostgreSQL/servers/databases/write	Allows creation or update of database resources under an Azure Database for PostgreSQL Single Server at the Azure Resource Manager level. Does not grant in-database (SQL) privileges.
Microsoft.DBforPostgreSQL/servers/databases/delete	Allows deletion of database resources under an Azure Database for PostgreSQL Single Server at the ARM level.
Microsoft.DBforPostgreSQL/flexibleServers/read	Allows read-only access to Azure Database for PostgreSQL Flexible Server resources, including configuration and properties.
Microsoft.DBforPostgreSQL/flexibleServers/databases/read	Allows read-only access to database resources under an Azure Database for PostgreSQL Flexible Server at the ARM level.
Microsoft.DBforPostgreSQL/flexibleServers/databases/write	Allows creation or update of database resources under an Azure Database for PostgreSQL Flexible Server at the ARM level.
Microsoft.DBforPostgreSQL/flexibleServers/databases/delete	Allows deletion of database resources under an Azure Database for PostgreSQL Flexible Server at the ARM level.

Permissions Required to Refresh Materialized Views in Azure PostgreSQL

To refresh materialized views automatically or manually, the database user must have the required schema and object permissions.

For materialized view refresh:

```
GRANT USAGE ON SCHEMA <schema_name> TO <refresh_user>;
GRANT SELECT ON ALL TABLES IN SCHEMA <schema_name> TO <refresh_user>;
GRANT ALL ON <materialized_view_name> TO <refresh_user>;
```

For dependency analysis:

```
GRANT SELECT ON pg_class, pg_namespace, pg_depend, pg_rewrite TO refresh_
user;
```

Where you must,

- replace <schema_name> with your actual schema name.
- replace <refresh_user> with the database user configured in your DataProtect source.
- replace <materialized_view_name> with each materialized view name.

> **Note:** Ensure to use manual refresh after recovery if the database user does not have the required permissions.

For manual execution of script:

```
GRANT CONNECT ON DATABASE <restore_database_name> TO <refresh_user>;
GRANT USAGE, CREATE ON SCHEMA <script_schema> TO <refresh_user>;
```

Where, <script_schema> is the first schema in the search_path (often public). This is required for the script to create and drop temporary helper views.

Manual execution of script requires connection to database and usage on the schema to create temporary views.

Permissions Required for Azure Storage Protection

The following table lists the permissions and privileges that must be included in a custom role.

Permission Type	Premission
Actions	Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/listkeys/action Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Resources/subscriptions/resourcegroups/read
DataActions	Microsoft.Storage/storageAccounts/blobServices/containers/blobs/*

Cohesity Cloud Edition

Action	Resource Provider	Operation Name
Cloud Edition creation using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/deployments/validate/action Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/subscriptions/locations/read Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write "Microsoft.Storage/storageAccounts/delete Microsoft.Storage/storageAccounts/listAccountSas/action Microsoft.Storage/storageAccounts/blobServices/containers/write
	Microsoft.Network	Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/write Microsoft.Network/networkInterfaces/delete Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/join/action
	Microsoft.Compute	Microsoft.Compute/disks/write Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/start/action

Action	Resource Provider	Operation Name
Add or remove a node using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/deployments/read Microsoft.Resources/deployments/write Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/read "Microsoft.Storage/storageAccounts/write "Microsoft.Storage/storageAccounts/delete Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/listKeys/action
	Microsoft.Network	Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/write Microsoft.Network/networkInterfaces/delete Microsoft.Network/networkInterfaces/join/action Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/join/action
	Microsoft.Compute	Microsoft.Compute/disks/write Microsoft.Compute/disks/delete Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/delete

Action	Resource Provider	Operation Name
Add or remove a node using an existing node in the cluster.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/deployments/read Microsoft.Resources/deployments/write Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/listKeys/action
	Microsoft.Network	Microsoft.Network/virtualNetworks/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action

Action	Resource Provider	Operation Name
Power on the cluster using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/blobServices/containers/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read

Action	Resource Provider	Operation Name
Power off the cluster using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/blobServices/containers/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read

Action	Resource Provider	Operation Name
Destroy the cluster using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/blobServices/containers/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read Microsoft.Compute/disks/delete

Required Ports for Azure Workloads

Firewall Ports for Azure Cosmos DB – Cassandra Protection

Azure Cosmos DB offers two modes of connection - gateway mode and direct mode.

Connection Mode	Supported Protocol	API/Service Port
Gateway	HTTPS	Cassandra - 10350
Direct	TCP (Encrypted via TLS)	When using public/service endpoints: ports in the 10000 through 20000 range. When using private endpoints: ports in the 0 through 65535 range.

You can connect to the Azure Cosmos DB - Cassandra account from the Cohesity host using the following methods under **Settings > Networking** in the Microsoft Azure portal:

- Public Access
  - All networks.
  - Selected networks (with Cohesity host CIDR range in the Firewall allowed host - moderate security).
- Private Access

  To use this method, you must create a Private EndPoint. For more information, see Create a private endpoint by using the Azure portal. Once the EndPoint is created,

  1. Login to the Cohesity cluster.
  2. Navigate to **Settings > Networking > Host Mapping**.
  3. Add the Cassandra Private Endpoint IP host mapping.

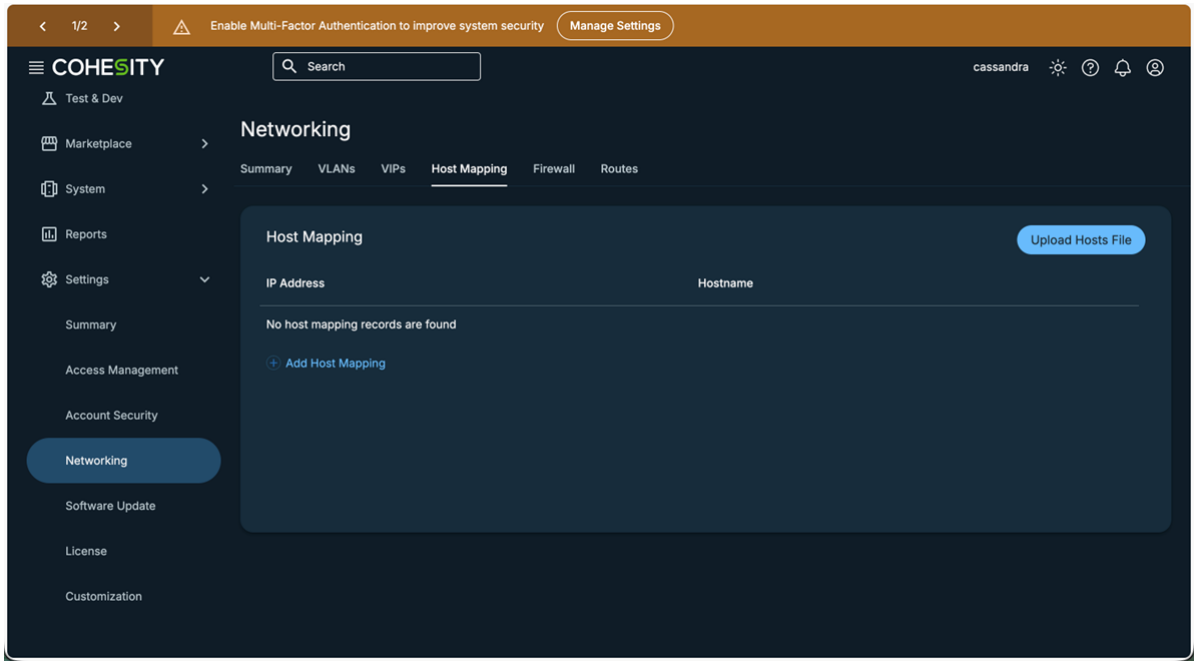

Firewall Ports for Azure Cosmos DB – NoSQL Protection

Azure Cosmos DB offers two modes of connection - gateway mode and direct mode.

Connection Mode	Supported Protocol	API/Service Port
Gateway	HTTPS	SQL (443), MongoDB (10255), Table (443), Cassandra (10350), Graph (443)
Direct	TCP (Encrypted via TLS)	When using public/service endpoints: ports in the 10000 through 20000 range. When using private endpoints: ports in the 0 through 65535 range.

For more information, see Azure Cosmos DB SQL SDK connectivity modes.

Firewall Ports for Azure SQL Database Protection

Ensure connectivity from Cohesity DataProtect to the Azure SQL database on port 1433.

You can connect to the Azure SQL databases from the Cohesity host using the following methods under **Settings > Networking** in the Microsoft Azure portal:

- Public Access
  - All networks.
  - Selected networks (with Cohesity host CIDR range in the Firewall allowed host - moderate security).
- Private Access
  - To use this method, you must create a Private EndPoint. For more information, see Create a private endpoint by using the Azure portal.
  - Once the EndPoint is created, login to the Cohesity cluster.
  - Navigate to **Settings > Networking > Host Mapping**.

- Add the SQL Private Endpoint IP host mapping.

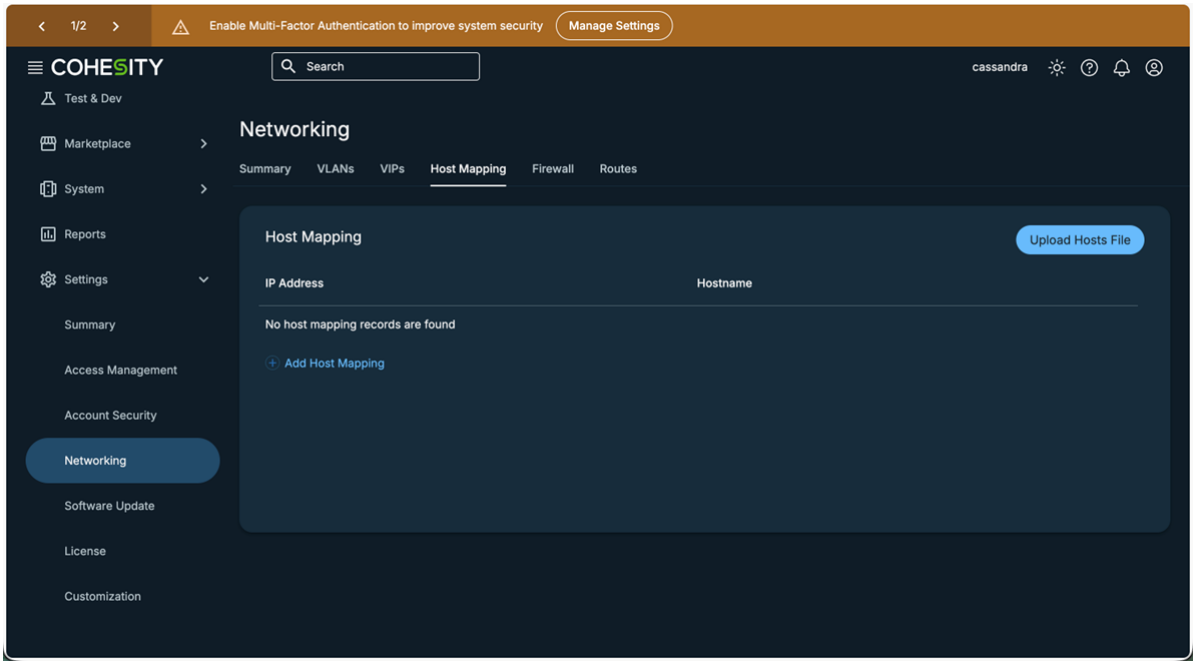

Firewall Ports for Azure SQL Managed Instance Protection

Ensure connectivity from Cohesity cluster to the Azure SQL Managed Instance server on port 1433.

Firewall Ports for Azure PostgreSQL Protection

Ensure connectivity from Cohesity cluster to the Azure PostgreSQL server on port 5432.

Register Azure Source

You can register Azure source with the Cohesity cluster using the following ways:

- Register an Azure Cloud Source at Subscription Level
- Register an Azure Source at Tenant Level

Registering Azure source at tenant level is supported only on Cohesity Cloud Edition on Azure.

Register an Azure Source at Tenant Level

You must register Azure as a data source at tenant level to protect your Azure workloads.

> **Note:** Registering Azure source at tenant level is supported only on Cohesity Cloud Edition on Azure.

Prerequisites for Source Registration Using Managed Identity

Cohesity DataProtect supports managed identity - User-Assigned Managed Identity (UMI) and System-Assigned Managed Identity (SMI), to access Azure APIs for Azure native data protection.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Ensure the following when registering an Azure source using the Managed Identity authentication methods:

- Cohesity DataProtect runs on an Azure VM (for example, NGCE node), so that Instance Metadata Service (IMDS) is available for SMI and UMI token acquisition.
- Managed Identity configurations are set for authentication method.
  - **SMI**: System-Assigned Managed Identity (SMI) is enabled on the Azure VM where Cohesity runs.
  - **UMI**: User-Assigned Managed Identity (UMI) is created and attached to the Azure VM. The Client ID of the managed identity is required during registration.
  - **Role-Based Access Control (RBAC)**: The managed identity has the required permissions for the specific workload. For more information, see the **Plan and Prepare** section of the respective Azure workload.

Register Azure as a Source

Before registering Azure as a source, ensure to check the requirements and considerations for the required workload.

To register Azure as a source:

1. Navigate to **Data Protection > Sources > Register**, and then select **Cloud Platforms > Azure**.

   The **Register Azure Source** window is displayed.
2. In the **Source Details** page,
   1. In the **Source Details** section, under the **Domain Name** field, provide the unique tenant ID assigned by Azure. For information on getting the tenant ID, see Get tenant ID.
   2. In the **Category** section, select the Azure cloud environment associated with the tenant:
      - **Standard** – For Microsoft Azure commercial (public cloud) tenants.
      - **Gov** – For Microsoft Azure Government environment tenants.

> **Note:** The category determines the Azure cloud environment used during source registration and cannot be changed after the source is registered.

3. In the **Azure Workloads** section, select the workloads you want to protect:

   Databases

   - MySQL
   - PostgreSQL
   - SQL
   - SQL MI
   - Table Storage

   Virtual Machine

   - Azure VM

   Storage

   - Blob and Data Lake Storage

   > **Note:** The available workload options depend on the selected category:
   > - If the selected category is **Standard**, all supported Azure workloads are available for selection.
   > - If the selected category is **Gov**, only Azure VM and Azure Blob Storage workloads are available for selection.

4. Click **Continue**.

3. In the **Azure Apps** page,under the **Authentication Method** section, select the method for Cohesity to authenticate with Microsoft Azure to access and manage Azure resources.

   Azure sources registered using managed identity-based authentication use Managed Identity to access Azure APIs for all the data protection services, including backup and recovery.

   > **Note:** Some authentication methods require additional configuration in the Microsoft Azure portal. For more information, see the **Plan and Prepare** section of the respective Azure workload.

1. Choose one of the following authentication methods:
   - **Azure App** method uses an Azure Active Directory application (service principal) for authentication.
   - In the **Application ID** field, enter the application ID assigned by Azure during the service principal creation process and in the **Application Secret Key** field, enter the application secret key for setting up authentication for the service principal.
   - **User Managed Identity** method uses a dedicated managed identity that you create and assign to the Azure resources.

     In the **Client ID** field, enter the client ID of the User Managed Identity.

     > **Note:** Ensure the managed identity has the required role assignments on the target Azure subscriptions. For more information, see the Plan and Prepare section of the respective Azure workload.

   - **System Managed Identity** method uses the system built-in managed identity automatically created for Cohesity DataProtect in Azure.

     > **Note:** This method requires role assignments on Azure subscriptions. For more information, see the Plan and Prepare section of the respective Azure workload.

   > **Note:** Managed Identity is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

2. Click **Continue**.

4. In the **Subscriptions** page,
   1. Under the **Add subscriptions for the source** option, in the **Subscription ID** field, enter the subscription IDs of the subscriptions where the database you want to protect belongs to.

      You can add multiple subscription IDs by clicking the **+** icon.

      > **Note:**
      > - Subscriptions must be manually added and will not be automatically

> discovered.

2. Click **Register**.

The source is registered and displayed in the **Data Protection > Sources** page, under the **Azure** section. You can click on the source to view the entity tree of the Azure servers and databases.

Register an Azure Cloud Source at Subscription Level

A Cohesity cluster supports registering an Azure Cloud Source for the following functionality:

- **Converting, cloning, and copying Snapshots of VMware VMs**—For this functionality, the registered Azure Cloud Source is where the VM Snapshots are cloned or copied to. Converting, cloning or copying VM Snapshots occurs during the following use cases:
  - Clone VMs from On-Premises to Azure Cloud Service (On Demand CloudSpin)
  - Clone VMs to Azure Cloud Using CloudSpin
  - Replicate and Clone VMs in Azure Cloud Edition
- **Backing up Azure infrastructure** — For this functionality, the registered Azure Cloud source contains the Azure workloads to back up.

To register an Azure Cloud Source:

> **Note:** Before registering the Azure Cloud Source with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations.

1. Select **Data Protection > Sources**.
2. Select **Register > Virtual Machines** and select **Azure: Azure Subscription** from the Source drop-down.
3. Enter the information described below and click **Register Source**.

Field Name	Description
Category	Select one of the following options: • **Standard**–Standard Azure. • **Gov**–Azure Government. • **Stack**–Azure Stack Hub. You must specify the authentication type, region, and the domain name. For more information, see Azure Stack Hub.

Field Name	Description
Subscription ID	This is the subscription ID from which you want to back up data. Log in to the Azure portal. In the left panel, click **Subscriptions**. Copy the SUBSCRIPTION ID from the table.
Application ID	The Application ID assigned by Azure during the service principal creation process. For instructions see Get application ID and authentication key in the Microsoft Azure documentation.
Application Key	The Azure Application key generated using Legacy App Registration during the service principal creation process that is used for authentication. For instructions see Get application ID and authentication key in the Microsoft Azure documentation.
Tenant ID	The unique tenant ID assigned by Azure. For instructions see Get tenant ID in the Microsoft Azure documentation. > **Note:** The **Tenant ID** field is not displayed for Azure Stack Hub Active Directory Federation Services (AD FS) authentication type.
Stack Type	If you select Stack, select the authentication type: • Azure Active Directory • Active Directory Federation Services Azure Stack Hub uses either Azure Active Directory or AD FS for authentication. AD FS is mostly used for disconnected deployments of Azure Stack Hub. For more information, see Microsoft Azure Stack Hub documentation.
Region	If you select **Stack**, specify the Azure Stack Hub region. For instructions, see Region management in Azure Stack Hub in the Microsoft Azure documentation.
Domain Name	If you select **Stack**, specify the Azure Stack Hub domain name.

Support for HTTP Proxy Server

Cohesity cluster supports HTTP proxy for Azure connectors for performing the following workflows:

- On-demand CloudSpin
- Policy-based CloudSpin
- Cloud-native backup and recovery
- Disaster recovery failback workflows

Manage Azure Source Registration

After registering Azure as a source, you can:

- Update the Azure source configuration.
- Refresh the Azure source details.
- Unregister the Azure source from Cohesity DataProtect.

Update Azure Source Configuration

You can update the registered Azure source to add or remove the Azure services, and subscriptions protected by Cohesity from your Azure source.

To edit an Azure source:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Azure sources and select **Edit**.
3. In the **Edit Azure Source** page, select or unselect the Azure services you need and click **Continue**.
4. Update the required **Authentication Method**:
   - **Azure App** - Azure Application ID, Azure secret key
   - **User Managed Identity** – Managed Identity Client ID
   - **System Managed Identity**

   > **Note:** Managed Identity is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

   Click **Continue**.
5. Add or remove **subscriptions**.
6. Once you edit the Azure source, click **Update**.

Refresh Azure Source

You can refresh the Azure source configuration and fetch the latest changes on your Azure environment.

To refresh the Azure source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Azure source and select **Refresh**.

Unregister Azure Source

If you plan to stop backing up your Azure databases, you can unregister the Azure source from Cohesity DataProtect.

> **Note:** Before you unregister an Azure source from Cohesity DataProtect, you must unprotect all the protected objects in that Azure source.

To unregister the Azure source:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Azure source and select **Unregister**.
3. In the **Unregister Source** dialog, click **Unregister**.

Azure Virtual Machines

Cohesity provides a simple and fast backup, recovery, and data management solution for Microsoft Azure Virtual Machines.

This chapter includes the following topics:

Plan and Prepare for Azure Protection

Review the following prerequisites and considerations of Azure.

Prerequisites

- Azure agent must be installed on the Windows VMs if you need to CloudSpin the VMs to Azure. For more information, see Windows VM Agent.
- The Cohesity cluster must be able to resolve the following URLs by name:
  - Login.windows.net
  - management.azure.com
  - *.blob.core.windows.net

    Ensure to whitelist *.blob.storage.azure.net.
- All storage accounts associated with the subscription must have the relevant permissions, not just the account used to register Azure.

Recovery Prerequisites

To recover files and folder:

- Firewall port 50051 must be open on the target VM. For more information, see Manage Firewall Ports
- The target VM must be reachable with private IP.
- The Cohesity agent must be installed on the target VM.

> **Note:** You can perform files and folder recovery to an Azure VM Linux instance only if the Azure VM Linux instance is of x86_64 architecture.

Considerations

General Considerations

- Cloning is not supported.
- App-consistent backup is not supported.
- LDM for OS disk is not supported.
- VMs with static IP addresses aren't restored with the same static IP.
- All disk sizes must be a multiple of 1 MB.

> **Note:** Adjusting disk size may cause failback issues to on-premises vCenter environments.

Independent Disk Considerations

- Independent disks aren't converted during snapshot-based operations.
- If the boot disk is an independent disk, conversion and cloning aren't supported.

This applies to the following use cases:

- Clone VMs from On-Premises to Azure Cloud Service (On Demand CloudSpin)
- Replicate and Clone VMs in Azure Cloud Edition
- Back Up VMs to On-Prem, Replicate and Failover to Azure Cloud
- Back Up Azure Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server
- Clone VMs to Azure Cloud Using CloudSpin

Tenant-level Registration Considerations

- Tenant-level registration of Azure Cloud Source is supported only on Cohesity Cloud Edition on Azure.
- This feature is not supported on On-premises hardware clusters, Virtual Edition clusters, or Cohesity Cohesity Cloud Edition (Legacy).

Supported Operating Systems (UEFI Boot)

Backup and restore of UEFI-enabled Azure VMs is supported for:

- Windows: 2012, 2016, 2019
- Ubuntu: 14, 16, 18, 20

- RHEL: 6.x, 7.x, 8.x
- CentOS: 6.x, 7.x

File-Level Recovery Considerations

File-level recovery (FLR) isn't supported in the following scenarios:

- Azure VMs using the XFSv5 file system.
- Encrypted VMs.
- When Life Cycle Management (LCM) is enabled on:
  - Azure storage.
  - External targets (Azure Hot Tier).
- End-of-Life (EOL) Linux distributions when required package repositories aren't reachable (for example, YUM repositories for CentOS or RHEL-based distributions). In this scenario:
  - The Cohesity agent required for FLR can't be installed on the target VM.
  - File-level recovery fails.

  > **Note:** Backup and recovery of VMs continue to work as expected. Only file-level recovery is affected.

- File/folder recovery to Azure Linux VMs is supported only on x86_64 architecture.
- Recovery of Azure files and folders for encrypted VMs isn't supported.

  When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.
- File download is supported.

Azure Native Snapshot Considerations

The following configurations aren't supported:

- VMs with Ephemeral disks.
- Shared managed disks.

Supported Disk Types (Native Snapshot API)

- Managed disks:
  - Standard_LRS
  - Premium_LRS
  - StandardSSD_LRS

    - Premium_ZRS
    - StandardSSD_ZRS
- Unmanaged disks

Recovery Considerations and Behavior

- Recovery may fail if a disk contains more than 50 Azure tags.
- Managed disk VMs in a powered-off state may show as 0 bytes, but backup and recovery are supported.
- Recovery of unmanaged disks depends on the target storage container SKU.
- Restore fails if required resources are deleted:
    - Resource group
    - Storage account
    - Storage container
    - Blobs
- VMs restored to a different location:
    - Don’t retain **Availability Set configuration.**
- Disk-level recovery considerations:
    - Supported only for VMs protected using native snapshot protection on Cohesity Cloud Edition on Azure.
    - For spanned volumes (such as LVM on Linux or Dynamic Disks/Storage Pools on Windows), Cohesity restores disks individually. To maintain consistency, recover all disks that belong to the volume in the same recovery job. Recovering only some disks can break the configuration and may require manual reassembly inside the guest OS.
    - If the OS disk of a Linux VM contains /etc/fstab entries for disks that do not use the nofail option, the restored VM may fail to restart

Authentication and Access After Restore

- The restored Azure VM does not support login using username and password present in the backed-up VM. Login using SSH keys is supported. When the restored VM boots up for the first time, Azure images are pre-installed with cloud-init software that locks user accounts as per the default behavior.

  Workaround:

  1. Create an admin user from the Azure portal. Or, log in using the root account (if enabled).
  2. Unlock accounts using:

     ```
     passwd -u <username>
     ```

Encryption Considerations

- Azure Disk Encryption (ADE):
  - Restore to a different location requires manual key replication.
- Azure Storage Service Encryption (SSE):
  - No such consideration.

Networking and Deployment Constraints

- Private endpoint considerations:
  - Not supported for VMs selected using auto-protect.
  - Use manual selection to enable private endpoint.
- Disk size limit with private endpoint:
  - Maximum supported size is 32 TB.

CloudSpin and Platform Notes

- Hyper-V to Azure CloudSpin:
  - Supports BIOS and UEFI boot.
  - UEFI requires Generation 2 VMs.
  - Supported only on managed disks.

Azure Availability Sets

Cohesity supports Azure availability sets. Availability sets provide VM redundancy and availability. For more information about availability sets, see Microsoft Azure documentation.

You can specify the availability set in the following workflows:

- Azure public cloud—Clone and recover VMs
- Azure public cloud using CloudSpin—Clone VMs
- Azure public cloud using Convert and Deploy—Clone VMs
- Azure Stack Hub—Recover VMs

If the VM was part of an availability set during backup and if the VM is cloned or restored to the same location then the new VM is created in the same availability set.

Related Topics

- Set Clone Options—Azure Subscription Source
- Recover VMs to a New Location—Azure
- Clone VMs to Cloud Using Policy Based CloudSpin—Launch VM from the Cohesity Dashboard:

Backup Azure Virtual Machines

This chapter includes the following topics:

Back Up and Recover VMs from Azure Cloud Edition

Cohesity Cloud Edition supports the backup and recovery of Virtual Machines (such as a VM in a vCenter server) to a Cohesity clusterCohesity Cloud Edition running in Azure Cloud service.

VMs can be recovered to a new location or back to the original location as shown in the following figures.

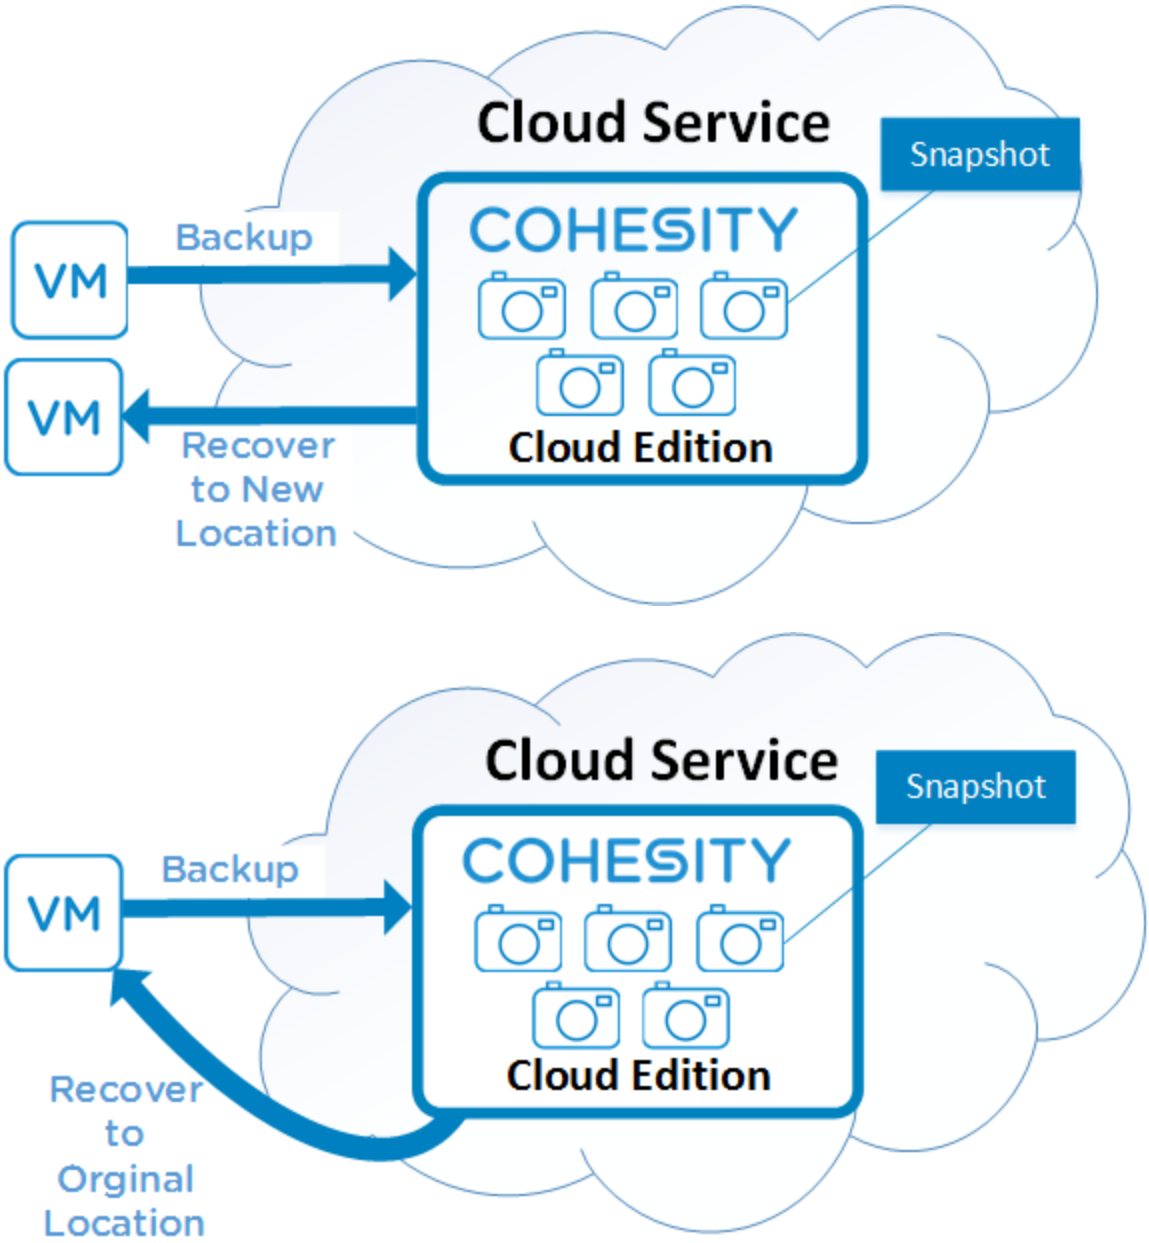

Back Up Azure Cloud VMs to On-Premises Cluster Using Native Snapshot API

As shown by the figure below, the Cohesity cluster supports backing up virtual servers using the native snapshot API of the cloud service. For example, a virtual server Protection Group

backs up VMs from Azure to an on-premises Cohesity cluster. The result of this backup is a snapshot on the on-premises Cohesity cluster. You can recover the VMs from the snapshot on the Cohesity cluster to the cloud service as described in Recover VMs.

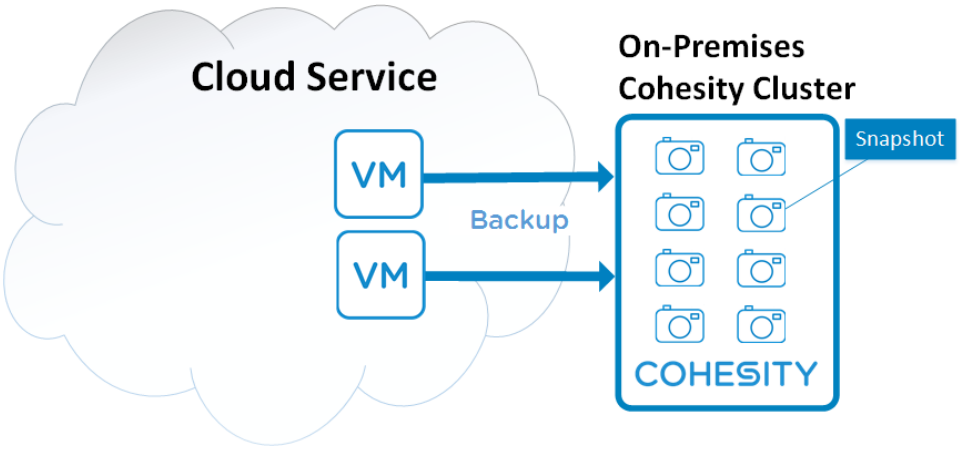

Configure the Backup of a Cloud VM Using Cloud Native APIs

Prerequisites

- Azure Commercial Native Snapshot backup requires the following URLs to be allowlisted in the firewall:
  - Login.windows.net
  - management.azure.com
  - *.blob.core.windows.net
- If you are using a disk private endpoint, then ensure to meet the following prerequisites:
  - The Cohesity cluster must have network connectivity to the subnet used for SAS URL private endpoints.
  - If your Cohesity cluster is located in an on-prem data center, then you must set up VPN or use the Azure express route. If you are using the Cohesity Cloud edition cluster, you must ensure network connectivity between the Cloud edition subnet and the subnet used for SAS URL private endpoints.
  - If you are using a proxy server for the backup of Azure VM using disk private endpoint, then you must perform the following steps:
    1. Create a private DNS zone and link to the virtual network of the proxy server.
    2. Contact Cohesity support and provide the Private DNS Zone name.

3. Azure does not support SAS URI creation of a snapshot with a private endpoint if the size of Azure disk is more than 32 TB.

To configure the backup of Azure VMs using cloud-native APIs, perform the following steps:

1. Register the Azure cloud source. For more information, see Register an Azure Cloud Source at Subscription Level.
2. Create a Protection Group and select **Native Snapshot** as the Protection Group type. While configuring the job, select a cloud source. For detailed instructions, see Add a Protection Group for Azure VM.
3. You can recover the VMs from the snapshot on the Cohesity cluster to the cloud service as described in Recover VMs.

Native Snapshot Process for Creating a Snapshot

Creating a Snapshot for Managed Disks

Default method:

The following is the default native snapshot process for creating a snapshot:

1. A snapshot of each managed disk is taken.
2. (*Applicable only for private endpoint-enabled backup*) A disk access resource with a private endpoint linked is created.
3. (*Applicable only for private endpoint-enabled backup*) The disk access is then attached to the snapshot of the managed disk.

> **Note:** The disk access resource can be shared only by five disks. For every backup, depending on the disk access availability, Cohesity creates a new disk access or uses an existing disk access to attach it to the disk snapshot.

4. A SAS URI (Uniform Resource Identifier) is generated for the created snapshot.
5. The snapshot is then copied to the Cohesity cluster.
   For managed disks, Azure does not provide APIs to perform incremental backups. The entire data is downloaded onto the Cohesity cluster each time you run a Protection Group.

Alternative method:

You can also use the following process for creating the snapshot:

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

1. A snapshot (S1) of each managed disk is taken.
2. (*Applicable only for private endpoint-enabled backup*) A disk access resource with a private endpoint linked is created and is then attached to the snapshot (S1).

   > **Note:** The disk access resource can be shared by up to five disks. For each backup, depending on disk access availability, Cohesity either creates a new disk access or reuses an existing one to attach to the disk snapshot.

3. A SAS URI is generated for the snapshot (S1).
4. Using the SAS URI, the snapshot (S1) is copied to the Cohesity cluster for the first backup, which will be a full backup.

   (*Applicable only for private endpoint-enabled backup*) The disk access resource attached to the snapshot (S1) is deleted.

   This snapshot (S1) on Azure is retained after the first protection run.
5. In the next protection run, a new snapshot (S2) of each managed disk is taken.
6. (*Applicable only for private endpoint-enabled backup*) A disk access resource with a private endpoint linked is created and is then attached to the new and the previous snapshots (S1 and S2).

   > **Note:** Since each disk access resource can be shared by only five disks, multiple disk access resources may be created if there are more than five snapshots.

7. A SAS URI is generated for the new and previous snapshots.
8. The changes between the previous snapshot (S1) and the new snapshot (S2) are calculated and then downloaded to the Cohesity cluster.
9. The SAS URI of both the snapshots is revoked.
10. The previous snapshot (S1) on Azure is deleted. The new snapshot (S2) is retained on Azure for the next incremental protection run.

> **Important:** Multiple Protection Groups can share disk access. However, if no Protection Group has used it for the past 20 minutes, Cohesity will delete it.

Creating a Snapshot for Unmanaged Disks

1. A snapshot of an unmanaged disk is stored as a page blob.
2. The incremental changes between the previous and the current snapshot are calculated and then downloaded to a Cohesity cluster.
3. The previous snapshot is then deleted.

> **Note:** Only one outstanding snapshot (which reflects the latest complete snapshot) is always maintained and therefore to take an incremental backup, the current snapshot is compared with the previous outstanding snapshot and the incremental changes are only considered for backup. The older snapshots are deleted so that the system always maintains only one outstanding snapshot at any time.

Considerations

Review the following considerations of Azure native snapshotting:

- App consistent backup is not supported.
- Cloning is not supported.
- Disks cannot be excluded when using Azure Native Snapshot API.

Back Up Azure Cloud VMs to Cohesity Cloud Edition Cluster Using Native Snapshot API

Cohesity Cohesity Cloud Edition supports backing up Virtual Servers using the native snapshot API of the cloud service. For example, a Virtual Server Protection Group backs up VMs in a cloud service to a Cohesity Cloud Edition Cohesity cluster running in the same cloud service. The result of this backup is Snapshots on the Cohesity Cloud Edition cluster. You can recover the VMs from the snapshot on the Cohesity cluster to the cloud service as described in Recover VMs.

Review the following considerations of Azure native snapshotting:

- App consistent backup is not supported.
- Cloning is not supported.
- Disks cannot be excluded when using Azure Native Snapshot API.

Configure the Backup of a Cloud VM Using Cloud Native APIs

Prerequisites

Verify the following requirements prior to backing up the VMs:

- Azure Commercial Native Snapshot backup requires the following URLs to be allowlisted in the firewall:
  - Login.windows.net
  - management.azure.com
  - *.blob.core.windows.net
- If you are using a disk private endpoint, then ensure to meet the following prerequisites:
  - The Cohesity cluster must have network connectivity to the subnet used for SAS URL private endpoints.
  - If your Cohesity cluster is located in an on-prem data center, then you must set up VPN or use the Azure express route. If you are using the Cohesity Cloud edition cluster, you must ensure network connectivity between the Cloud edition subnet and the subnet used for SAS URL private endpoints.
  - If you are using a proxy server for the backup of Azure VM using disk private endpoint, then you must perform the following steps:
    1. Create a private DNS zone and link to the virtual network of the proxy server.
    2. Contact Cohesity support and provide the Private DNS Zone name.
    3. Azure does not support SAS URI creation of a snapshot with a private endpoint if the size of Azure disk is more than 32 TB.

To configure the backup of Azure VMs using cloud native APIs , perform the following steps:

1. Register the Azure cloud source. For more information, see Register an Azure Cloud Source at Subscription Level.
2. Create a Protection Group and select **Native Snapshot** as the Protection Group type. While configuring the job, select a cloud source. For detailed instructions, see Add a Protection Group for Azure VM.
3. You can recover the VMs from the snapshot on the Cohesity cluster to the cloud service as described in Recover VMs.

Native Snapshot Process for Creating a Snapshot

> **Note:** By default, every backup run is a full backup. Please contact Cohesity Support to enable incremental backup.

Creating a Snapshot for Managed Disks

Default method:

The following is the default native snapshot process for creating a snapshot:

1. A snapshot of each managed disk is taken.
2. (*Applicable only for private endpoint-enabled backup*) A disk access resource with a private endpoint linked is created.
3. (*Applicable only for private endpoint-enabled backup*) The disk access is then attached to the snapshot of the managed disk.
4. A SAS URI (Uniform Resource Identifier) is generated for the created snapshot.
5. The snapshot is then copied to the Cohesity cluster.
   For managed disks, Azure does not provide APIs to perform incremental backups. The entire data is downloaded onto the Cohesity cluster each time you run a Protection Group.

Alternative method:

You can also use the following process for creating the snapshot:

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

1. A snapshot (S1) of each managed disk is taken.
2. (*Applicable only for private endpoint-enabled backup*) A disk access resource with a private endpoint linked is created and is then attached to the snapshot (S1).
3. A SAS URI is generated for the snapshot (S1).
4. Using the SAS URI, the snapshot (S1) is copied to the Cohesity cluster for the first backup, which will be a full backup.

   (*Applicable only for private endpoint-enabled backup*) The disk access resource attached to the snapshot (S1) is deleted.

   This snapshot (S1) on Azure is retained after the first protection run.
5. In the next protection run, a new snapshot (S2) of each managed disk is taken.
6. (*Applicable only for private endpoint-enabled backup*) A disk access resource with a private endpoint linked is created and is then attached to the new and the previous snapshots (S1 and S2) .

   > **Note:** Multiple disk access resource may exist if the there are more than five snapshots.
7. A SAS URI is generated for the new and previous snapshots.
8. The changes between the previous snapshot (S1) and the new snapshot (S2) are calculated and then downloaded to the Cohesity cluster.
9. The SAS URI of both the snapshots are revoked.

10. The previous snapshot (S1) on Azure is deleted. The new snapshot (S2) is retained on Azure for the next incremental protection run.

Creating a Snapshot for Unmanaged Disks

1. A snapshot of an unmanaged disk is stored as a page blob.
2. The incremental changes between the previous and the current snapshot are calculated and then downloaded to a Cohesity cluster.
3. The previous snapshot is then deleted.

> **Note:** Only one outstanding snapshot (which reflects the latest complete snapshot) is always maintained and therefore to take an incremental backup, the current snapshot is compared with the previous outstanding snapshot and the incremental changes are only considered for backup. The older snapshots are deleted so that the system always maintains only one outstanding snapshot at any time.

Back Up Azure Cloud VMs Using Cloud Snapshot Management

Cohesity cluster uses Cloud Snapshot Management (CSM) APIs to backup the Azure VMs (Managed VMs) and AWS Cloud services to the same cloud. To enable this feature, you must select the Snapshot Manager as the job type while creating the Virtual Server Protection Group. When this Protection Group is executed, the VMs' snapshots are created on the same cloud service. The CSM snapshots are stored in the same Azure subscription, but can be used to recover to a different subscription.

Cloud Snapshot Management Process for Creating the Snapshot

For Azure Managed Disks:

1. Fetch the information about the VM using the GET API.

   This API takes `subscriptionId`, `resourceGroupName`, and `vmName` as an argument and returns the information about the VM (such as disks, network interfaces, and so on).
2. After the VM information is fetched, snapshots of the VM's disk is created on the same cloud service using the Snapshots API.

You can also recover these snapshots. For more information, see Recover VMs.

Considerations

- The backup of unmanaged VMs is not supported.
- Incremental snapshot of managed disks is not supported

- App consistent backup is not supported.
- Cloning is not supported.
- File download is not supported.

Configure the Back Up of a Cloud VM Using CSM

1. Register the Azure cloud source. For more information, see Register an Azure Cloud Source at Subscription Level.
2. Create a Protection Group and select **Snapshot Manager** as the Protection Group type. While configuring the job, select a cloud source. For detailed instructions, see Add a Protection Group for Azure VM.
3. You can recover the VMs from the snapshot on the Cohesity cluster to the cloud service as described in Recover VMs.

Cohesity creates a snapshot of VMs for Azure on the same Cloud service.

Add a Protection Group for Azure VM

In Cohesity, Azure VMs are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group.

To protect Azure VMs:

1. Select **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **Virtual Machines**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the Azure source you want to back up. You can also register a new Azure source with the Cohesity cluster by clicking **Register Source**. For more information, see Register an Azure Cloud Source at Subscription Level.
4. In the **Objects** section, choose the Azure VM using one of the following views:
   - **List View**: Use this option to select the Azure VMs for protection from a flat, searchable list.
   - **Hierarchy View**: Use this option to select Azure VMs for protection at various hierarchy levels (Tenant level, Subscription level, resource group level).

     Optionally:

     a. You can click the auto-protect icon at the **Tenant** level, **Subscription level, Resource group** level to automatically protect newly added VMs.

b. You can exclude VMs from auto-protection within an account, region, or cluster by clicking the auto-protected icon ( ) next to the VM. The VMs excluded from auto-protection are displayed with the exclude ( ) icon.

- **Tags**: Use this option to select the VMs for protection based on tags. Select the checkbox of a tag, and then select Select All Child Objects to include all tables with that tag for protection.

  Optionally:

  - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added VMs with this tag.
  - You can exclude the Azure VMs from auto-protection within a tag by clicking the auto-protect icon ( ) next to the table. The VMs excluded from auto-protection are displayed with the exclude ( ) icon.
  - You can exclude VMs with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the excluded icon.

5. (*Supported only on Cohesity Cloud Edition on Azure*) Optionally, you can exclude any disk of the Azure VM from protection. To exclude a disk:
   1. Click the **All Volumes** edit icon displayed next to the VM of which you want to exclude the disk.
   2. Unselect the disks that you want to exclude from protection.

      > **Note:** You cannot exclude the root disk from protection.

   3. Click **Save** > **Continue**.

   > **Note:**
   > - You cannot exclude disk at the object level for auto-protected VMs. To exclude disk for auto-protect VMs, use the **Volume Exclusion Settings** option in the Additional Settings.
   > - Exclusion of disk is supported only for the **Native Snapshot** Protection Group Type.

6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.

7. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
8. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
9. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, click **Protect**.
10. Select a **Protection Group Type**:
    - **Native Snapshot**—Back up VMs in a cloud service using the native snapshot API of the cloud service.
    - **Snapshot Manager**—Backup VMs in a cloud service using the Cloud Snapshot Manager APIs of the cloud service.
11. Optionally, from the **Policy** drop-down list, you can change the policy you selected earlier policy or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
12. In the **Storage Domain** field, select the storage domain where you want to store the snapshots.
13. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
14. In the **SLA** field, select how long the administrator expects a Protection Group run to take.
    - **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
    - **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run.
15. If you need to change any of the advanced settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
16. Click **Protect**.

The protection of Azure VMs is now active and running.

Additional Settings

If you select..	Then..
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.

If you select..	Then..
End Date	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

If you select..	Then..

Azure SAS URL Type
 	
Note: This option is available only if you have selected the **Native Snapshot** option as the Protection Group Type.
	For the backup of Azure VMs, Cohesity uses the snapshots of Azure managed disks that are accessible through SAS URLs. Select one of the following endpoint options to access data from the SAS URL:
	• **Public Endpoint**: Select this option to access snapshots over a public network.
	• **Private Endpoint**: Select this option to access snapshots using a private IP address from your virtual network. Provide the following details for configuring the private endpoint:
	  • **Region**: Select the region of the VM you have selected for backup.
	  • **Virtual Network**: Select a virtual network in Azure. Ensure this virtual network is connected to the Cohesity cluster.
	  • **Subnet**: Select a subnet of the virtual network.
	
	Note:
	• In the native snapshot backup process for managed disks, a disk access resource is created and attached to the snapshot. Since the disk access requires an IP address, ensure that the selected subnet has sufficient IP availability. For more see Back Up Azure Cloud VMs to On-Premises Cluster Using Native Snapshot API
	• If you have selected VMs from multiple regions, then you must select at least one virtual network and subnet per region. Click the add icon to add regions and their virtual networks and subnets.
	• The Private Endpoint option is not supported for the VMs selected using auto-protect. To use a private endpoint, you must manually select the VMs for protection.
	• The disk size limit of an Azure VM using a private endpoint is 32 TB.
	
	Cohesity recommends using dedicated subnets for Azure SAS URL endpoint.

If you select..	Then..
QoS	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. Backup SSD policy and reduces the SSD wear-out compared to the Backup SSD policy. • **Backup Auto**: (Applicable only for C6K Platform) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the • **TestAndDev High**: The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
Indexing	This option is toggled on by default to index the data of the Protection Group. Indexing improves files and folders search. Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.

If you select..	Then..
Volume Exclusion Settings (*Supported only on Cohesity Cloud Edition on Azure*) > **Note:** This option is available only if you have selected the **Native Snapshot** option as the Protection Group Type.	Use the following tag keys to construct a boolean expression for excluding disks based on disk tags: • Name • automation • cohesity::helios::uuid • email • environment • expiry • location • non-business-hrs • purpose • rigel • squad • test • user • weekend For example, the expression, `(name = net_volume AND environment IN (qa, dev)) OR sqaud = eng`, excludes all disks from protection for which the disk tags meet the above condition. > **Note:** Exclusion of disks enabled at the object level will take precedence over the exclusion configured here.
Cancel Run at Quite Start Time	Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. If the option is toggled off, then after a protection run starts, it continues to execute even when a quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period.

If you select..	Then..
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: a. High b. Medium c. Low
Description	Enter a brief description about the Protection Group.

Azure Stack Hub

Cohesity supports image-level backups of VMs in Azure Stack Hub. Cohesity supports the following Azure Stack Hub versions:

- 2102
- 2008
- 2005

Note the following terminology:

- Source cluster—Cohesity cluster on which the Azure Stack Hub VM is registered and protected.
- Target cluster—Remote Cohesity cluster to which the Azure Stack Hub VM backup data is replicated. For more information about replication, see About Replication Setup.

You can perform the following tasks:

- Back up and restore Azure Stack Hub VMs.
- Restore an individual file from the backup. You must install the Cohesity agent on the VM to perform file-level recovery.
- Replicate and archive backups. You can restore VMs, files, or folders only from replicated or archived copies.
- Back up and restore Azure Stack Hub infrastructure data.
- Restore Azure Stack Hub VMs
    - On one Cohesity cluster, you can restore VMs from:
        - Azure Stack Hub to Azure Standard
        - Azure Standard to Azure Stack Hub
        - Azure Stack Hub that uses AD to Azure Stack Hub that uses AD FS
        - Azure Stack Hub that uses AD FS to Azure Stack Hub that uses AD
    - You can restore VMs from:
        - Azure Stack Hub in the source cluster to Azure Standard in the target cluster
        - Azure Standard in the source cluster to Azure Stack Hub in the target cluster
        - Azure Stack Hub that uses AD in the source cluster to Azure Stack Hub that uses AD FS in the target cluster
        - Azure Stack Hub that uses AD FS in the source cluster to Azure Stack Hub that uses AD in the target cluster
- Failover Azure Stack Hub VMs
    - On one Cohesity cluster, you can failover VMs from:
        - Azure Stack Hub to Azure Standard
        - Azure Standard to Azure Stack Hub
    - You can failover VMs from:
        - Azure Stack Hub in the source cluster to Azure Standard in the target cluster
        - Azure Standard in the source cluster to Azure Stack Hub in the target cluster
- Migrate Azure Stack Hub VMs
    - On one Cohesity cluster, you can migrate VMs from:
        - Azure Stack Hub to Azure Standard
        - Azure Standard to Azure Stack Hub

- You can migrate VMs from:
  - Azure Stack Hub in the source cluster to Azure Standard in the target cluster
  - Azure Standard in the source cluster to Azure Stack Hub in the target cluster

When registering or editing an Azure Cloud source, specify the authentication type, region, and domain name for Azure Stack Hub. For more information, see the following topics:

- Register an Azure Cloud Source at Subscription Level
- Azure

Protect Azure Stack Hub Infrastructure

Cohesity enables you to manage Azure Stack Hub infrastructure backups. You can backup and restore the data to an existing Azure Stack Hub or another Azure Stack Hub in case of a disaster. Azure Stack Hub administrators can backup Azure Stack Hub infrastructure data to an SMB share. For more information, see Recover data in Azure Stack Hub with the Infrastructure Backup Service in the Microsoft Azure documentation.

To protect Azure Stack Hub infrastructure data, complete the following steps:

1. Create a View to backup Azure Stack Hub infrastructure data. For more information, see Create Views Using a Predefined Template.
2. Create a protection job for the View. For more information, see Add or Edit a Protection Group for a View.

Considerations

- Incremental backup is not supported.
- Direct upload to Azure managed disks is not supported.
- Generation 2 VMs and encrypted VMs are not supported.
- Cloning of Hyper-V VMs to Azure Stack Hub is not supported.

Recover Azure Virtual Machines

This chapter includes the following topics:

Recover Files or Folders for Azure Native Snapshots to the New Server

Prerequisites

- Firewall port 50051 must be open on the target VM. For more information, see Manage Firewall Ports
- The target VM must be reachable with private IP.
- The Cohesity agent must be installed on the target VM.

> **Note:** You can perform files and folder recovery to an Azure VM Linux instance only if the Azure VM Linux instance is of x86_64 architecture.

Recover Files and Folders to New Server

Perform the following steps to recover files and folders to the new server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on Cohesity.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on Cohesity.

If you select..	Then..
Files and Folders	a. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. b. From the list of search results, select the required files or folders to recover. c. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: i. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. ii. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. iii. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. d. Go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display

If you select..	Then..
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. 8. Go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Server Type	Volume Type	

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next: Recover Options**: If you select this option, then continue with the step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task completes, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':', it is replaced with '_COLON_' in the archive.
     Example:

- The original name is "a:b" then the archived name is "a_COLON_b"
- The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
Example:

- The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **New Server** to recover the files or folders to a new server.
6. **Source**: From the **Source** drop-down menu, select a registered Azure source to which you want to recover the files and folders. You can also register a new source by clicking **Register Source**.
7. **Target**: From the **Target** drop-down menu, select an Azure VM to which the files and files are to be recovered.
8. **Recover to**: In the **Recover to** field, change the default location to which the files or folders are recovered. By default, the files and folders are recovered to the /**tmp**/ directory for VMs.

   For NAS, an example path on an alternate location could be /tmp/recover. You can recover encrypted files/folders to their original location only.

   Example for Windows alternate location:

   - D:\cohesity_restore\1.txt
   - D:\cohesity_restore\bar

   User Account Requirements for VM Protection may need to be met to recover files to a Windows VM.

   If recovering to a standalone Hyper-V host or cluster, you must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.
9. Optional. Edit the default settings of the following **Recovery options**:

   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

- C:\restore\1.txt
- C:\restore\cohesity<taskid>_0\1.txt
- C:\restore\bar
- C:\restore\cohesity<taskid>_1\bar

If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

- **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.
- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

10. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Files or Folders for Azure Native Snapshots to the Original Server

Prerequisites

- Firewall port 50051 must be open on the target VM. For more information, see Manage Firewall Ports
- The target VM must be reachable with private IP.
- The Cohesity agent must be installed on the target VM.

> **Note:** You can perform files and folder recovery to an Azure VM Linux instance only if the Azure VM Linux instance is of x86_64 architecture.

Recover Files and Folders to Original Server

Perform the following steps to recover files and folders to the original server:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Files or Folders**.

The Recover Files or Folders page appears.

3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on Cohesity.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on Cohesity.

If you select..	Then..
Files and Folders	a. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. b. From the list of search results, select the required files or folders to recover. c. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: i. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. ii. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. iii. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. d. Go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display

If you select..	Then..
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. 8. Go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Server Type	Volume Type	

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next: Recover Options**: If you select this option, then continue with the step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task completes, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':', it is replaced with '_COLON_' in the archive.
     Example:

- The original name is "a:b" then the archived name is "a_COLON_b"
- The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
Example:

- The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **Original Server** to recover the files or folders to a new server.
6. **Recover to Original Path**: To recover to a different location in the original server, disable the **Recover to Original Path** option, and then provide the location to which the files or folders are to be recovered in the **Recover To** field. By default, the files and folders will be recovered to the original location.
7. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

     - C:\restore\1.txt
     - C:\restore\cohesity<taskid>_0\1.txt
     - C:\restore\bar
     - C:\restore\cohesity<taskid>_1\bar

     If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
   - **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.

- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

8. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Azure VMs to a New Location

You can recover the VM files to an alternate datastore and create new instances of the VMs in a new location of a registered Azure source.

Before you can recover VMs to its new location, Snapshots of those volumes must exist on the Cohesity cluster. To create Snapshots, create and run a Protection Group.

To recover VMs to a new location:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > VMs**.

   The Virtual Machines page appears.
3. Search for Protected Object (VM) or Protection Group that contains the snapshots to recover from by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more VMs or Protection Groups that contain the snapshots to recover.

> **Note:**
> - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
> - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

5. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the Protection Run drop-down menu, select the snapshot from which you want to recover the VMs. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a..	Then..
VM	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.

   From the **Recover To** options, select **New Location** and provide the following details to recover the VM to an alternate datastore and create new instances of the VMs in a new location of a registered Source.

   a. **Source:** Select a registered Source to recover the VMs to.
   b. **Resource Group, Storage Account:** Select the corresponding Resource Group and Storage Account to which the recovery has to be done. The storage account you select must be in the same region as the resource group. Any storage account not in the same region as the parent resource group will not be displayed as an option in the recovery workflow.
   c. **Storage Container:** Select the storage container in which page blobs will be created.
   d. **Compute:** Select the VM size type such as **Standard_D1** to use when recovering the VMs.
   e. **Availability Set:** Optionally, select the availability set to which the VM has to be recovered. For more information, see Azure Availability Sets.
   f. **Virtual Network:** Select the virtual network where the VM must be placed.
   g. **Subnets:** Subnets in the virtual network where VMs must be placed.

7. Optional. Edit the default settings of the following **Recovery Options**:

Recovery Options	Description
Rename	Add Prefix and/or Suffix strings to the names of the new VMs created by this task. For example, you could specify Prefix Test_ and Suffix _QA to the VM name WinOS and when the VM is recovered, the new VM that is created is called Test_WinOS_QA.

Recovery Options	Description

Azure SAS URL Type	For recovering Azure VMs, Cohesity uses the snapshots of Azure managed disks that are accessible through SAS URLs. Select one of the following endpoint options to access data from the SAS URL:
	• **Public Endpoint**: Select this option to access snapshots over a public network
	• **Private Endpoint**: Select this option to access snapshots using a private IP address from your virtual network. Provide the following details for configuring the private endpoint:
	  • **Region**: Select the region of the VM you have selected for recovery.
	  • **Virtual Network**: Select a virtual network in Azure. Ensure this virtual network is connected to the Cohesity cluster.
	  • **Subnet**: Select a subnet of the virtual network.
	Cohesity recommends using dedicated subnets for the Azure SAS URL endpoint.

Power State	To ensure that the restored Azure VMs are functioning properly, the VMs are powered on by default after the restoration process. However, you can enable the **Power Off** option if you prefer to keep the VMs powered off once Cohesity verifies their functionality. If you are recovering VMs from a Storage Domain with Cloud Tier enabled, Cohesity recommends keeping the VMs powered off.
Continue on Error	Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cohesity cluster will still attempt to create the other VMs. If this option is disabled, then the recovery task is discontinued if VM creation fails. However, if the VM creation succeeds, but VM storage migration fails, the task continues. By default this option is disabled.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
Task Name	Change the default name of the recovery task.

8. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Azure VMs to Original Location

Before you can recover VMs to its original location, snapshots of those volumes must exist on the Cohesity cluster. To create Snapshots, create and run a Protection Group.

To recover VMs to its original location:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > VMs**.

   The Virtual Machines page appears.
3. Search for a Protected Object (VM) or Protection Group that contains the snapshots to recover from by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more VMs or Protection Groups that contain the snapshots to recover.

   > **Note:**
   > - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
   > - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.
5. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the Protection Run drop-down menu, select the snapshot from which you want to recover the VMs. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a..	Then..
VM	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the VM to the original location.
8. Optional. Edit the default settings of the following **Recovery Options**:

Recovery Options	Description
Rename	Add Prefix and/or Suffix strings to the names of the new VMs created by this task. For example you could specify Prefix `Test_ and Suffix _QA` to the VM name WinOS and when the VM is recovered, the new VM that is created is called `Test_WinOS_QA`.

Azure SAS URL Type	For recovering Azure VMs, Cohesity uses the snapshots of Azure managed disks that are accessible through SAS URLs. Select one of the following endpoint options to access data from the SAS URL:
	• **Public Endpoint:** Select this option to access snapshots over a public network.
	• **Private Endpoint**: Select this option to access snapshots using a private IP address from your virtual network. You can either enable the "**Use the same SAS URL details a Protection**" option or provide the following details for configuring the private endpoint:
	  • **Region**: Select the region of the VM you have selected for recovery.
	  • **Virtual Network**: Select a virtual network in Azure. Ensure this virtual network is connected to the Cohesity cluster.
	  • **Subnet**: Select a subnet of the virtual network.
	If you have selected VMs from multiple regions, then you must select at least one virtual network and subnet per region. Click the add icon to add regions and their virtual networks and subnets.
	Cohesity recommends using dedicated subnets for the Azure SAS URL endpoint.

Power State	To ensure that the restored Azure VMs are functioning properly, the VMs are powered on by default after the restoration process. However, you can enable the **Power Off** option if you prefer to keep the VMs powered off once Cohesity verifies their functionality. If you are recovering VMs from a Storage Domain with Cloud Tier enabled, Cohesity recommends keeping the VMs powered off.

Recovery Options	Description
Continue on Error	Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cohesity cluster will still attempt to create the other VMs. If this option is disabled, then the recovery task is discontinued if VM creation fails. However, if the VM creation succeeds, but VM storage migration fails, the task continues. By default this option is disabled.
Task Name	Change the default name of the recovery task.

9. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Azure Disk

You can recover Azure VM disks in the following ways:

- **Recover to the original VM**: If you are recovering the virtual disk to the original VM, you can overwrite the disk in its original VM or attach it as a new disk.
- **Recover to a different VM**: If recovering the virtual disk to a different VM, the disk is attached as a new disk in the VM you select

> **Note:** Disk-level recovery for Azure VMs is supported only for VMs protected using the native snapshot protection method on a Cohesity Cloud Edition cluster deployed on Azure. For the list of permissions required for disk-level recovery, see Required Permissions for Azure Workloadss.

Before you initiate the recovery of Azure disks, perform the prerequisites review the considerations.

To recover virtual disks:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > Virtual Disks**.

   The Virtual Disks page appears.
3. Search for the Azure VM whose virtual disks you want to recover by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

A list of VMs displays.

You can optionally filter the search results by source, Protection Group, Storage Domain, server type, or date.

4. From the list of search results, select the VM whose virtual disks you want to recover.

   By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the edit icon.

   The *Edit recovery point for <VM_name>* page appears.

   In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in **Timeline** mode by clicking on the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:
     1. Select the **List** tab from the top-right corner of the *Edit recovery point for <VM_name>* page.
     2. By default, the page displays the snapshots available in the last one month.
     3. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
     4. Select the snapshots to be recovered.

        By default, the snapshot is recovered from the Azure Storage associated with the Cohesity storage domain. You can also select snapshots archived to the cloud by clicking on the cloud icon if available in the Location option.
     5. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the **Timeline** mode:
     1. From the **Choose a date** drop-down, select a date.
     2. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
     3. Select the snapshot to be recovered by clicking on the corresponding blue dot.

        By default, the snapshot is recovered from the Azure Storage associated with the Cohesity storage domain. You can also select snapshots archived

to the cloud by clicking on the cloud icon if available in the Location option.

4. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on where you want to restore the disk and then provide the required information:
   - **Original Location**: Recover the disk to the original virtual machine.
   - **New Location**: Recover the disk to a different virtual machine.
     a. From the **Source** drop-down, select a registered source having the VM on to which you want to recover the virtual disks. You can also register a new source by clicking **Register Source** from the **Source** drop-down menu.
     b. From the **Target VM** drop-down, select a VM on which you want to recover the disk.
8. From the **Target Resource Group**, select the resource group where the recovered disks will be created.
   - This field is not required if all selected disks are recovered in overwrite mode.
   - In a mixed scenario (some disks overwrite and some do not), the target resource group applies only to non-overwrite disks.
9. Under **Source Disks**, select one or more disks to recover. Each disk represents a managed disk from the source VM. You can select both OS disk and data disks.
10. (Applies only for recovering the VM disks to **Original Location**) For each selected disk:
    - Set **Overwrite Original** to **Yes** to overwrite the existing disk.
    - Set **Overwrite Original** to **No** to attach the recovered disk as a new disk.

Note:

- If one or more disks are configured to overwrite the original disk, the VM is automatically powered off before the recovery starts.
- When an OS disk is recovered:
  - If **Overwrite Original** set to **Yes**, then:
    - The VM boots from the recovered OS disk.
    - The recovered OS disk is created in the same resource group as the original OS disk.
    - By default, Azure retains the original OS disk in the subscription while the VM runs using the recovered OS disk.
  - If **Overwrite Original** set to **No**, then:
    - The OS disk is restored as a new, bootable managed disk in the target resource group.
    - The disk is not automatically attached to the VM.
    - To use the recovered OS disk, you must manually swap the OS disk from the Azure portal or attach it to a new VM.

11. For each selected disk, provide a unique name in Restored Disk Name that complies with Azure managed disk naming rules.
    - You do not need to provide a name for disks that are overwritten.
    - Overwritten data disks retain the original disk name.
    - Overwritten OS disks receive an automatically generated name.
12. In the **Recovery Options** section, configure the following options:
    1. **Power State**: Enable the **Power off VMs before recovery** option to ensure that the VM is powered off before the virtual disks are recovered. Enable the **Power on VMs after restore** option to power on the VM after the virtual disks are recovered.
    2. **Azure SAS URL Type**: For recovering Azure VM disks, Cohesity uses the snapshots of Azure managed disks that are accessible through SAS URLs. Select one of the following endpoint options to access data from the SAS URL:

- **Public Endpoint**: Select this option to access snapshots over a public network
- **Private Endpoint**: Select this option to access snapshots using a private IP address from your virtual network. Provide the following details for configuring the private endpoint:
    i. **Region**: Select the region of the VM you have selected for recovery.
    ii. **Virtual Network**: Select a virtual network in Azure. Ensure that this virtual network is connected to the Cohesity cluster.
    iii. **Subnet**: Select a subnet of the virtual network.

Cohesity recommends using dedicated subnets for the Azure SAS URL endpoint.

13. **Task Name**: (*Optional*) Review or modify the Task Name. By default, a task name is auto generated (for example, *Recover_Azure_Disk_<date_time>*)
14. Click **Recover**.

Disaster Recovery of VMs Using Azure Cohesity Cloud Edition

To provide disaster recovery failover and failback, Cohesity supports the Cloud Migration of VMs between hypervisors servers (VMware vCenter), Azure Cloud Services, and Cohesity Cloud EditionCohesity clusters. For Hyper-V, Cohesity supports only failover to Azure.

Disaster recovery with Cohesity Cloud Edition consists of the following two use cases:

- Back Up VMs to On-Prem, Replicate and Failover to Azure Cloud
- Back Up Azure Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server

Back Up VMs to On-Prem, Replicate and Failover to Azure Cloud

To provide disaster recovery failover and failback, Cohesity supports the Cloud Migration of VMs between hypervisors servers (VMware vCenter), Azure Cloud Services and Cohesity Cloud EditionCohesity clusters. For Hyper-V, Cohesity supports only failover to Azure. Cohesity Cohesity Cloud Edition supports the replication of VMs captured on an on-premises Cohesity cluster to Cohesity Cloud Edition as shown by the example figure below. Starting from the left in the following figure, the use case shows how Protection Groups running on the on-premises Cohesity cluster captures Snapshots of VMs and stores these Snapshots on the on-premises Cluster. The Protection Policies associated with these Jobs, are configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity Cloud EditionCohesity cluster.

From the Snapshots stored on the Cohesity Cloud Edition, you can failover VMs. The failover process creates new VMs running in the Azure Cloud Service. As part of the Failover process, the VMs are converted and spun up in the Cloud Service.

Cohesity supports the failover of VMware VMs with UEFI boot mode enabled to Azure.

Failover to Azure supports the following UEFI boot mode enabled operating systems:

- Windows Server 2012, 2016, 2019

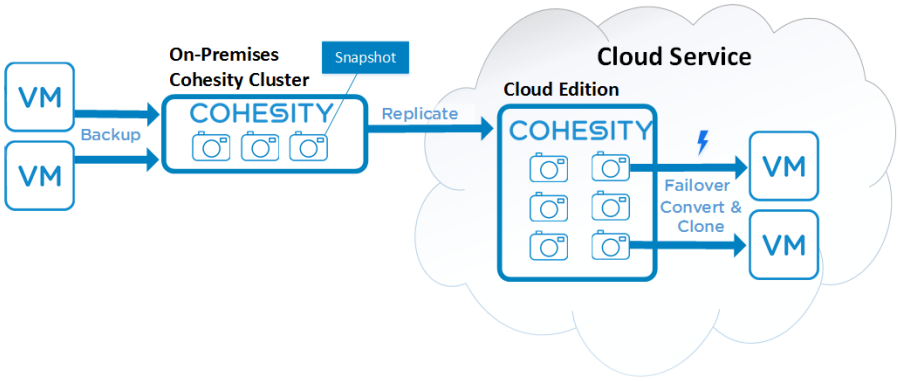

> **Important:** To support the failback (recovery) of these VMs back to the original Hypervisor server (such as VMware vCenter), Cohesity agents must be installed on the Windows VMs prior to backing up the snapshots on the on-premises Cohesity cluster. Failback is described in Back Up Azure Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server.

Verify the following requirements prior to backing up the VMs:

- If you are planning to perform failback (recovery) of the backed up VMs back to the Hypervisor server from the Cloud Service, then:
  - Before registering the VMs of the Hypervisor server with the on-premises Cohesity cluster, you must install Cohesity agent on the VMs if you are planning to back up the cloud VM back to the Cohesity cluster using the **Cohesity Protection Service** Protection Group type.
  - You do not need to install the Cohesity agent on the VMs if you are planning to backup the cloud VM back to the Cohesity cluster using native snapshots.

    > **Note:** You must install Cohesity agent on the Windows VMs before registering the VMs of the Hypervisor server with the on-premises Cohesity cluster. The Cohesity agent is required to bring all the disks on the VMs online.

- Failover and Failback of VMware VM to Azure supports the following operating systems:
  - Microsoft Windows Server 2012, Microsoft Windows Server 2016, and Microsoft Windows Server 2019, Microsoft Windows Server 2022
  - CentOS 6.0+ 64-bit, 7.0 - 7.9 64-bit, 8.0 64-bit

    - RHEL 6.7+ 64-bit, 7.0 - 7.9 64-bit, 8.0 - 8.2 64-bit
- VMware vSphere 5.5, 6.0, 6.5, and 6.7 is supported for cloning to Azure.
- Cloning Snapshots to an Azure Cloud is supported for VMware and Hyper-V VMs with the following operating systems:
    - Linux CentOS 7 on LVM
        - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
        - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

          ```
          UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
          defaults,nofail 0 0
          ```

          For detailed instructions, see Edit /etc/fstab to use UUID.
    - Windows
        - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
        - For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
        - When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. The Cohesity agents installed must be of the same version as the Cohesity Cluster version. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.
        - Due to a Windows issue in builds 14393 and build 15063, a deployed Windows 10 or Windows 2016 VM may not be connected to the network. Resolve this by installing fix KB4073562 for build 14393 or KB4016240 for build 15063 or by updating Windows. For more information, see Cannot connect remotely to a Windows 10 or Windows Server 2016 VM in Azure because of netvsc.sys.
        - Make the disks online by default using the following SAN policy commands:
            - `diskpart`
            - `san policy=OnlineAll`

          For more information, see the SAN Policy documentation.

- Red Hat Enterprise Linux
  - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
  - For the list of supported RHEL versions, see Supported Software.
  - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

    ```
    UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4 defaults,nofail 0 0
    ```

    For detailed instructions, see Edit /etc/fstab to use UUID.
- The disk size of the VMware VMs must be a multiple of 1 GB.
- Cohesity does not support the CloudSpin of VMs whose OS Disk is formatted with XFS version 5 to Azure. This limitation applies to both Hyper-V and VMware VMs, regardless of their operating system type.
- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to Azure.

Cloning Hyper-V VMs to Azure

- VMs that have one to five disks can be cloned. Each disk must be smaller than 2 TB.
- SUSE Linux VMs are *not* supported.
- All the Hyper-V versions supported by Cohesity for backup is supported for cloning to Azure. For the list of supported Hyper-V versions, see Supported Software.
- Generation 2 VMs are not supported by all VM sizes in Azure. For information on Generation 2 supported VM series, see Microsoft document.
- Cloning of Hyper-V VMs to Azure Stack Hub is not supported.

- For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.

To back up VMs to an on-premises Cluster, replicate and failover to Cohesity Cloud Edition:

1. Register the Hypervisor server with the on-premises Cohesity cluster
2. Set up replication between the on-premises Cohesity cluster and the Cohesity Cloud Edition Cohesity cluster. For instructions, see Replication and Remote Access Setup.

3. On the on-premises Cohesity cluster, create a Policy that replicates to a Cohesity Cloud Edition Cohesity cluster as shown in the top figure. For instructions, see Create or Edit a Standard Policy.
4. On the on-premises Cohesity cluster, create a Virtual Server Protection Group (such as `Win2k12`) that has **Cloud Migration** enabled. For detailed instructions, see Add or Edit a Protection Group for Virtual Servers. Associate the Policy created in the previous step with this Protection Group.

   When selecting Objects to be protected by a Protection Group, you can only select VMs with VMware tools installed, Remote Desktop enabled, and for Windows VM Cohesity agents also must be installed.

   When you create a Protection Group on the on-premises Cohesity cluster that replicates to a Cohesity Cloud Edition Cohesity cluster, after the first Snapshot is replicated to the Cohesity Cloud Edition Cluster, an Inactive copy of the Protection Group is created on the Cohesity Cloud Edition Cluster that received the Snapshot. The new Inactive Protection Group on the Cohesity Cloud Edition Cohesity cluster has the same name as the Protection Group on the on-premises Cohesity cluster, for example: `Win2k12`. The VM must be powered on and connected to the network when the Protection Group is running.
5. You can failover from the Snapshots that were replicated to the Cohesity Cloud Edition Cluster. The failover process creates new VMs running in the cloud service (such as AWS or Azure). As part of the Failover process, the VMs are converted and spun up in the cloud service.

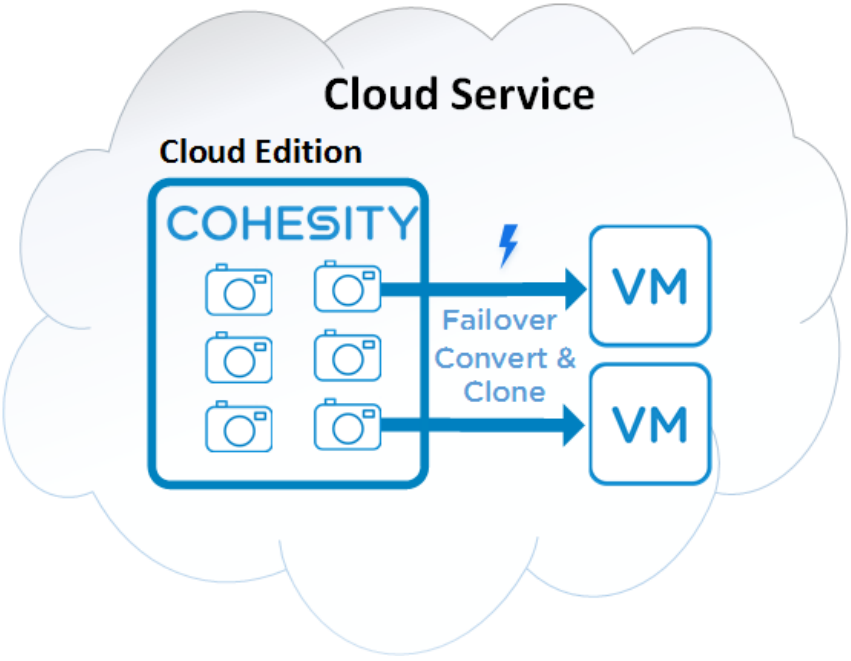

To failover from the Snapshots:

a. On the Cohesity Cloud Edition Cohesity cluster that contains the replicated Snapshots, select **Data Protection > Protection** and browse for the Inactive Protection Group that was automatically created.

b. Click on the inactive Protection Group name.

c. Click the action icon (⋮) next to the Protection Group and select **Failover**.

d. From the **Failover to Source** drop-down, select the name of the registered Cloud Source where the VMs will be cloned to.

e. Click **Continue to Clone**.

f. Follow the instructions provided in Set Clone Options. You can also add custom tags to all the resources created during failover.

A clone task is created and the VM is created in the Azure Cloud Service and is powered up if that option was selected.

6. Follow the steps in the next part of the Disaster Recovery workflow which is documented in Back Up Azure Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server.

You can also use CloudSpin (test/dev clone) workflow to failover individual VMs within the Protection Group instead of failovering all VMs in the Protection Group. For more information, see Test & Dev Using CloudSpin.

Back Up Azure Cloud VMs, Replicate to On-Prem and Failback to Hypervisor Server

To provide disaster recovery failover and failback, Cohesity supports the Cloud Migration of VMs between hypervisors servers (VMware vCenter), Azure Cloud Services and Cohesity Cloud EditionCohesity clusters.

> **Note:**
> - For Hyper-V, Cohesity supports only failover to Azure. Failback is not supported.
> - Failover of VMs that originated in Azure (Azure-native VMs) to an on-premises VMware vCenter is not supported. This topic describes the failback of VMs that were originally backed up from an on-premises VMware vCenter server and subsequently failed over to Azure using the process described in Back Up VMs to On-Prem, Replicate and Failover to Azure Cloud.

This topic describes the second use case of disaster recovery using Cloud Edition—how to backup cloud VMs to the Cloud Edition cluster, replicate to the on-premise cluster, and then failback to the Hypervisor server.

For information on how the VMs are backed up in the Azure cloud service, see the first process of disaster recovery using Cloud Edition documented in Back Up VMs to On-Prem, Replicate and Failover to Azure Cloud.

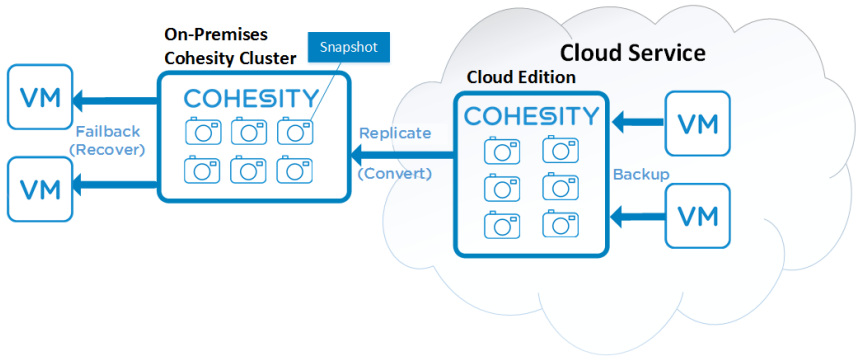

As illustrated in the figure above, on the Cohesity Cloud Edition running in the Cloud Service, a Protection Group using native snapshots or Cohesity Protection Service backs up the cloud VMs (as indicated by the backup arrow in the figure). The policy associated with the Protection Group specifies that these Snapshots are replicated to the on-premises Cohesity cluster (as indicated by the Replicate arrow). From those Snapshots stored on the on-premises Cohesity cluster, you can failback (recover) the VMs to their original on-premises VMware vCenter server or an alternate VMware vCenter server (as indicated by the Failback (Recover) arrow in the figure). This restores the VMs to the on-premises hypervisor environment from which they originally failed over to the Azure Cloud Service.

To back Up Cloud VMs to Cohesity Cloud Edition, replicate to a on-premises Cohesity cluster and failback to a hypervisor server:

Verify the following requirements prior to failback:

- The VMs that are backed up to the Cloud Service must have VMware Tools and Remote Desktop installed before the original Protection Group captures the Snapshots from the VMware vCenter Server as described in the Back Up VMs to On-Prem, Replicate and Failover to Azure Cloud.
- If you are backing up the Cloud VMs to the Cohesity Cloud Edition Cluster using the Cohesity Protection Service Protection Group type, then the VMs that are backed up to the Cloud Service must also have Cohesity agent installed.

Steps:

1. If not already configured, set up replication between the Cohesity Cloud Edition Cluster and the on-premises Cohesity cluster. For instructions, see Replication and Remote Access Setup.
2. On the Cohesity Cloud Edition Cluster, create a Policy that replicates from the Cohesity Cloud Edition Cluster to the on-premises Cohesity cluster as shown in the figure. For instructions, see Create or Edit a Standard Policy.
3. On the Cohesity Cloud Edition Cluster, refresh the Object hierarchy for the Source of the VMs that will be backed up:

a. Select **Data Protection > Sources**.

b. Click the action icon (⋮) next to the Source and select **Refresh**.

4. On the Cohesity Cloud Edition Cluster, create a Virtual Server Protection Group (such as `BackToPrem`) with Protection Group Type option as **Native Snapshots** or **Cohesity Protection Service**. Associate the Policy created in the previous step with this Protection Group. For detailed instructions, see Add or Edit a Protection Group for Virtual Servers. If **Cohesity Protection Service** is selected as the Protection Group type, then the VMs that are backed up to the Cloud Service (VMs that failed over to the cloud) must have Cohesity agent installed.

   When selecting Windows VMs to failback (recover), you can only select Snapshots of VMs that had the Cohesity agent installed when they were created with the Back Up VMs to On-Prem, Replicate and Failover to Azure Cloud process. The VM should be powered on when the Protection Group is running.

   When you create a Protection Group on the Cohesity Cloud Edition Cluster that replicates to an on-premises Cohesity cluster, after the first Snapshot is replicated to the on-premises Cluster, an Inactive copy of the Job is created on the on-premises Cluster that received the Snapshot. The new Inactive Job on the on-premises Cohesity cluster has the same name as the Protection Group on the Cohesity Cloud Edition Cluster, for example: `BackToPrem`.

5. You can failback (recover) the VMs to the original Hypervisor server or an alternate Hypervisor server:

   a. On the on-premises Cohesity cluster that contains the replicated Snapshots, select **Data Protection > Protection** and browse for the Inactive Job that was automatically created.

   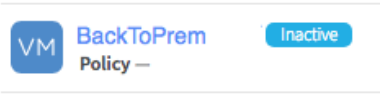

   b. Click on the inactive Job name.

   c. Browse for protection run that captured a Snapshot that you want to recover.

   d. Click the action icon (⋮) next to the Job and select **Recover VM**.

   e. Follow the instructions provided in Recover VMs.

   The VM is created in the Hypervisor server and is powered up if that option was selected. As part of the convert and clone process, the VMware Tools will be uninstalled. You must reinstall the VMware Tools.

   > **Important:** Do not change the size of the partition or the disk of the recovered VM.

Disaster Recovery of Views in Azure Cohesity Cloud Edition

Cohesity Cohesity Cloud Edition supports the disaster recovery of Views as shown by the following example figures.

Views can be used to store data such as files and directories. For example, other backup products can store data on a view that is hosted on the Cohesity cluster. To access a view from your system, you must mount the view.

On the left of the first figure, shows systems accessing a View hosted on an on-premises Cohesity cluster using the NFS or SMB protocols. A Protection Group running on the on-premises Cohesity cluster captures Snapshots of the View and stores these Snapshots on the Cluster. The Protection Policy associated with this Job, is configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity cluster Cohesity Cloud Edition. From these Snapshots stored in the Cohesity Cloud Edition, you can clone a new View that can be accessed by systems using the NFS or SMB protocols as shown on the right of the first figure.

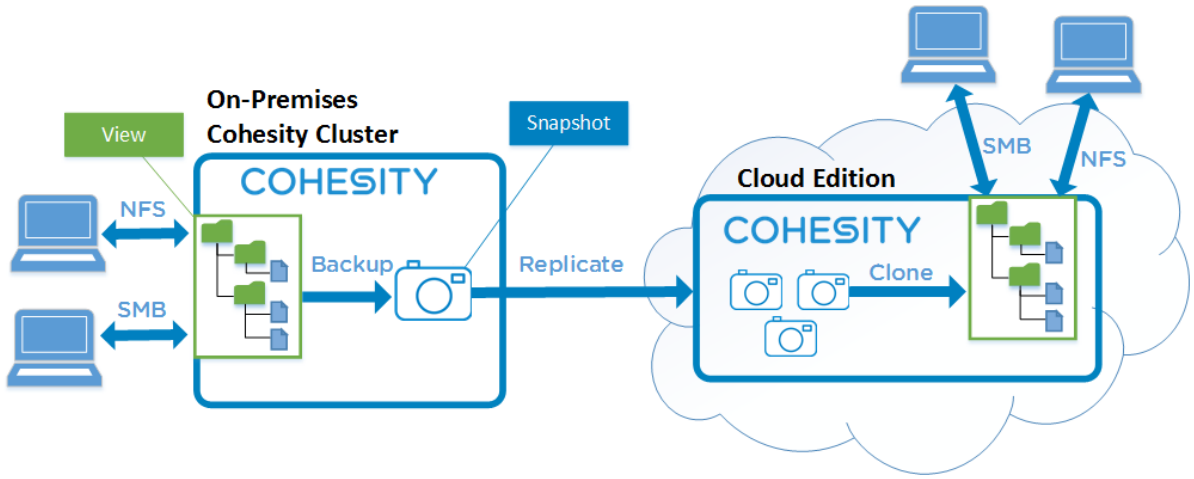

You can backup Views hosted on the Cohesity Cloud Edition as shown on the right of the following figure. A Protection Group captures a Snapshot of a View and stores the Snapshot on the Cohesity cluster Cohesity Cloud Edition. The Protection Policy associated with this Job, is configured to replicate (or copy) Snapshots from the Cohesity cluster Cohesity Cloud Edition to the on-premises Cohesity cluster. From these View Snapshots stored in the on-premises Cohesity cluster, you clone a View and make it available for mounting as shown on the left of the following figure.

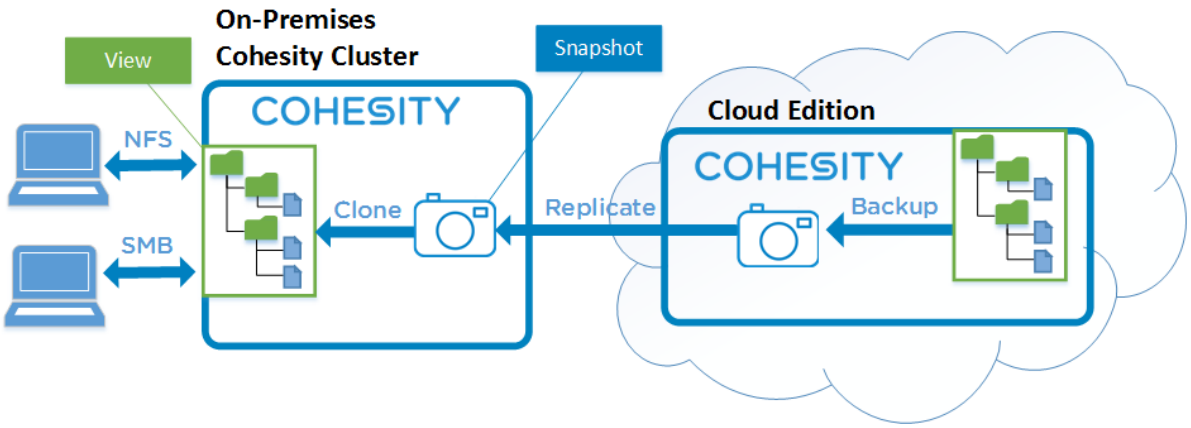

Clone Azure Virtual Machines

This chapter includes the following topics:

Clone VMs from On-Premises to Azure Cloud Service (On Demand CloudSpin)

Cohesity supports conversion of snapshots of VMs from an on-premises Cohesity cluster to the Azure Cloud service. After conversion, the VM runs in the Cloud service.

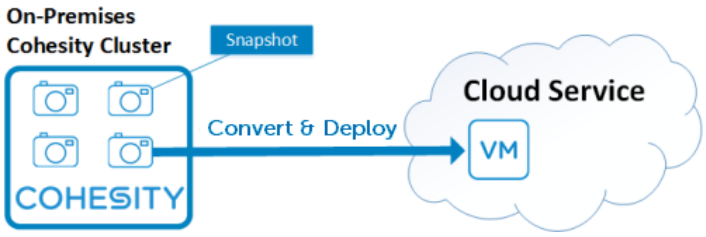

Prerequisites and Considerations

- For the list of VMs that supports CloudSpin to Azure, see Supported Software.
- Cohesity supports the CloudSpin of VMware VMs with UEFI boot mode enabled.
- Cohesity supports both on-demand and policy-based CloudSpin for snapshots stored on the Cohesity cluster (primary copy).
- Cohesity does not support on-demand and policy-based CloudSpin for archived or CloudRetrieved snapshots.
- Cohesity supports on-demand CloudSpin for snapshots that are replicated to another Cohesity cluster.
- Cloning VMware and Hyper-V VMs to Azure
    - Linux CentOS 7 on LVM
        - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.

- When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

  ```
  UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
  defaults,nofail 0 0
  ```

  For detailed instructions, see Edit /etc/fstab to use UUID.

- Windows
  - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
  - For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
  - When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. The Cohesity agents installed must be of the same version as the Cohesity Cluster version. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.
  - Due to a Windows issue in builds 14393 and build 15063, a deployed Windows 10 or Windows 2016 VM may not be connected to the network. Resolve this by installing fix KB4073562 for build 14393 or KB4016240 for build 15063 or by updating Windows. For more information, see Cannot connect remotely to a Windows 10 or Windows Server 2016 VM in Azure because of netvsc.sys.
  - Make the disks online by default using the following SAN policy commands:
    - `diskpart`
    - `san policy=OnlineAll`

    For more information, see the SAN Policy documentation.
- Red Hat Enterprise Linux
  - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
  - For the list of supported RHEL versions, see Supported Software.
  - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed

up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

```
UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
defaults,nofail 0 0
```

For detailed instructions, see Edit /etc/fstab to use UUID.

- The disk size of the VMware VMs must be a multiple of 1 GB.
- Cohesity does not support the CloudSpin of VMs whose OS Disk is formatted with XFS version 5 to Azure. This limitation applies to both Hyper-V and VMware VMs, regardless of their operating system type.
- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to Azure.

- Cloning Hyper-V VMs to Azure
  - VMs that have one to five disks can be cloned. Each disk must be smaller than 2 TB.
  - SUSE Linux VMs are *not* supported.
  - All the Hyper-V versions supported by Cohesity for backup is supported for cloning to Azure. For the list of supported Hyper-V versions, see Supported Software.
  - Generation 2 VMs are not supported by all VM sizes in Azure. For information on Generation 2 supported VM series, see Microsoft document.
  - Cloning of Hyper-V VMs to Azure Stack Hub is not supported.

Convert and Deploy a VM from On-Premises to Azure

Verify that the above prerequisites are met prior to backing up the VMs.

1. Create a Protection Group to capture Snapshots of VMs. For more information, Add a Protection Group for Azure VM.
2. Register the Cloud service that the VMware VMs are cloned **to**. For detailed instructions, see either Register an Azure Cloud Source at Subscription Level.
3. Log in to the Cohesity Dashboard of the Cohesity cluster that contains the Snapshots to clone from. For example, log in to the on-premises Cohesity cluster.
4. To clone and convert the Virtual Server (VM), follow the instructions provided in Clone VMs and select an Azure Subscription Source.
5. After the Clone task completes, log in to the new VM in the Cloud service.

   For Azure:

   a. Log in to Microsoft Azure Portal.
   b. Search for the new Virtual Machine name.
   c. Click on the VM name.

   d. Under the **Subscription name**, copy the IP address of the new VM.
   e. Log in to the VM using the IP address. If Remote Desktop was installed in the VM that was cloned, you can log in using the Remote Desktop.

6. Multiple disks are automatically mounted on the Windows VM, if the following conditions are true before the Protection Group captures the Snapshot of a Windows VM:
   - A Cohesity agent is installed on the VM.
   - VMware tools are installed on the VM.
   - VM is running

   If any of the conditions were not true when the Job captured the Snapshot, the `C:` drive is mounted but you must mount the other drives such as `D:` as described here:
   a. If Remote Desktop is installed on the new cloned VM, log in to the new VM, start Remote Desktop Connection, enter the IP of the new cloned VM.
   b. In the new VM, launch the Windows Disk Management.
   c. Right click on a disk that is not mounted and select **Online**. Repeat for all other offline disks.

Tear Down a Launched VM from Cohesity Dashboard

1. In the Cohesity Dashboard, select **Test & Dev**.
2. Click the actions menu (⋮) next to the clone task you want to tear down and select **Tear Down**.

Clone VMs to Azure Cloud Using CloudSpin

On a Cohesity cluster with a Protection Group that captures Snapshots of VMs with a CloudSpin schedule, the cluster clones VMs to Azure cloud service. You can then launch the VMs to a running state.

What happens during the CloudSpin process is dependent on the cloud service type:

- Cohesity supports CloudSpin of Hyper-V to Azure.
- Cohesity supports CloudSpin of VMware to Azure.
- During cloning to an Azure Cloud using Cloudspin, the Cohesity cluster converts the snapshot to VHD format and then copies the VHD image to a storage container in the Azure Cloud. Managed disks, including SSD and HDD, are also supported.
- For policy-based CloudSpin to Azure, the Storage Account must support page blobs.

For the list of VMs that supports CloudSpin to Azure, see Supported Software.

Cohesity supports CloudSpin on both the on-premises Cohesity cluster and Cohesity Cloud Edition Cohesity clusters:

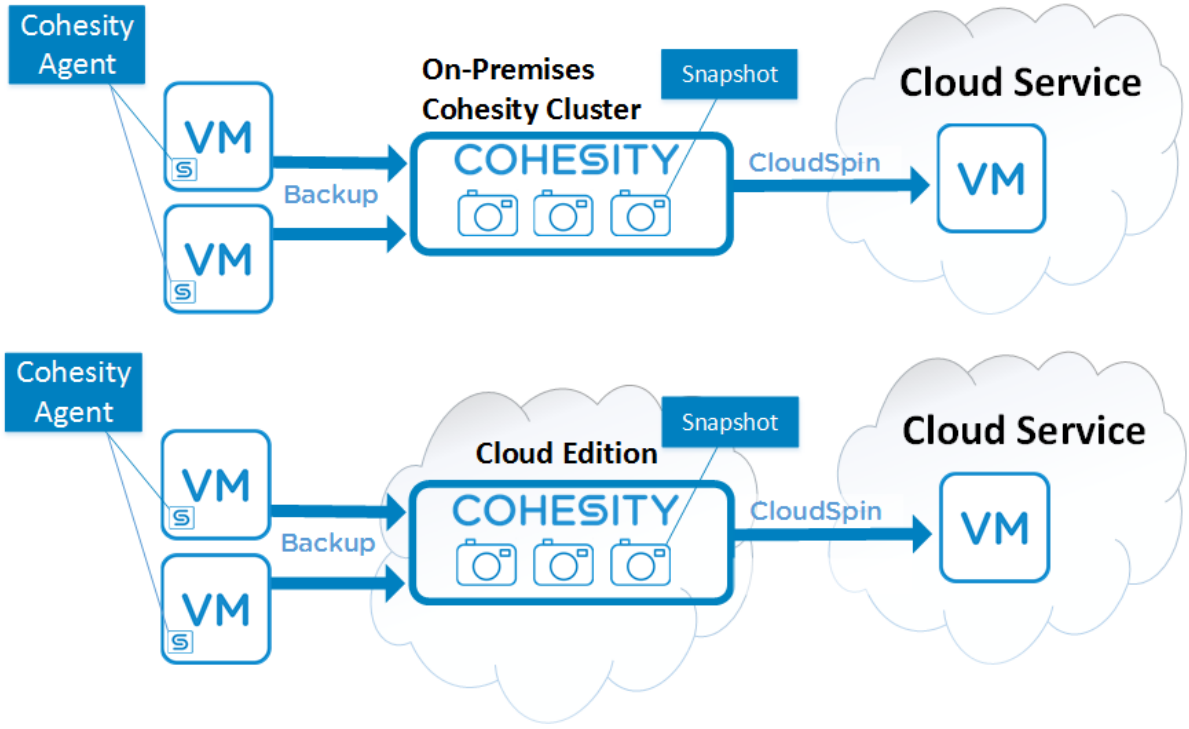

To use CloudSpin to clone Snapshots to Azure :

1. Verify the following requirements prior to backing up the VMs:
    - All the VMware vSphere versions supported by Cohesity for backup are supported for cloning to Azure. For the list of supported VMware vSphere versions, see Supported Software.
    - Cohesity does not support the CloudSpin of VMs whose OS Disk is formatted with XFS version 5 to Azure.
    - Cloning Snapshots to an Azure Cloud is supported for VMware and Hyper-V VMs with the following Operating Systems:
        - Linux CentOS 7 on LVM
            - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
            - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

              ```
              UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4 defaults,nofail 0 0
              ```

              For detailed instructions, see Edit /etc/fstab to use UUID.

- Windows
  - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
  - For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
  - When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. The Cohesity agents installed must be of the same version as the Cohesity Cluster version. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.
  - Due to a Windows issue in builds 14393 and build 15063, a deployed Windows 10 or Windows 2016 VM may not be connected to the network. Resolve this by installing fix KB4073562 for build 14393 or KB4016240 for build 15063 or by updating Windows. For more information, see Cannot connect remotely to a Windows 10 or Windows Server 2016 VM in Azure because of netvsc.sys.
  - Make the disks online by default using the following SAN policy commands:
    - `diskpart`
    - `san policy=OnlineAll`

    For more information, see the SAN Policy documentation.
- Red Hat Enterprise Linux
  - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
  - For the list of supported RHEL versions, see Supported Software.
  - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

    ```
    UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
    ```

```
defaults,nofail 0 0
```

For detailed instructions, see Edit /etc/fstab to use UUID.

- The disk size of the VMware VMs must be a multiple of 1 GB.
- Cohesity does not support the CloudSpin of VMs whose OS Disk is formatted with XFS version 5 to Azure. This limitation applies to both Hyper-V and VMware VMs, regardless of their operating system type.
- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to Azure.

- Cloning Hyper-V VMs to Azure
  - VMs that have one to five disks can be cloned. Each disk must be smaller than 2 TB.
  - SUSE Linux VMs are *not* supported.
  - All the Hyper-V versions supported by Cohesity for backup is supported for cloning to Azure. For the list of supported Hyper-V versions, see Supported Software.
  - Generation 2 VMs are not supported by all VM sizes in Azure. For information on Generation 2 supported VM series, see Microsoft document.
  - Cloning of Hyper-V VMs to Azure Stack Hub is not supported.

2. Register a Cloud Source. Select **Data Protection > Sources** and from the **Register Source** drop-down, select **Azure Cloud**. A cloud source is a cloud service instance where the snapshots are converted and copied **to**. For details, see the appropriate Source topic:
   - Register an Azure Cloud Source at Subscription Level
3. For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
4. Create a Policy with a CloudSpin schedule that converts and copies Snapshots to the Cloud Source registered in the previous step. For detailed instructions, see Create or Edit a Standard Policy. Protection Groups that use this policy will automatically convert and copy Snapshots backed up by the Protection Group to a Cloud Source using the specified CloudSpin schedule.
5. Create a Protection Group that uses the Policy created in the previous step.

   After the Protection Group run and CloudSpin Copy Task are successfully completed, the VM can be launched.
6. Launch the VM created in the Cloud Service:

   Launch VM from the Cohesity Dashboard:

a. Select **Data Protection > Protection**.
b. Click the Job that backed up the VM to launch in the Cloud.
c. Browse for a protection run that contains a Snapshot to convert to a VM. The protection run must have the Status of **Success** and the associated CloudSpin Copy task must be successful .
d. Click the date link of the protection run in the Start Time column to view the protection run details.
e. Click the **ClouldSpin Task** tab.
f. To launch, complete one of the following actions:

- Click the actions menu (⋮) and select **Launch Instance**. The following fields are displayed:
  - **Source**—Displays the registered source.
  - **Resource Group**—Displays the resource pool that stores the cloned VMs.
  - **Storage Account**—Displays the VM size used when creating the cloned VMs.
  - **Storage Container**—Displays the storage container used to store the cloned VMs.
  - **Compute**—Select the Microsoft Azure VM size type to use when creating the cloned VMs.
  - **Availability Set**—Optionally, select the availability set to which the VM has to be cloned. For more information, see Azure Availability Sets.
  - **Virtual Network**—Select the Microsoft Azure Virtual Network for the cloned VMs.
  - **Subnet**—Select the Microsoft Azure Virtual subnet for the cloned VMs.
  - **Use Azure Managed Disks**—To use Azure managed disks, enable the toggle and select the OS Disk Type and Data Disk Type.

    Click **Submit** to save the changes.
- Click **Launch All Instances**.

To tear down a launched VM, see Tear Down a Launched VM from Cohesity Dashboard.

Launch from Azure using PowerShell:

Follow these steps from the Create a Windows VM from a specialized disk using PowerShell topic in the Microsoft Azure Documentation:

a. Install the newest version of Powershell as described in Before you begin.
b. Follow the steps in Option 2: Upload a specialized VHD.
c. Follow the steps in Create the new VM.

To prevent a teardown of a cloned CloudSpin VM, see the How to prevent a tear down of a cloned CloudSpin VM KB article.

Tear Down a Launched VM from Cohesity Dashboard

- In the Cohesity Dashboard, select **Test & Dev**.
- For the clone task to tear down, click the actions menu (⋮) and select **Tear Down**.

Consideration

- Cohesity does not support CloudSpin of Active Directory VM directly to Azure. If you CloudSpin Active Directory VM directly to Azure, the VM does not boot and the **Choose an option Recovery** screen is displayed in **Boot Diagnostics**.

  Workaround: You need to first remove the Active Directory Domain Services in the on-premises VM, take a backup and CloudSpin it to Azure, and then add it to the Domain Controller on Azure. For more information, see Microsoft Documentation.
- Azure does not support booting UEFI VMs with unmanaged disks.
- The maximum data disk size supported for an Azure VM for CloudSpin is 8191 GiB.
- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to Azure.
- Cohesity supports both on-demand and policy-based CloudSpin for snapshots stored on the Cohesity cluster (primary copy).
- Cohesity does not support on-demand and policy-based CloudSpin for archived or CloudRetrieved snapshots.
- Cohesity supports on-demand CloudSpin for snapshots that are replicated to another Cohesity cluster.

Edit /etc/fstab to use UUID

You must edit the `/etc/fstab` file and change the (non-LVM) entries to use UUIDs to identify the disk partitions. These changes must be made to the `/etc/fstab` file before the Protection Group (which captures the Snapshots to be cloned) is run.

For example, the following entry in the `/etc/fstab` table must be changed to use the UUID to identify the partition instead of the device file name of `/dev/sbd1`.

```
/dev/sdb1 /mnt/disk1    ext4    defaults        0 0
```

> **Note:**
> - Use labels or UUIDs instead of device names, as device names may change when other disks are added or removed.
> - Kernel name descriptors must not be used in configuration files as they are not persistent and may change during each boot. UUIDs or PARTUUIDs are preferred over labels.
> - Using only the device name in the absence of UUID or label might cause issues as these names are not unique. This can cause overwrites and eventual loss of an entire volume indexing itself.
> - Each volume must be uniquely identified by an identifier like UUID or label. **A missing unique identifier (UUID/label) can result in dummy volume names being used instead of actual volumes.**

You can determine the UUID by running the `blkid` command, for example:

```
# blkid /dev/sdb1
/dev/sdb1 UUID="12345cac-77bc-46b1-9f79-cb180e438c12" TYPE="ext4"
```

Edit the `/etc/fstab` as root and for this entry replace the device file name with the returned UUID, for example:

```
12345cac-77bc-46b1-9f79-cb180e438c12 /mnt/disk1 ext4    defaults        0 0
```

> **Note:** LVM volume entries should not be updated. LVM volumes start with the `/dev/mapper` prefix, for example: `/dev/mapper/centos-root` and `/dev/mapper/centos-swap`.

Repeat these steps for all entries in the `/etc/fstab` that are not LVM volume entries.

Replicate Azure Virtual Machines

This chapter includes the following topics:

Replicating VMs in Azure Cloud Edition

Cohesity Cohesity Cloud Edition supports the replication of VMs captured on an on-premises Cohesity cluster to Cohesity Cloud Edition as shown by the example figure below.

Starting from the left in the following figure, the use case shows how Protection Groups running on the on-premises Cohesity cluster captures Snapshots of VMs and stores these Snapshots on the on-premises Cluster. The Protection Policies associated with these Jobs,

are configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity Cloud EditionCohesity cluster.

Starting from the left in the following figure, the bottom workflow shows how the Cohesity Cloud Edition can directly backup a VM that is not in the Cloud. For example a VM in a vCenter server can be backed up directly to Cohesity Cloud Edition running in a Azure Cloud service.

You can recover VMs from the Snapshots stored in the Cohesity Cloud Edition running in the Azure Cloud service.

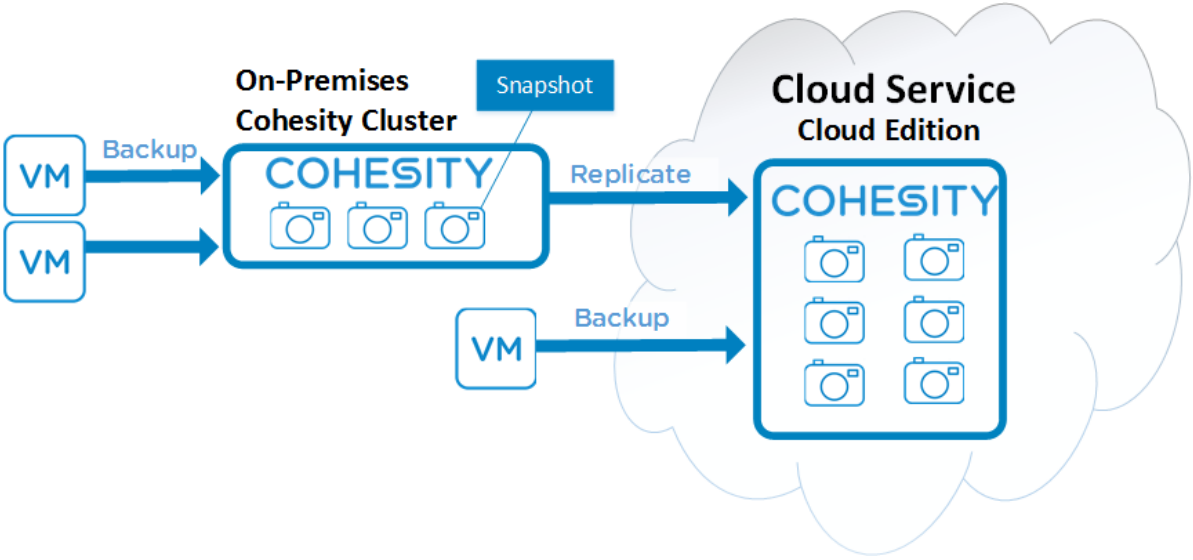

Replicate and Clone VMs in Azure Cloud Edition

CohesityCohesity Cloud Edition supports Test/Dev (cloning) as shown by the following example figure.

Starting from the left in the following figure, the top workflow shows how Protection Groups running on the on-premises Cohesity cluster captures Snapshots of the VMs and stores these Snapshots on the on-premises Cluster. The Protection Policies associated with these Jobs, are configured to replicate (or copy) Snapshots from the on-premises Cohesity cluster to the Cohesity clusterCohesity Cloud Edition.

Starting from the left in the following figure, the bottom workflow shows how the Cohesity Cloud Edition can directly backup a VM that is not in the Cloud. For example a VM in a vCenter server can be backed up directly to Cohesity Cloud Edition running in a Azure Cloud service.

From the Snapshots stored in the Cohesity Cloud Edition on the Azure Cloud service, you can spin up test/dev environments using cloning.

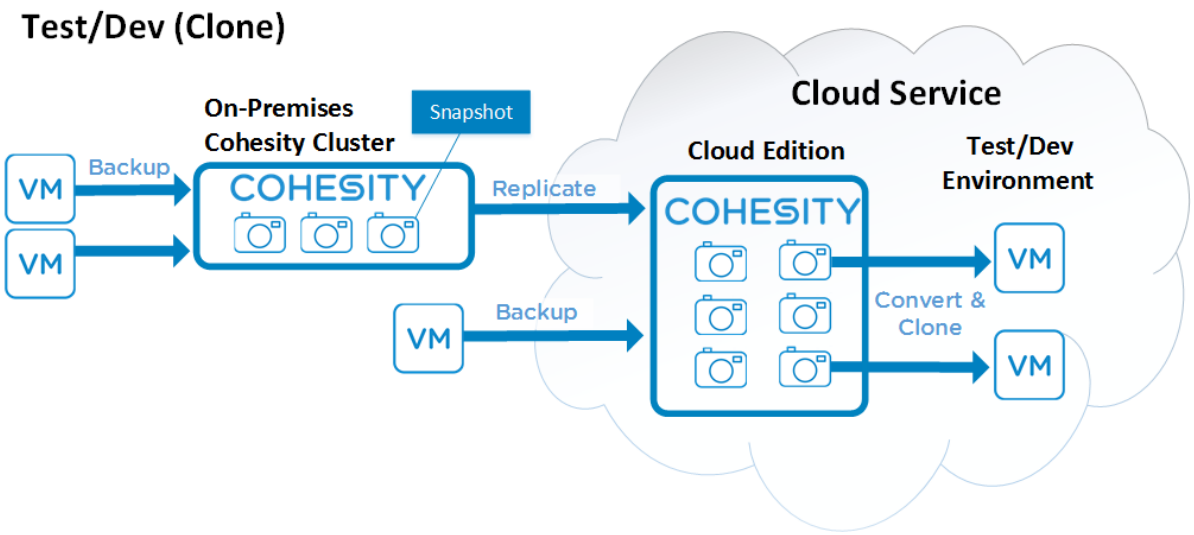

To convert and clone VMs to a Cloud service:

1. Verify the following requirements prior to backing up the VMs:
   - For the list of VMs that supports CloudSpin to Azure, see Supported Software.
   - All the VMware vSphere versions supported by Cohesity for backup is supported for cloning to Azure. For the list of supported VMware vSphere versions, see Supported Software.

   Cloning Snapshots to an Azure Cloud is supported for VMware and Hyper-V VMs with the following operating systems:
   - Linux CentOS 7 on LVM
     - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
     - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

       ```
       UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
       defaults,nofail 0 0
       ```

       For detailed instructions, see Edit /etc/fstab to use UUID.
   - Windows
     - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
     - For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.

- When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. The Cohesity agents installed must be of the same version as the Cohesity Cluster version. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.
- Due to a Windows issue in builds 14393 and build 15063, a deployed Windows 10 or Windows 2016 VM may not be connected to the network. Resolve this by installing fix KB4073562 for build 14393 or KB4016240 for build 15063 or by updating Windows. For more information, see Cannot connect remotely to a Windows 10 or Windows Server 2016 VM in Azure because of netvsc.sys.
- Make the disks online by default using the following SAN policy commands:
  - `diskpart`
  - `san policy=OnlineAll`

  For more information, see the SAN Policy documentation.

- Red Hat Enterprise Linux
  - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
  - For the list of supported RHEL versions, see Supported Software.
  - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

    ```
    UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
    defaults,nofail 0 0
    ```

    For detailed instructions, see Edit /etc/fstab to use UUID.
- The disk size of the VMware VMs must be a multiple of 1 GB.
- Cohesity does not support the CloudSpin of VMs whose OS Disk is formatted with XFS version 5 to Azure. This limitation applies to both Hyper-V and VMware VMs, regardless of their operating system type.
- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to Azure.

Cloning Hyper-V VMs to Azure

- VMs that have one to five disks can be cloned. Each disk must be smaller than 2 TB.

- SUSE Linux VMs are *not* supported.
- All the Hyper-V versions supported by Cohesity for backup is supported for cloning to Azure. For the list of supported Hyper-V versions, see Supported Software.
- Generation 2 VMs are not supported by all VM sizes in Azure. For information on Generation 2 supported VM series, see Microsoft document.
- Cloning of Hyper-V VMs to Azure Stack Hub is not supported.

2. Create a Protection Group to capture Snapshots of VMs. For more information, see Add a Protection Group for Azure VM.
3. Register the Cloud service that the VMware VMs are cloned **to**. For detailed instructions, see either Register an Azure Cloud Source at Subscription Level .
4. Log in to the Cohesity Dashboard of the Cohesity cluster that contains the Snapshots to clone from. That is, log in to the Cohesity cluster running in Azure.
5. Clone the VM in the Azure Cloud service. For detailed instructions, see Clone VMs.
6. After the Clone task completes, log in to the new VM in the Cloud service.

   For Azure:

   a. Log in to Microsoft Azure Portal.
   b. Search for the new Virtual Machine name.
   c. Click on the VM name.
   d. Under the **Subscription name**, copy the IP address of the new VM.
   e. Log in to the VM using the IP address. If Remote Desktop was installed in the VM that was cloned, you can log in using the Remote Desktop.
7. Multiple disks are automatically mounted on the Windows VM, if the following conditions are true before the Protection Group captures the Snapshot of a Windows VM:

   - A Cohesity agent is installed on the VM.
   - VMware tools are installed on the VM.
   - VM is running

   If any of the conditions were not true when the Job captured the Snapshot, the `C:` drive is mounted but you must mount the other drives such as `D:` as described here:

   a. If Remote Desktop is installed on the new cloned VM, log in to the new VM, start Remote Desktop Connection, enter the IP of the new cloned VM.
   b. In the new VM, launch the Windows Disk Management.
   c. Right click on a disk that is not mounted and select **Online**. Repeat for all other offline disks.

Migrate Azure Virtual Machines

This chapter includes the following topics:

Migrate Cloud Edition Azure Nodes to Managed Disks

The nodes (VMs) in a Cohesity Cloud Edition Azure cluster can use Azure-managed disks or unmanaged disks. Managed disks are described in Introduction to Azure Managed Disks.

Cohesity recommends that all nodes in a cluster use the same type of disk. You can migrate nodes from unmanaged disks to managed disks as described below.

Before Migrating the Nodes

Do the following before you migrate the nodes to use managed disks.

1. Take a snapshot of every node (VM) in the cluster as a backup precaution.
2. SSH into any node in the cluster.

   ```
   ssh cohesity@<ip_address>
   ```

3. Run the following commands and record the command output to verify disk utilization after the migration.

   ```
   df -h
   ```

   ```
   lsblk
   ```

Migrate the Nodes

To migrate the nodes to use managed disks, repeat the following steps for each node in the cluster.

1. Make sure there are no outstanding replication, archival or CloudSpin tasks active for the cluster.
2. In the Azure portal, click the Cloud Shell button shown below and activate the PowerShell environment.

   For more information, see Quickstart for PowerShell in Azure Cloud Shell.
3. In the PowerShell window, run the following commands to deallocate the node (VM).

   Replace *myResourceGroup* and *myVM* with the names of the resource group and VM instance name you want to deallocate.

   ```
   $rgName = "myResourceGroup"
   ```

   ```
   $vmName = "myVM"
   ```

   ```
   Stop-AzVM -ResourceGroupName $rgName -Name $vmName -Force
   ```

4. After the above commands run successfully, run the following command to initiate the migration.

```
ConvertTo-AzVMManagedDisk -ResourceGroupName $rgName -VMName $vmName
```

5. Repeat the previous steps for each node in the cluster.

Verify the Migration

1. SSH to any node in the cluster.
2. Run the following Cohesity CLI command to verify that all services are up and running.

   ```
   iris_cli cluster status
   ```

3. Run the following commands and compare the results to the same command output that you recorded previously. The disk utilization should be the same as it was before the migration.

   ```
   df -h
   ```

   ```
   lsblk
   ```

Azure MySQL

Cohesity DataProtect provides a simple, fast, and cost-effective backup, recovery, and data management solution for Microsoft Azure MySQL databases.

Azure Database for MySQL is a fully managed relational database service offered by Microsoft Azure. It allows users to deploy, manage, and scale MySQL databases in the cloud without the need to manage the underlying infrastructure.

> **Important:** The backup and recovery of Azure MySQL is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).

Plan and Prepare for Azure MySQL Protection

Before protecting Azure MySQL with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure MySQL with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Authentication for Azure MySQL Protection

You must authenticate the Cohesity cluster to discover the Azure MySQL databases using one of the following methods:

- **Username** and **password** of the Azure MySQL instance.
- **User Managed Identity**.

Configure Database User Associated with User Managed Identity

To configure a database user associated with User Managed Identity:

1. Get the client ID for managed identity using the following commands in Azure CLI:
   ```
   # az ad sp list --display-name <managed_identity_name> --query
   [*].appId
   --out tsv
   ```
2. Login to the MySQL server and create a user using the following commands:
   ```
   # SET aad_auth_validate_oids_in_tenant = OFF;
   # CREATE AADUSER '<db_user>' IDENTIFIED BY
   '<Generated_client_id>';
   # GRANT <Permissions> ON *.* TO '<db_user>'@'%'
   ```
3. Create a normal database user or one associated with the relevant User Managed Identity and grant the required permissions to that user using the following command:
   GRANT SELECT , INSERT , UPDATE, DELETE, CREATE, DROP , RELOAD,
   ```
   PROCESS, REFERENCES, INDEX, ALTER, SHOW DATABASES, LOCK
   TABLES, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE,
   SHOW_ROUTINE, EVENT , TRIGGER, SET_USER_ID ON *.* TO `<dbuser>`@'%'
   WITH GRANT OPTION;
   ```

   > **Note:** Replace <*dbuser*> in the above command with the actual name of the database user.

Firewall Ports for Azure MySQL Protection

- Ensure connectivity from Cohesity DataProtect to the Azure MySQL server on port 3306.

Considerations for Azure MySQL Protection

- **Full backup support** - Only full backup is supported; incremental backup is not available for Azure MySQL.
- **Azure MySQL database size display** - The size of the database under Azure MySQL server will only be populated once the credentials are specified for the server entity. Otherwise, it will be reported as 0.
- **Attach metadata and credentials** - A database instance moved to a different resource group is considered a new entity. Hence, ensure that you attach the new metadata and credentials.

- **Azure MySQL version compatibility** - You can recover an Azure MySQL database to an alternate instance with an Azure MySQL version different than the backup instance, depending on the Azure MySQL version compatibility.
- **Database naming convention** -Backup and recovery of databases with names containing the characters **#** and **/** are not supported. Ensure the database names adhere to the naming conventions suggested by the cloud service.
- **Special character in password** - Server or database passwords do not support **;** (semicolon).
- **Database Level Protection** - Protection occurs at the database level and not server level; backup and recovery are performed for the databases on the server. Cohesity does not change or create an Azure MySQL server in the workflows. Hence, the properties of the server like tags are not recovered.

Backup Azure MySQL Database

After registering your Azure MySQL source in Cohesity DataProtect, Cohesity uses the cloud credentials provided to authenticate with Azure.

Once authenticated, Cohesity automatically discovers all the subscriptions that the service account can access, as well as all supported workloads such as Azure MySQL databases within those subscriptions. The discovered workloads appear under **Data Protection > Sources**. You can select the **MySQL** tab to view all the databases discovered.

Hierarchy View displays the entity hierarchy of the Azure MySQL objects, allowing you to navigate from the resource group-level down to individual databases.

Azure MySQL Entity Hierarchy

The Azure MySQL entity hierarchy follows this structure:

Structure	Description
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Database Server	Fully managed MySQL database instance hosted within the Azure cloud.
Database	Fully managed relational database service in Microsoft Azure.

Backup Scope by Azure MySQL Hierarchy

Cohesity cluster allows you to backup Azure MySQL databases at different levels of the hierarchy.

You can choose to back up:

- **Subscription** – Protect all Azure MySQL databases within the subscription.
- **Resource Group** – Protect all Azure MySQL databases within the resource group.
- **Database Server** – Protect all Azure MySQL databases within the database server.
- **Database** – Protect individual Azure MySQL database.

Supported Backup Types

Cohesity supports only full backups for Azure MySQL. Incremental backups are not supported.

Protect Azure MySQL

Azure MySQL databases are protected using a Protection Group and a Protection Policy.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites to Protect Azure MySQL

Set Azure MySQL Server Credentials

For Backup, the database credentials are used to authenticate and connect to the MySQL server. Once the connection is established, the relevant database can be dumped and protected.

For Azure MySQL, you can set credentials at the Azure MySQL Server level. This information is used to connect to the database during backup and recovery. You can also add credentials at any higher level in the hierarchy, such as the subscription. Credentials added at a higher level automatically apply to all objects beneath it.

The Cohesity cluster searches for credentials in a bottom-up manner, starting from the specific entity (for example, a server) and moving up the hierarchy until valid credentials are found. When multiple credentials exist, the credentials closest to the MySQL server take precedence.

For example:

- If **region-1** has a set of credentials and **server-1** under that region has its own, the Cohesity cluster uses the credentials on **server-1** to authenticate to **server-1**.
- Another server, **server-2**, under the same region, the Cohesity cluster uses the region-level credentials for authentication because that is the closest level where credentials are specified.

To set the database credentials:

1. In **DataProtect**, navigate to **Data Protection > Sources**, find the registered Azure source, and click on it.
2. Click on the **MySQL** tab.
3. Select the checkbox next to the required Azure MySQL Server.
4. Click the **Database Credentials** button.
5. You can select one of the following options:
    1. **Username and Password** - Provide the database admin username and password. These credentials are required for leaf-level databases. The account must have access to all databases within the MySQL server.
    2. **User Managed Identity (UMI)** - Uses authentication automatically managed by Azure. For more information, see Configure Database User Associated with User Managed Identity.

        > **Note:**
        > - User Managed Identity must be attached to the Azure VM on which Cohesity DataProtect is deployed.
        > - UMI authentication method does not work across tenants.
        > - Authentication fails if the database credentials provided are incorrect. Ensure that the username/password/UMI ClientID is correct, and/or the UMI associated database user is properly configured.
6. Click **Save**.

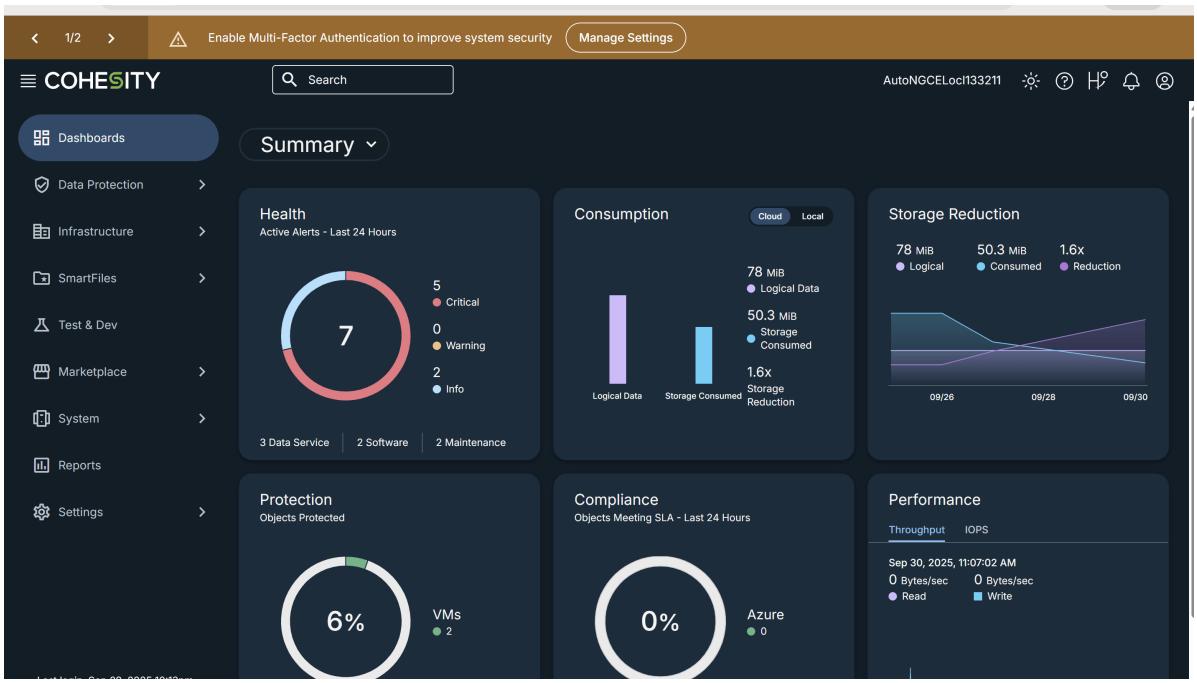

Protect Azure MySQL Databases

To protect your Azure MySQL Server databases:

1. In **DataProtect** navigate to **Data Protection > Sources**.
2. Find the registered Azure MySQL source and click on it.
3. Click on the **MySQL** tab.
4. Use the checkboxes to select the objects for protection. You can select the tenants , subscriptions , resource groups , databases server , or databases under them. For more information, see Azure MySQL Entity Hierarchy.
   You can also use the **List View** to select the instances for protection from a flat, searchable list or **Tags View** to select the instances based on tags.
   - To protect all objects in the source, click the checkbox next to **Object**.
   - To protect all the child objects under the Azure MySQL Server, click the checkbox next to the Azure MySQL Server and click the **Select All Child Objects** option.
   - To auto-protect the Azure MySQL Server, click the checkbox next to the Azure MySQL Server object and select the **Auto-Protect This Azure MySQL Server** option.

> **Note:** Object-level exclusion is not supported during Azure MySQL protection.

5. Click **Protect**.
6. Choose a policy to specify backup frequency and retention. For more information on policy, see Create or Edit a Standard Policy.

> **Note:** Periodic full is not recommended for Azure MySQL since all incremental backups are converted to full backups. If periodic full is part of the policy, two full backups will run sequentially when new databases are onboarded.

7. Optionally, to configure any of the advanced settings, select **More Options**.
8. Click **Protect**.

   You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific Azure MySQL shows the history of all protection runs, including the one in progress.

   Cohesity DataProtect starts backing up the databases you selected.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
	End Date	If you need to end protection on a specific date, enable this option to select the date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
	Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.

Section	Setting	Description
	Priority	Select a priority for the Object Protection Group Execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Azure MySQL Backup Configuration

In the Cohesity Platform, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.

4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.

3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Azure MySQL Databases

After you protect your Azure MySQL databases, you can recover them from Cohesity DataProtect to their original or a new location.

To recover protected Azure MySQL databases:

1. In **DataProtect**, navigate to **Data Protection > Sources**.
2. Click the required Azure MySQL source and click the **MySQL** tab.
3. Select the **Protection Status** as **Protected**.
4. Use the filters, search box, and views to locate and select the required databases.
5. Click **Recover** at the top to open the **New Recovery** form with the **Latest** snapshot (protection run).

   If you need to recover from an earlier snapshot, click the **Edit** icon to open the **Recovery Point** page.
   Click **List** to view the available recovery points by timestamp and click **Select Recovery Point**.
6. Under **Recover To**, select either of the following:
   1. **Original PaaS Instance** - To recover the Azure MySQL database to the Azure MySQL Server instance or Managed Instance from which the database was backed up.
   2. **Alternate PaaS Instance**. Select a target Azure MySQL Server instance or Managed Instance to which the Azure MySQL database will be recovered.
7. Select your **Recovery Options**:
   1. **Rename**. Add a Prefix and/or Suffix to the full name of the recovered Azure MySQL database.
   2. **Task Name**. Change the default name of the recovery task.
8. Click **Recover**.

   You can monitor the status of the recovery on the **Activity** page.

   Cohesity DataProtect starts recovering the selected Azure MySQL databases.

Azure Cosmos DB - Cassandra

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Azure Cosmos DB - Cassandra keyspaces.

Azure Cosmos DB - Cassandra is a fully managed NoSQL database service that enables you to run Cassandra workloads in the cloud with minimal changes to your application code.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Plan and Prepare for Azure Cosmos DB – Cassandra Protection

Before protecting Azure Cosmos DB – Cassandra with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure Cosmos DB – Cassandra with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Authentication for Azure Cosmos DB - Cassandra Keyspace Protection

You must authenticate to Azure Cosmos DB - Cassandra keyspace from Cohesity DataProtect using one of the following methods:

- Username and Password
- User Managed Identity (UMI)
- System Managed Identity (SMI)

Username and Password

This method uses the credentials of the admin user for the Azure Cosmos DB - Cassandra keyspace. This is the user account configured during server creation.

For backup, you can create a custom user. In the CLI and run the following commands:

1. To create login user on the server:

```
#CREATE LOGIN [newlogin] WITH PASSWORD = 'somepassword;
```

2. To create same user under database which should be backed up by user **newlogin**:

```
User <database>;
#CREATE USER [newlogin] FOR LOGIN [newlogin];
ALTER ROLE db_owner ADD MEMBER [newlogin];
```

3. Run the following command on the master system database:

```
CREATE USER [newlogin] FOR LOGIN [newlogin];
```

For recovery, the default admin credentials created during server creation are required if the custom user credentials were set for authentication before backup. During recovery, you must change the credentials back to the default admin credentials.

User Managed Identity (UMI)

This method uses a user-assigned managed identity to authenticate to the Azure Cosmos DB - Cassandra keyspace.

Perform the following to use the UMI authentication:

1. Create a user-assigned managed identity. For more information, see Manage user-assigned managed identities using the Azure portal.
2. Assign the **Managed Identity Operator** role to the UMI identity. For more information, see User-assigned managed identity.

3. Set the UMI as the identity for all the cluster nodes. Ensure that the UMI is from the same subscription as the one in which the keyspace exists. For more information, see Assign a user-assigned managed identity to an existing VM.
4. Assign this User Managed Identity as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.
5. Assign the required user managed identity to Azure Cosmos DB - Cassandra keyspace under the Identity section - User assigned managed identity on the Azure Portal.

System Managed Identity (SMI)

This method uses a unique identity that is automatically created by Azure when you enable SMI on the node. This identity belongs only to the specific node and is deleted if the node is deleted.

Each node in the cluster includes its own SMI (multiple distinct identities). Since the Azure Cosmos DB - Cassandra keyspace allows only one Microsoft Entra admin per node, having multiple nodes, each with a different SMI may cause issues.

To mitigate this issue, you can perform the following:

1. Create a Microsoft Entra group.
2. Add SMI identities as group members.
3. Assign this group as Microsoft Entra Admin.

This ensures that all the nodes can be authenticated using their respective SMIs through the group.

> **Note:** SMIs are not created separately as a resource. You must enable it on the cluster nodes from the Azure portal.

Perform the following to use the SMI authentication:

1. Enable system-assigned managed identity on all the nodes of the cluster. For more information, see Enable system-assigned managed identity on an existing VM.
2. Create a Microsoft Entra basic group and add SMI's of all the nodes created in previous step as members. The SMI name is the same as the VM (node) name, which can be easily searched for while finding members to include in the group. For more information, see Create a basic group and add members.
3. Assign this Microsoft Entra basic group as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.

Considerations for Azure Cosmos DB - Cassandra Keyspace Protection

Review and understand the following limitations and considerations before you protect Azure Cosmos DB - Cassandra Keyspace:

- **Full back up support** - Only full backup is supported; incremental backup is not available for Azure Cosmos DB - Cassandra protection.
- **Attach metadata and credentials** - When you recover a keyspace to a different account, ensure that you attach the metadata and credentials to the destination account.
- **Rename keyspace during recovery** - Keyspace names can have up to 48 alpha-numeric characters and can contain underscores; only letters and numbers are supported as the first character.
- **Keyspace-Level Protection** - Protection occurs at the keyspace level and not account level; backup and recovery are performed for the keyspaces on the account. Cohesity does not change or create an Azure Cosmos DB-Cassandra account in the workflows. Hence, the properties of the account like tags are not recovered.
- **Ingest rate during recovery** - Ingest rate during recovery can be changed to a value between 0.1 and 1.0. The default value is 0.4, which is 40% of the maximum theoretical ingest rate capacity. This helps in resolving rate limit errors during recovery.
- **Serverless vs provisioned throughput** - Keyspace belonging to an account with Serverless mode cannot be recovered to the one with provisioned throughput capacity mode and vice versa.
- **SSR (Server Side Retry)** - During backup and recovery, SSR must be enabled on the source and destination Cassandra account. This is mandatory for the backup and recovery to be successful.
- **TTL (Time-to-live) on recovered tables** - When recovering Azure Cosmos DB-Cassandra tables, the default_time_to_live (TTL) setting is preserved from the source table and applied to the recovered table. As a result, any newly inserted rows in the recovered table will automatically expire after a duration of Row creation time + default_time_to_live (in seconds).

  This ensures consistency with the TTL-based data expiration behavior of the original table.

  > **Note:** Row-level TTL settings are not retained during backup. When data is recovered, the TTL for those rows is set to the table's default_time_to_live value.

- **Index Name on Recovery** - Index names are not retained during the backup and recovery of Azure Cosmos DB-Cassandra tables. While index names are accepted syntactically by the service, they are not functionally relevant in Azure Cosmos DB-Cassandra. To ensure consistency and compatibility, index names are regenerated during the recovery process.

- **Materialized views support** - If the source table includes materialized views, during recovery the destination account must have Materialized Views settings enabled. This is mandatory for successful recovery of materialized views. As per Azure, materialized views are not recommended for production workloads. For more information, see Materialized Views in Azure Cosmos DB for Apache Cassandra.
- **Replication regions** - To get optimum performance during recovery, it is recommended to disable multi-region replication. After recovery, you can re-enable multi-region replication.
- **Version Downgrade not supported**: Recovering a backup taken from a higher database version to a lower version is not supported.

Backup Azure Cosmos DB - Cassandra Keyspace

After registering your Azure source in the Cohesity cluster, Cohesity uses the cloud credentials provided to authenticate with Azure.

Once authenticated, Cohesity automatically discovers all the subscriptions that the service account can access, as well as all supported workloads such as Azure Cosmos DB - Cassandra keyspaces within those subscriptions. The discovered workloads appear under **Data Protection > Sources**. You can select the **Cosmos DB - Cassandra** tab to view all the keyspaces discovered.

Hierarchy View displays the entity hierarchy of the Azure Cosmos DB - Cassandra objects, allowing you to navigate from the resource group-level down to individual keyspaces.

Azure Cosmos DB - Cassandra Entity Hierarchy

The Azure Cosmos DB - Cassandra entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Cosmos DB Cassandra Account	Fully managed NoSQL database service hosted within the Azure cloud.

Structure	Description
Keyspace	The top-level container for your data which holds tables. It is analogous to a database in RDBMS.

Backup Scope by Azure Cosmos DB - Cassandra Hierarchy

Cohesity allows you to backup Azure Cosmos DB - Cassandra keyspaces at different levels of the hierarchy.

You can choose to back up:

- **Tenant** – Protect all Azure Cosmos DB - Cassandra keyspaces within the tenant.
- **Subscription** – Protect all Azure Cosmos DB - Cassandra keyspaces within the subscription.
- **Resource Group** – Protect all Azure Cosmos DB - Cassandra keyspaces within the resource group.
- **Cosmos DB Cassandra Account** – Protect all Azure Cosmos DB - Cassandra keyspaces within the Cosmos DB Account.
- **Keyspace** – Protect individual Azure Cosmos DB - Cassandra keyspace.

Supported Backup Types

Cohesity supports full backups for Azure Cosmos DB - Cassandra. Incremental backups are not supported.

Protect Azure Cosmos DB - Cassandra Keyspaces

Azure Cosmos DB - Cassandra keyspaces are protected using a Protection Group and a Protection Policy.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites to Protect Azure Cosmos DB - Cassandra

Set Azure Cosmos DB - Cassandra Database Credentials

For Backup, the account credentials are used to authenticate and connect to the Cosmos DB - Cassandra Account. Once the connection is established, the relevant keyspace can be dumped and protected.

For Azure Cosmos DB - Cassandra, you can set credentials at the account, resource group, subscription, and tenant level. This information is used to connect to the keyspace during backup and recovery.

It is recommended to apply credentials at the Azure Cosmos DB - Cassandra account level. If you apply credentials at a higher level in the entity hierarchy, they will be applied to all child-level databases. If the credentials do not match any of the databases at the child-level, the backup of the databases may fail.

To set the keyspace credentials:

1. Navigate to **Data Protection > Sources**, find the registered Azure source, and click on it.
2. Click on the **Cosmos DB - Cassandra** tab.
3. Select the checkbox next to the required Azure Cosmos DB - Cassandra Account.
4. Click the **Database Credentials** button and select one of the following:
    - **Username and Password** - Provide the keyspace admin username and password. These credentials are required for leaf-level keyspaces. The account must have access to all databases within the Azure Cosmos DB - Cassandra account. For more information, see Username and Password.
    - **User Managed Identity** - Provide the Managed Identity Client ID. This uses authentication managed by Azure. For more information, see User Managed Identity (UMI).
      User Managed Identity must be attached to the Azure VM on which Cohesity DataProtect is deployed. UMI cannot be used to protect the assets across tenants.
    - **System Managed Identity** - Select the option.
      System Managed Identity must be enabled on the Azure VM on which Cohesity DataProtect is deployed and does not work across tenants. For more information, see System Managed Identity (SMI).
5. Click **Save**.

Protect Azure Cosmos DB - Cassandra Keyspaces

To protect your Azure Cosmos DB - Cassandra Account keyspaces:

1. Navigate to **Data Protection > Protection > Protect**, and then select **Azure > Cosmos DB - Cassandra**.
2. In the **New Protection** page, click **Add Objects** and select your registered source.

   Use the checkboxes to select the objects for protection. You can select the tenants, subscriptions, resource groups, Cosmos DB Cassandra accounts, or keyspaces under them. For more information, see Azure Cosmos DB - Cassandra Entity Hierarchy.

You can also use **List View** to select the instances for protection from a flat, searchable list.

You can auto-protect all the objects under a resource group, subscription, or tenant using the AutoProtect icon.

To exclude an auto-protected object, click on the AutoProtect icon next to the object.

3. Click **Continue**.
4. Click **Protection Group** and provide a name for your protection group.
5. Click **Policy** and choose a policy to specify backup frequency and retention. For more information on policy, see Create or Edit a Standard Policy.
6. Optionally, to configure any of the settings or additional settings, select **More Options**.
7. Click **Protect**.

You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific Azure Cosmos DB - Cassandra shows the history of all protection runs, including the one in progress.

Cohesity starts backing up the keyspaces you selected.

Settings

Setting	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Additional Settings

Pause Future Runs	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
End Date	If you need to end protection on a specific date, enable this option to select the date.
QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.
Priority	Select a priority for the Object Protection Group Execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Azure Cosmos DB - Cassandra Backup Configuration

Backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Cassandra.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ⋮ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Cassandra.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.

4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Cassandra.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Cassandra.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Azure Cosmos DB - Cassandra Keyspaces

After you protect your Azure Cosmos DB - Cassandra Keyspaces, you can recover them from Cohesity to their original or a new location.

To recover protected Azure Cosmos DB - Cassandra keyspaces:

1. Navigate to **Data Protection > Recoveries > Recover**, and then select **Azure > Cosmos DB - Cassandra**.
2. In the **New Recovery** page, use the filter, search box, and views to locate and select the required keyspaces.

    The latest snapshot will be displayed (protection run).

    If you need to recover an earlier snapshot, click the **Edit** icon to open the **Recovery Point** page.

    Click **List** to view the available recovery points by timestamp and click **Select Recovery Point**.
3. Click **Next: Recover Options**.
4. Under **Recover To**, select either of the following:
    - **Original Location** - To recover the Azure Cosmos DB - Cassandra keyspace to the Azure Cosmos DB - Cassandra instance or Managed Instance from which

the keyspace was backed up.

Recovering to the original location will replace any existing database with the same name in the selected instance. Type **YES** to confirm.

- **New Location** - Select a target Azure Cosmos DB - Cassandra instance or Managed Instance to which the Azure Cosmos DB - Cassandra keyspace will be recovered.

5. Select your **Recovery Options**:
   - **Rename**. Add a Prefix and/or Suffix to the full name of the recovered Azure Cosmos DB - Cassandra keyspace.
   - **Task Name**. Change the default name of the recovery task.
   - **Ingest Rate Cap**. Limit the data ingestion rate to 40% of theoretical maximum to avoid exceeding the provisioned throughput (RU/s) of the Azure Cosmos DB - Cassandra workload. Recommended for preventing rate limit errors.
     Ingest rates can be between 0.1 and 1.0. The default value is 0.4, which is 40% of the maximum theoretical ingest rate capacity.
6. Click **Recover**.

You can monitor the status of the recovery on the **Activity** page.

Cohesity starts recovering the selected Azure Cosmos DB - Cassandra keyspaces.

Azure Cosmos DB - MongoDB

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Azure Cosmos DB - MongoDB databases.

Azure Cosmos DB offers an API for MongoDB, which allows you to use Azure Cosmos DB like a native MongoDB database. This compatibility allows the use of existing MongoDB applications, tooling, drivers, and SDKs.

> **Important:** The backup and recovery of Azure PostgreSQL is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).

Plan and Prepare for Azure Cosmos DB - MongoDB Protection

Before protecting Azure Cosmos DB – MongoDB with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure Cosmos DB – MongoDB with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Authentication for Azure Cosmos DB - MongoDB Database Protection

You must authenticate to Azure Cosmos DB - MongoDB database from Cohesity DataProtect using one of the following methods:

- Username and Password
- User Managed Identity (UMI)
- System Managed Identity (SMI)

Username and Password

This method uses the credentials of the admin user for the Azure Cosmos DB - MongoDB database. This is the user account configured during server creation.

For backup, you can create a custom user. In the CLI and run the following commands:

1. To create login user on the server:

   ```
   #CREATE LOGIN [newlogin] WITH PASSWORD = 'somepassword;
   ```

2. To create same user under database which should be backed up by user **newlogin**:

   ```
   User <database>;
   #CREATE USER [newlogin] FOR LOGIN [newlogin];
   ALTER ROLE db_owner ADD MEMBER [newlogin];
   ```

3. Run the following command on the master system database:

   ```
   CREATE USER [newlogin] FOR LOGIN [newlogin];
   ```

For recovery, the default admin credentials created during server creation are required if the custom user credentials were set for authentication before backup. During recovery, you must change the credentials back to the default admin credentials.

User Managed Identity (UMI)

This method uses a user-assigned managed identity to authenticate to the Azure Cosmos DB - MongoDB database.

Perform the following to use the UMI authentication:

1. Create a user-assigned managed identity. For more information, see Manage user-assigned managed identities using the Azure portal.
2. Assign the **Managed Identity Operator** role to the UMI identity. For more information, see User-assigned managed identity.
3. Set the UMI as the identity for all the cluster nodes. Ensure that the UMI is from the same subscription as the one in which the database exists. For more information, see Assign a user-assigned managed identity to an existing VM.
4. Assign this User Managed Identity as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.

5. Assign the required user managed identity to Azure Cosmos DB - MongoDB database under the Identity section - User assigned managed identity on the Azure Portal.

System Managed Identity (SMI)

This method uses a unique identity that is automatically created by Azure when you enable SMI on the node. This identity belongs only to the specific node and is deleted if the node is deleted.

Each node in the cluster includes its own SMI (multiple distinct identities). Since the Azure Cosmos DB - MongoDB database allows only one Microsoft Entra admin per node, having multiple nodes, each with a different SMI may cause issues.

To mitigate this issue, you can perform the following:

1. Create a Microsoft Entra group.
2. Add SMI identities as group members.
3. Assign this group as Microsoft Entra Admin.

This ensures that all the nodes can be authenticated using their respective SMIs through the group.

> **Note:** SMIs are not created separately as a resource. You must enable it on the cluster nodes from the Azure portal.

Perform the following to use the SMI authentication:

1. Enable system-assigned managed identity on all the nodes of the cluster. For more information, see Enable system-assigned managed identity on an existing VM.
2. Create a Microsoft Entra basic group and add SMI's of all the nodes created in previous step as members. The SMI name is the same as the VM (node) name, which can be easily searched for while finding members to include in the group. For more information, see Create a basic group and add members.
3. Assign this Microsoft Entra basic group as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.

Considerations for Azure Cosmos DB - MongoDB Database Protection

Review and understand the following limitations and considerations before you protect Azure Cosmos DB - MongoDB database:

- **Account with vCore/DocumentDB cluster** - Discovery, protection, and recovery are not supported if the account is configured using the vCore/DocumentDB cluster.
- **Account with customized key** - Backup and recovery are not supported if the account is configured with a customized key.
- **MongoDB 3.2** - Azure Cosmos DB for MongoDB version 3.2 is not supported.

- **Naming rules** - The following are the rules for naming the databases:
  - The length of the database names must be between 3 and 63 characters.
  - Database names support all characters except #, /, ?,&, <, >, =, }, $, {,], [,", ', ., \ .
- **Version Downgrade not supported**: Recovering a backup taken from a higher database version to a lower version is not supported.

Backup Azure Cosmos DB - MongoDB Database

After registering your Azure source in the Cohesity cluster, Cohesity uses the cloud credentials provided to authenticate with Azure.

Once authenticated, Cohesity automatically discovers all the subscriptions that the service account can access, as well as all supported workloads such as Azure Cosmos DB - MongoDB databases within those subscriptions. The discovered workloads appear under **Data Protection > Sources**. You can select the **Cosmos DB - MongoDB** tab to view all the databases discovered.

Hierarchy View displays the entity hierarchy of the Azure Cosmos DB - MongoDB objects, allowing you to navigate from the resource group-level down to individual databases.

Azure Cosmos DB - MongoDB Entity Hierarchy

The Azure Cosmos DB - MongoDB entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Cosmos DB MongoDB Account	Fully managed Cosmos DB - MongoDB database service hosted within the Azure cloud.
Database	Fully managed NoSQL database in Microsoft Azure. The size of the databases is displayed only after valid credentials are set at the database account level and the source is refreshed.

Backup Scope by Azure Cosmos DB - MongoDB Hierarchy

Cohesity cluster allows you to backup Azure Cosmos DB - MongoDB databases at different levels of the hierarchy.

You can choose to back up:

- **Tenant** – Protect all Azure Cosmos DB - MongoDB databases within the tenant.
- **Subscription** – Protect all Azure Cosmos DB - MongoDB databases within the subscription.
- **Resource Group** – Protect all Azure Cosmos DB - MongoDB databases within the resource group.
- **Cosmos DB MongoDB Account** – Protect all Azure Cosmos DB - MongoDB databases within the account.
- **Database** – Protect individual Azure Cosmos DB - MongoDB database.

Supported Backup Types

Cohesity supports full backups for Azure Cosmos DB - MongoDB. Incremental backups are not supported.

Protect Azure Cosmos DB - MongoDB Databases

Azure Cosmos DB - MongoDB databases are protected using a Protection Group and a Protection Policy.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites to Protect Azure Cosmos DB - MongoDB

Set Azure Cosmos DB - MongoDB Database Credentials

For Backup, the database credentials are used to authenticate and connect to the Cosmos DB - MongoDB account. Once the connection is established, the relevant database can be dumped and protected.

For Azure Cosmos DB - MongoDB, you can set credentials at the Azure Cosmos DB - MongoDB account level. This information is used to connect to the database during backup and recovery.

To set the database credentials:

1. Navigate to **Data Protection > Sources**, find the registered Azure source, and click on it.
2. Click on the **Cosmos DB - MongoDB** tab.
3. Select the checkbox next to the required Azure Cosmos DB - MongoDB account.
4. Click the **Database Credentials** button and select one of the following:
   - **Username and Password** - Provide the database admin username and password. These credentials are required for leaf-level databases. The account must have access to all databases within the Azure Cosmos DB - MongoDB account. For more information, see Username and Password.
   - **User Managed Identity** - Provide the Managed Identity Client ID. This uses authentication managed by Azure.For more information, see User Managed Identity (UMI).
     User Managed Identity must be attached to the Azure VM on which Cohesity DataProtect is deployed. UMI cannot be used to protect the assets across tenants.
   - **System Managed Identity** - Select the option.
     System Managed Identity must be enabled on the Azure VM on which Cohesity DataProtect is deployed and does not work across tenants. For more information, see System Managed Identity (SMI).
5. Click **Save**.

Enable Server Side Retry (SSR)

You must enable Server Side Retry (SSR) to automate operation retries. This is required for handling rate limiting errors and addressing failure when the rate limit is reached.

Perform the following to enable SSR:

1. Sign in to the Azure portal.
2. Navigate to your Azure Cosmos DB - MongoDB account.
3. Go to the **Features** pane under the **Settings** section.
4. Select **Server Side Retry**.
5. Click **Enable** to enable this feature for all collections in your account.

For more information, see Prevent rate-limiting errors for Azure Cosmos DB for MongoDB operations.

Protect Azure Cosmos DB - MongoDB Databases

To protect your Azure Cosmos DB - MongoDB databases:

1. Navigate to **Data Protection > Protection > Protect**, and then select **Azure > Cosmos DB - MongoDB**.
2. In the **New Protection** page, click **Add Objects** and select your registered source.
3. Use the checkboxes to select the objects for protection. You can select the tenants, subscriptions, resource groups, Cosmos DB MongoDB accounts, or databases under them. For more information, see Azure Cosmos DB - MongoDB Entity Hierarchy

   You can also use **List View** to select the database entities for protection from a flat, searchable list.

   You can auto-protect all the objects under a resource group, subscription, or tenant using the AutoProtect icon .

   To exclude an auto-protected object, click on the AutoProtect icon next to the object.
4. Click **Continue**.
5. Click **Protection Group** and provide a name for your protection group.
6. Click **Policy** and choose a policy to specify backup frequency and retention. For more information on policy, see Create or Edit a Standard Policy.
7. Optionally, to configure any of the settings or additional settings, select More Options.
8. Click **Protect**.

You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific Azure Cosmos DB - MongoDB shows the history of all protection runs, including the one in progress.

Cohesity starts backing up the databases you selected.

Settings

Setting	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.

Setting	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Additional Settings

Pause Future Runs	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
End Date	If you need to end protection on a specific date, enable this option to select the date.
QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.

Priority	Select a priority for the Object ProtectionGroupxecution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Azure Cosmos DB - MongoDB Backup Configuration

Backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - MongoDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - MongoDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - MongoDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - MongoDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Azure Cosmos DB - MongoDB Databases

After you protect your Azure Cosmos DB - MongoDB databases, you can recover them from Cohesity to their original or a new location.

Prerequisites to Recover Azure Cosmos DB - MongoDB

Enable Server Side Retry (SSR)

You must enable Server Side Retry (SSR) to automate operation retries. This is required for handling rate limiting errors and addressing failure when the rate limit is reached.

Perform the following to enable SSR:

1. Sign in to the Azure portal.
2. Navigate to the target Azure Cosmos DB - MongoDB account.
3. Go to the **Features** pane under the **Settings** section.
4. Select **Server Side Retry**.
5. Click **Enable** to enable this feature for all collections in your account.

For more information, see Prevent rate-limiting errors for Azure Cosmos DB for MongoDB operations.

Recover Azure Cosmos DB - MongoDB Databases

To recover protected Azure Cosmos DB - MongoDB databases:

1. Navigate to **Data Protection > Recoveries > Recover**, and then select **Azure > Cosmos DB - MongoDB**.
2. In the **New Recovery** page, use the filter, search box, and views to locate and select the required databases.

   The latest snapshot will be displayed (protection run).

   If you need to recover an earlier snapshot, click the **Edit** icon to open the **Recovery Point** page.
3. Click **List** to view the available recovery points by timestamp and click **Select Recovery Point**.
4. Click **Next: Recover Options**.
5. Under **Recover To**, select either of the following:
   - **Original Location** - To recover the Azure Cosmos DB - MongoDB database to the same Azure Cosmos DB - MongoDB account from which the database was backed up.

     If overwrite is enabled, recovering to the original location will replace any existing database with the same name in the selected account. Type **YES** to confirm.
   - **New Location** - Select a target Azure Cosmos DB - MongoDB account to which the Azure Cosmos DB - MongoDB database will be recovered. Toggle on **Overwrite existing database(s)** to allow Cohesity to overwrite the existing databases in the target location. Type **YES** to confirm.
6. Select your **Recovery Options**:
   - **Rename**. Add a Prefix and/or Suffix to the full name of the recovered Azure Cosmos DB - MongoDB database.
   - **Task Name**. Change the default name of the recovery task.
7. Click **Recover**.

You can monitor the status of the recovery on the **Activity** page.

Cohesity starts recovering the selected Azure Cosmos DB - MongoDB databases.

Azure Cosmos DB - NoSQL

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Azure Cosmos DB - NoSQL databases.

Azure Cosmos DB - NoSQL is a fully managed, serverless, and globally distributed NoSQL database service offered by Microsoft Azure. It is designed for modern application development, including AI applications and agents, providing high performance, availability, and scalability for handling massive volumes of NoSQL and vector data.

> **Important:** The backup and recovery of Azure PostgreSQL is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).

Plan and Prepare for Azure Cosmos DB – NoSQL Protection

Before protecting Azure Cosmos DB – NoSQL with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure Cosmos DB – NoSQL with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Authentication for Azure Cosmos DB - NoSQL Database Protection

You must authenticate to Azure Cosmos DB - NoSQL database from Cohesity DataProtect using one of the following methods:

- Username and Password
- User Managed Identity (UMI)
- System Managed Identity (SMI)

Username and Password

This method uses the credentials of the admin user for the Azure Cosmos DB - NoSQL database. This is the user account configured during server creation.

For backup, you can create a custom user. In the CLI and run the following commands:

1. To create login user on the server:

   ```
   #CREATE LOGIN [newlogin] WITH PASSWORD = 'somepassword;
   ```

2. To create same user under database which should be backed up by user **newlogin**:

```
User <database>;
#CREATE USER [newlogin] FOR LOGIN [newlogin];
ALTER ROLE db_owner ADD MEMBER [newlogin];
```

3. Run the following command on the master system database:

```
CREATE USER [newlogin] FOR LOGIN [newlogin];
```

For recovery, the default admin credentials created during server creation are required if the custom user credentials were set for authentication before backup. During recovery, you must change the credentials back to the default admin credentials.

User Managed Identity (UMI)

This method uses a user-assigned managed identity to authenticate to the Azure Cosmos DB - NoSQL database.

Perform the following to use the UMI authentication:

1. Create a user-assigned managed identity. For more information, see Manage user-assigned managed identities using the Azure portal.
2. Assign the **Managed Identity Operator** role to the UMI identity. For more information, see User-assigned managed identity.
3. Set the UMI as the identity for all the cluster nodes. Ensure that the UMI is from the same subscription as the one in which the database exists. For more information, see Assign a user-assigned managed identity to an existing VM.
4. Assign this User Managed Identity as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.
5. Assign the required user managed identity to Azure Cosmos DB - NoSQL database under the Identity section - User assigned managed identity on the Azure Portal.

System Managed Identity (SMI)

This method uses a unique identity that is automatically created by Azure when you enable SMI on the node. This identity belongs only to the specific node and is deleted if the node is deleted.

Each node in the cluster includes its own SMI (multiple distinct identities). Since the Azure Cosmos DB - NoSQL database allows only one Microsoft Entra admin per node, having multiple nodes, each with a different SMI may cause issues.

To mitigate this issue, you can perform the following:

1. Create a Microsoft Entra group.
2. Add SMI identities as group members.
3. Assign this group as Microsoft Entra Admin.

This ensures that all the nodes can be authenticated using their respective SMIs through the group.

> **Note:** SMIs are not created separately as a resource. You must enable it on the cluster nodes from the Azure portal.

Perform the following to use the SMI authentication:

1. Enable system-assigned managed identity on all the nodes of the cluster. For more information, see Enable system-assigned managed identity on an existing VM.
2. Create a Microsoft Entra basic group and add SMI's of all the nodes created in previous step as members. The SMI name is the same as the VM (node) name, which can be easily searched for while finding members to include in the group. For more information, see Create a basic group and add members.
3. Assign this Microsoft Entra basic group as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.

Considerations

Review and understand the following limitations and considerations before you protect Azure Cosmos DB - NoSQL Database:

- **Full backup support** - Only full backup is supported; incremental backup is not available for Azure Cosmos DB - NoSQL protection.
- **Attach metadata and credentials** - When you recover a database to a different account, ensure that you attach the metadata and credentials to the destination account.
- **Rename database during recovery** - Database names can have up to 48 alpha-numeric characters and can contain underscores; only letters and numbers are supported as the first character.
- **Database-Level Protection** - Protection occurs at the database level and not account level; backup and recovery are performed for the databases on the account. Cohesity does not change or create an Azure Cosmos DB-NoSQL account in the workflows. Hence, the properties of the account like tags are not recovered.
- **Serverless vs provisioned throughput** - Database belonging to an account with Serverless mode cannot be recovered to the one with provisioned throughput capacity mode and vice versa.
- **Partition key configuration** - During recovery, the partition key configuration of the source container must match the destination. The partition key cannot be changed during recovery operations.
- **Indexing policy preservation** - The indexing policy from the source container is preserved and applied to the recovered container. Custom indexing policies are maintained during recovery.

- **Unique key constraints** - Unique key constraints defined on the source container are preserved during recovery. Ensure the destination account supports the same unique key configurations.
- **TTL (Time-to-live) on recovered containers** - When recovering Azure Cosmos DB-NoSQL containers, the default TTL setting is preserved from the source container and applied to the recovered container. As a result, any newly inserted items in the recovered container will automatically expire after a duration of Item creation time + default TTL (in seconds).

  This ensures consistency with the TTL-based data expiration behavior of the original container.

  > **Note:** Item-level TTL settings are not retained during backup. When data is recovered, the TTL for those items is set to the container's default TTL value.

- **Stored procedures, triggers, and UDFs** - Stored procedures, triggers, and user-defined functions (UDFs) defined on the source container are backed up and recovered to the destination container. Ensure the destination account has the necessary permissions to execute these server-side logic components.
- **Change feed configuration** - Change feed settings are not preserved during recovery. After recovery, you may need to reconfigure the change feed processors and their lease containers.
- **Analytical store** - If the source container has an analytical store enabled, the destination container must also support analytical store. This feature must be enabled at the account level before recovery.
- **Replication regions** - To get optimum performance during recovery, it is recommended to disable multi-region replication. After recovery, you can re-enable multi-region replication.
- **Custom role permissions** - For backup and recovery using custom roles, ensure the service principal has the following minimum data plane permissions:
  - **Backup**: readMetadata, sqlDatabases/read, containers/read, items/read, executeQuery, readChangeFeed.
  - **Recovery**: readMetadata, sqlDatabases/read, sqlDatabases/write, containers/read, containers/write, items/create, items/upsert.
- **Rate limiting and throttling** - During backup and recovery, monitor the HTTP 429 (Too Many Requests) responses. Adjust the ingest rate or increase provisioned throughput temporarily to avoid throttling.
- **Container throughput** - Ensure the destination container has sufficient provisioned throughput (RU/s) to handle recovery. Insufficient throughput will result in throttling and extended recovery times.

- **Hierarchical partition keys** - If the source container uses hierarchical partition keys (multi-level partitioning), the destination account must support this feature. Hierarchical partition keys are preserved during recovery.
- **Version Downgrade not supported**: Recovering a backup taken from a higher database version to a lower version is not supported.

Backup Azure Cosmos DB - NoSQL Database

After registering your Azure Cosmos DB - NoSQL source in Cohesity, Cohesity uses the cloud credentials provided to authenticate with Azure.

Once authenticated, Cohesity automatically discovers all the subscriptions that the service account can access, as well as all supported workloads such as Azure Cosmos DB - NoSQL databases within those subscriptions. The discovered workloads appear under **Data Protection > Sources**. You can select the **Cosmos DB - NoSQL** tab to view all the databases discovered.

Hierarchy View displays the entity hierarchy of the Azure Cosmos DB - NoSQL objects, allowing you to navigate from the resource group-level down to individual databases.

Azure Cosmos DB - NoSQL Entity Hierarchy

The Azure Cosmos DB - NoSQL entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Cosmos DB NoSQL Account	Fully managed Cosmos DB - NoSQL database service hosted within the Azure cloud.
Database	Fully managed NoSQL database in Microsoft Azure.

Backup Scope by Azure Cosmos DB - NoSQL Hierarchy

Cohesity cluster allows you to backup Azure Cosmos DB - NoSQL databases at different levels of the hierarchy.

You can choose to back up:

- **Tenant** – Protect all Azure Cosmos DB - NoSQL databases within the tenant.
- **Subscription** – Protect all Azure Cosmos DB - NoSQL databases within the subscription.
- **Resource Group** – Protect all Azure Cosmos DB - NoSQL databases within the resource group.
- **Cosmos DB - NoSQL Account** – Protect all Azure Cosmos DB - NoSQL databases within the Cosmos DB – NoSQL account.
- **Database** – Protect individual Azure Cosmos DB - NoSQL database.

Supported Backup Types

Cohesity supports only full backups for Azure Cosmos DB - NoSQL. Incremental backups are not supported.

Protect Azure Cosmos DB - NoSQL Databases

Azure Cosmos DB - NoSQL databases are protected using a Protection Group and a Protection Policy.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites to Protect Azure Cosmos DB – NoSQL

Set Azure Cosmos DB - NoSQL Database Credentials

For Backup, the database credentials are used to authenticate and connect to the Cosmos DB - NoSQL account. Once the connection is established, the relevant database can be dumped and protected.

For Azure Cosmos DB - NoSQL, you can set credentials at the Azure Cosmos DB - NoSQL account level. This information is used to connect to the database during backup and recovery.

To set the database credentials:

1. Navigate to **Data Protection > Sources**, find the registered Azure source, and click on it.
2. Click on the **Cosmos DB - NoSQL** tab.
3. Select the checkbox next to the required Azure Cosmos DB - NoSQL account.
4. Click the **Database Credentials** button and select one of the following:
   - **Username and Password** - Provide the database admin username and password. These credentials are required for leaf-level databases. The account must have access to all databases within the Azure Cosmos DB - NoSQL account. For more information, see Username and Password.
   - **User Managed Identity** -Provide the Managed Identity Client ID. This uses authentication managed by Azure. For more information, see User Managed Identity (UMI).
     User Managed Identity must be attached to the Azure VM on which Cohesity DataProtect is deployed. UMI cannot be used to protect the assets across tenants.
   - **System Managed Identity** - Select the option.
     System Managed Identity must be enabled on the Azure VM on which Cohesity DataProtect is deployed and does not work across tenants. For more information, see System Managed Identity (SMI).

   > **Note:** Always ensure to add credentials to the Cosmos DB - NoSQL account. Do not add credentials to individual databases.
5. Click **Save**.

Protect Azure Cosmos DB - NoSQL Databases

To protect your Azure Cosmos DB - NoSQL databases:

1. Navigate to **Data Protection > Protection > Protect**, and then select **Azure > Cosmos DB - NoSQL**.
2. In the **New Protection** page, click **Add Objects** and select your registered source.
3. Use the checkboxes to select the objects for protection. You can select the tenants, subscriptions, resource groups, Cosmos DB - NoSQL account, or databases under them. For more information, see Azure Cosmos DB - NoSQL Entity Hierarchy.

   You can also use **List View** to select the instances for protection from a flat, searchable list.

   You can auto-protect all the objects under a resource group, subscription, or tenant using the AutoProtect icon .

To exclude an auto-protected object, click on the AutoProtect icon next to the object.

4. Click **Continue**.
5. Click **Protection Group** and provide a name for your protection group.
6. Click **Policy** and choose a policy to specify backup frequency and retention. If you do not have a policy, you can easily create one.
7. Optionally, to configure any of the settings or additional settings, select More Options.
8. Click **Protect**.

You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific Azure Cosmos DB - NoSQL shows the history of all protection runs, including the one in progress.

Cohesity starts backing up the databases you selected.

Settings

Setting	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Additional Settings

Pause Future Runs	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
End Date	If you need to end protection on a specific date, enable this option to select the date.

QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.
Priority	Select a priority for the Object ProtectionGroupxecution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Azure Cosmos DB - NoSQL Backup Configuration

Backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - NoSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ⋮ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - NoSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - NoSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - NoSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Azure Cosmos DB - NoSQL Databases

After you protect your Azure Cosmos DB - NoSQL databases, you can recover them from Cohesity to their original or a new location.

To recover protected Azure Cosmos DB - NoSQL databases:

1. Navigate to **Data Protection > Recoveries > Recover**, and then select **Azure > Cosmos DB - NoSQL**.
2. In the **New Recovery** page, use the filter, search box, and views to locate and select the required databases.

    The latest snapshot will be displayed (protection run).

    If you need to recover an earlier snapshot, click the **Edit** icon to open the **Recovery Point** page.
3. Click **List** to view the available recovery points by timestamp and click **Select Recovery Point**.
4. Click **Next: Recover Options**.
5. Under **Recover To**, select either of the following:
    - **Original Location** - To recover the Azure Cosmos DB - NoSQL database to the Azure Cosmos DB - NoSQL account from which the database was backed up.

      Recovering to the original location will replace any existing database with the same name in the selected instance. Type YES to confirm.
    - **New Location** - Select a target Azure Cosmos DB - NoSQL account to which the Azure Cosmos DB - NoSQL database will be recovered.
6. Select your **Recovery Options**:
    - **Rename**. Add a Prefix and/or Suffix to the full name of the recovered Azure Cosmos DB - NoSQL database.
    - **Task Name**. Change the default name of the recovery task.
7. Click **Recover**.

You can monitor the status of the recovery on the **Activity** page.

Cohesity starts recovering the selected Azure Cosmos DB - NoSQL databases.

Azure Cosmos DB - Table

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Azure Cosmos DB - Table.

Azure Cosmos DB -Table is a fully managed NoSQL database service that provides a highly scalable and globally distributed store for key-value and tabular data, using the familiar Azure Table storage APIs. It is ideal for applications that require fast, high-performance storage for structured, non-relational data, like device registries, user profiles, or configuration data.

> **Important:** The backup and recovery of Azure PostgreSQL is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).

Plan and Prepare for Azure Cosmos DB - Table Protection

Before protecting Azure Cosmos DB - Table with the Cohesity cluster, ensure you have reviewed the considerations described in this topic.

Before You Begin

To protect Azure Cosmos DB - Table with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Considerations for Azure Cosmos DB - Table Protection

- **Recovery Failure from Replicated Cluster** - Recovery from a replicated cluster of Azure Cosmos DB-Table may fail with the following error: "kStale: File previously opened with higher session ID: <xxx_yyy>".
- **Recovery Validation UI Warning** - During recovery validation, the following warning may be displayed in the UI: "Metadata view RecordIO client not available". This warning does not impact the actual recovery process.
- **Version Downgrade not supported**: Recovering a backup taken from a higher database version to a lower version is not supported.

Backup Azure Cosmos DB - Table

After registering your Azure Cosmos DB - Table source, Cohesity discovers all the subscriptions that the service account can access, as well as all supported workloads such as Azure Cosmos DB - Table tables within those subscriptions. The discovered workloads appear under **Data Protection > Sources**. You can select the **Cosmos DB - Table** tab to view all the tables discovered.

Hierarchy View displays the entity hierarchy of the Azure Cosmos DB - Table objects, allowing you to navigate from the resource group-level down to individual tables.

Azure Cosmos DB - Table Entity Hierarchy

The Azure Cosmos DB - Table entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Cosmos DB Table Account	Fully managed Cosmos DB - Table database service hosted within the Azure cloud.
Table	Fully managed Table tables in Microsoft Azure.

Backup Scope by Azure Cosmos DB - Table Hierarchy

Cohesity cluster allows you to backup Azure Cosmos DB - Table at different levels of the hierarchy.

You can choose to back up:

- **Tenant** – Protect all Azure Cosmos DB - Table tables within the tenant.
- **Subscription** – Protect all Azure Cosmos DB - Table tables within the subscription.
- **Resource Group** – Protect all Azure Cosmos DB - Table tables within the resource group.
- **Cosmos DB Table Account** – Protect all Azure Cosmos DB - Table tables within the Cosmos DB Account.
- **Tables** – Protect individual Azure Cosmos DB - Table tables.

Supported Backup Types

Cohesity supports full backups for Azure Cosmos DB - Table. Incremental backups are not supported.

Protect Azure Cosmos DB - Table

Azure Cosmos DB - Table tables are protected using a Protection Group and a Protection Policy.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Protect Azure Cosmos DB - Table

To protect your Azure Cosmos DB - Table tables:

1. Navigate to **Data Protection > Protection**, click **Protect**, and then select **Azure > Cosmos DB - Table**.
2. In the **New Protection** page, click **Add Objects** and select your registered source.
3. Use the checkboxes to select the objects for protection. You can select the tenants, subscriptions, resource groups, Cosmos DB – Table account, or tables under them. For more information, see Azure Cosmos DB - Table Entity Hierarchy.

   You can also use **List View** to select the instances for protection from a flat, searchable list.

   You can auto-protect all the objects under a resource group, subscription, or tenant using the AutoProtect icon.

   To exclude an auto-protected object, click on the AutoProtect icon next to the object.
4. Click **Continue**.
5. Click **Protection Group** and provide a name for your protection group.
6. Click **Policy** and choose a policy to specify backup frequency and retention. For more information on policy, see Create or Edit a Standard Policy.
7. Optionally, to configure any of the settings or additional settings, select **More Options**.
8. Click **Protect**.

You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific Azure Cosmos DB - Table shows the history of all protection runs, including the one in progress.

Cohesity starts backing up the tables you selected.

Settings

Setting	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Additional Settings

Pause Future Runs	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
End Date	If you need to end protection on a specific date, enable this option to select the date.
QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.

Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.
Priority	Select a priority for the Object ProtectionGroupxecution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Azure Cosmos DB - Table Backup Configuration

Backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Table.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ⋮ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Table.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Table.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure Cosmos DB - Table.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.

4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Azure Cosmos DB - Table

After you protect your Azure Cosmos DB - Table, you can recover them from Cohesity to their original or a new location.

To recover protected Azure Cosmos DB - Table tables:

1. Navigate to **Data Protection > Recoveries > Recover**, and then select **Azure > Cosmos DB - Table**.
2. In the **New Recovery** page, use the filters, search box, and views to locate and select the required tables.

    The latest snapshot will be displayed (protection run).

    If you need to recover an earlier snapshot, click the **Edit** icon to open the **Recovery Point** page.

    Click **List** to view the available recovery points by timestamp and click **Select Recovery Point**.
3. Click **Next: Recover Options**.
4. Under **Recover To**, select either of the following:
    - **Original Location** - To recover the Azure Cosmos DB - Table tables to the Azure Cosmos DB - Table instance or Managed Instance from which the table was backed up.

        Recovering to the original location will replace any existing table with the same name in the selected instance. Type YES to confirm.
    - **New Location** - Select a target Azure Cosmos DB - Table instance or Managed Instance to which the Azure Cosmos DB - Table will be recovered.
5. Select your **Recovery Options**:
    - **Rename**. Add a Prefix and/or Suffix to the full name of the recovered Azure Cosmos DB - Table table.
    - **Task Name**. Change the default name of the recovery task.
6. Click **Recover**.

You can monitor the status of the recovery on the **Activity** page.

Cohesity starts recovering the selected Azure Cosmos DB - Table tables.

Azure SQL Database

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Azure SQL databases.

Azure SQL database is a fully managed, intelligent, relational database-as-a-service (DBaaS) offered by Microsoft Azure. It provides a cloud-based SQL Server database engine, optimized for modern cloud applications and websites.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Plan and Prepare for Azure SQL Protection

Before protecting Azure SQL databases with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure SQL databases with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Authentication for Azure SQL Database Protection

You must authenticate to Azure SQL database from the Cohesity cluster using one of the following methods:

- Username and Password
- User Managed Identity (UMI)
- System Managed Identity (SMI)

Username and Password

This method uses the credentials of the admin user for the SQL server. This is the user account configured during server creation.

For backup, you can create a custom user. Connect to the SQL server using SQL management studio or SQLCMD and run the following commands:

1. To create login user on the server:

```
#CREATE LOGIN [newlogin] WITH PASSWORD = 'somepassword;
```

2. To create same user under database which should be backed up by user **newlogin**:

```
User <database>;
#CREATE USER [newlogin] FOR LOGIN [newlogin];
ALTER ROLE db_owner ADD MEMBER [newlogin];
```

3. Run the following command on the master system database:

```
CREATE USER [newlogin] FOR LOGIN [newlogin];
```

For recovery, the default admin credentials created during server creation are required if the custom user credentials were set for authentication before backup. During recovery, you must change the credentials back to the default admin credentials.

User Managed Identity (UMI)

This method uses a user-assigned managed identity to authenticate to the Azure SQL database.

Perform the following to use the UMI authentication:

1. Create a user-assigned managed identity. For more information, see Manage user-assigned managed identities using the Azure portal.
2. Assign the **Managed Identity Operator** role to the UMI identity. For more information, see User-assigned managed identity.
3. Set the UMI as the identity for all the cluster nodes. Ensure that the UMI is from the same subscription as the one in which the SQL database exists. For more information, see Assign a user-assigned managed identity to an existing VM.
4. Assign this User Managed Identity as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.
5. Assign the required user managed identity to Azure SQL instance/server under the Identity section - User assigned managed identity on the Azure Portal.

System Managed Identity (SMI)

This method uses a unique identity that is automatically created by Azure when you enable SMI on the node. This identity belongs only to the specific node and is deleted if the node is deleted.

Each node in the cluster includes its own SMI (multiple distinct identities). Since the Azure SQL database allows only one Microsoft Entra admin per node, having multiple nodes, each with a different SMI may cause issues.

To mitigate this issue, you can perform the following:

1. Create a Microsoft Entra group.
2. Add all the nodes' SMI identities as group members.
3. Assign this group as Microsoft Entra Admin.

This ensures that all the nodes can be authenticated using their respective SMIs through the group.

> **Note:** SMIs are not created separately as a resource. You must enable it on the cluster nodes from the Azure portal.

Perform the following to use the SMI authentication:

1. Enable system-assigned managed identity on all the nodes of the cluster. For more information, see Enable system-assigned managed identity on an existing VM.
2. Create a Microsoft Entra basic group and add SMI's of all the nodes created in previous step as members. The SMI name is the same as the VM (node) name, which can be easily searched for while finding members to include in the group. For more information, see Create a basic group and add members.
3. Assign this Microsoft Entra basic group as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.

Considerations for Azure SQL Protection

Review and understand the following limitations and considerations before you protect Azure SQL database:

- **Database-level configurations** - Configuring credentials at database level is not supported.
- **Column-level encryption** - Database protection does not support column-level encryption.
- **Database Tables** - Backup and recovery of the following tables in the databases are not supported:
  - GRAPH_TABLE,
  - TEMPORAL_TABLE,
  - LEDGER_TABLE,
  - MEMORY_OPTIMIZED_TABLE
- **Backup failure with invalid object name error**: Subsequent full/incremental backups of an Azure SQL Server backup may fail with the following error: "Invalid object name 'cdc.<schema_name>_<table_name>_CT'"

  This issue occurs while executing the backup query script with a transfer object in the query (to transfer table from one schema to another).

  This may cause a mismatch between the user tables and their associated CDC tables. The mismatch typically arises after the schema changes or when database objects are moved between schemas after CDC is enabled.

  Workaround:

  Perform the following:

1. Disable CDC on the database before running the backup.
2. Login to the SQL database using the following command:

   ```
   sqlcmd -S hostname -U user -d database_name
   ```

3. Run the following command on the SQL database:

   ```
   # EXEC sys.sp_cdc_disable_db
   ```

4. Perform a full backup. The CDC will be automatically enabled during the next backup.

- **Incremental backup fall back to full backup**: Incremental backups fall back to full backup in the following cases:
  - Schema changes are detected during backup.
  - The state file from an earlier backup is not available.
  - CDC is enabled/disabled on the database/tables.
  - Previous full backup does not exist.
- **Insufficient tempDB storage** - If a backup fails due to insufficient tempDB storage on the SQL server, perform the following:
  1. Login to the Azure portal.
  2. Navigate to **Protection Group > Additional Setting > CDC processing location > SQLDB server or Client**.
  3. Select the **Client-Side** processing for CDC processing.
- **VNet settings for Azure SQL** - An Azure VM used as a DataProtect node must be present in the same VNet as that of the Azure-managed instance. Alternatively, if the DataProtect node and SQL-managed instance are in different VNet, then both the VNets must peer to access the database instance.
- **Backup failure due to Readlock** - Backup fails when Readlock is placed on the database or resource group.
- For Azure databases, if a backup is generated by the native backup database workflow, using TDE enabled by Customer-managed Keys or TDE disabled, these types of tables are supported.
- **Database diagrams** - Database diagrams are not recovered.
- **Mandatory client certificate** - Backup and recovery are not supported for the databases that mandatorily require client certificates for connection to DataProtect.
- **Default ports** - Only default ports are supported.
- **Database names with special characters** - Backup and recovery of databases with names containing the characters '#' and '/' are not supported. Also, the database name should adhere to the naming conventions suggested by the cloud vendors.

- **Passwords with special characters** – Semicolon (;) is not supported in server or database passwords.
- **Subscriptions of different tenants** - Azure Managed Identity is not supported across subscriptions of different tenants.
- **Version Downgrade not supported**: Recovering a backup taken from a higher database version to a lower version is not supported.
- **Duplicate Records**: Recovered database may contain duplicate records if the source database does not include primary or foreign keys."
- **Incremental Backup Failure**: Incremental backup for Azure SQL database fails when the **tempdb** storage is exhausted due to a large number of CDC transactions on the database. To avoid this, use the batch count to limit the number of transactions being processed at a time.
- **CDC Job Cleanup**: Full backups can fail during CDC processing if CDC jobs are not cleaned at the SQL database level. This issue occurs when a database is deleted and then recreated with the same name.

  Workaround:

  You can manually perform the cleanup and disable CDC on database before backup. For more details, see Cleanup job.

  > **Note:** The CDC cleanup job removes all data from the change tracking tables based on the configured retention period.

- **CDC Performance**: For optimal CDC performance, ensure the following:
  - Do not use DDL triggers that block schema changes, like CREATE/DROP TABLE, CREATE/DROP FUNCTION.
  - Ensure that the SQL Server Agent service is running.
  - Use an account with db_owner or sysadmin privileges for CDC setup.
  - To avoid capture_instance conflicts, check if the table is already CDC-enabled.
- **Database Tables**: If the databases contain any of the following types of tables, the backup fails due to CDC limitations:
  - Graph tables
  - Temporal tables
  - Ledger tables (updatable ledger)
  - Memory optimized table (business critical tier only)
  - These types of tables are supported for Azure SQL databases, if the backup

uses the native backup database workflow, using TDE enabled by Customer-Managed Keys or TDE disabled.

- **SqlPackage Utility**: Consider the following for the SqlPackage utility:
  - SqlPackage may fail to export the following SQL Server objects: Logins, Jobs (SQL Agent Jobs), Linked Servers, Server-level permissions, Server roles, Credentials, Database Mail configuration, Replication settings, SSIS packages, Service Broker objects, CLR assemblies, Data in FILESTREAM/FILETABLE, Encrypted objects (if keys/certs are not exported), Cross-database dependencies, extended stored procedures, and some system objects.
  - SqlPackage primarily exports database-level objects and data, but not server-level or agent-level configurations. For a complete migration, use additional tools or scripts to handle these objects.

Backup Azure SQL Database

After registering your Azure source in the Cohesity cluster, Cohesity uses the cloud credentials provided to authenticate with Azure.

Once authenticated, Cohesity automatically discovers all the subscriptions that the service account can access, as well as all supported workloads such as Azure SQL databases within those subscriptions. The discovered workloads appear under **Data Protection > Sources**. You can select the **SQL** tab to view all the databases discovered.

Hierarchy View displays the entity hierarchy of the Azure SQL objects, allowing you to navigate from the resource group-level down to individual databases.

Azure SQL Database Entity Hierarchy

The Azure SQL entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.

Structure	Description
Azure SQL Server	Fully managed SQL database instance hosted within the Azure cloud.
Database	Fully managed relational database service in Microsoft Azure.

Backup Scope by Azure SQL Hierarchy

Cohesity cluster allows you to backup Azure SQL databases at different levels of the hierarchy.

You can choose to back up:

- **Tenant** – Protect all Azure SQL databases within the tenant.
- **Subscription** – Protect all Azure SQL databases within the subscription.
- **Resource Group** – Protect all Azure SQL databases within the resource group.
- **Azure SQL Server** – Protect all Azure SQL databases within the database server.
- **Database** – Protect individual Azure SQL database.

Supported Backup Types

Cohesity supports full backups and Change Data Capture (CDC)-based incremental backups for Azure SQL.

Protect Azure SQL Databases

Azure SQL databases are protected using a Protection Group and a Protection Policy.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites to Protect Azure SQL

Set Azure SQL Database Credentials

Before backing up the Azure SQL databases that are discovered in Cohesity DataProtect, you must configure credentials for the databases. These credentials enable secure authentication and connectivity between Cohesity cluster and Azure SQL databases.

Consider the following:

- You must use the credentials that you have already configured for the Managed Instances in your Azure environment. Once these credentials are validated and the connection is established, Cohesity can access the databases within the instance to perform operations such as dumping, protecting, and recovering.
- You can set credentials for the SQL database at the following levels in the entity hierarchy: Tenant, Subscription, Resource group, Azure SQL server. You cannot set credentials for individual databases.
- The credentials you set are applied to all child objects within the selected object.
- It is recommended to apply credentials at the Azure SQL server level. If you apply credentials at a higher level in the entity hierarchy (Resource group, Subscription, or Tenant), they will be applied to all child-level databases. If the credentials do not match any of the databases at the child-level, the backup of the databases may fail.

To set the database credentials:

1. Navigate to **Data Protection > Sources**, find the registered Azure source, and click on it.
2. Click on the **SQL** tab.
3. Select the checkbox next to the required Azure SQL Server.
4. Click the **Database Credentials** button and select one of the following:
    - **Username and Password** - Provide the database admin username and password. These credentials are required for leaf-level databases. The account must have access to all databases within the SQL server. For more information, see Username and Password.
    - **User Managed Identity** - Provide the Managed Identity Client ID and Subscription ID. This uses authentication managed by Azure. The Subscription ID of UMI is optional. If not specified, the UMI is assumed to belong to the same subscription as the database server. For more information, see User Managed Identity (UMI).
      User Managed Identity must be attached to the Azure VM on which Cohesity DataProtect is deployed. UMI cannot be used to protect the assets across tenants.
    - **System Managed Identity** - Select the option.
      System Managed Identity must be enabled on the Azure VM on which Cohesity DataProtect is deployed and does not work across tenants. For more information, see System Managed Identity (SMI).
5. Click **Save**.

Protect Azure SQL Databases

To protect your Azure SQL databases:

1. Navigate to **Data Protection > Protection > Protect**, and then select **Azure > SQL**.
2. In the **New Protection** page, click **Add Objects** and select your registered source.
3. Use the checkboxes to select the objects for protection. You can select the tenants, subscriptions, resource groups, SQL servers, or databases under them. For more information, see Azure SQL Database Entity Hierarchy.

   You can also use List View to select the instances for protection from a flat, searchable list.

   You can auto-protect all the objects under a resource group, subscription, or tenant using the AutoProtect icon .

   To exclude an auto-protected object, click on the AutoProtect icon next to the object.
4. Click **Continue**.
5. Click **Protection Group** and provide a name for your protection group.
6. Click **Change Data Capture (CDC) for Backups** and toggle on the **Use CDC During Backups** option to use CDC data for database backups.

   > **Note:** CDC data is required (mandatory) for database backups. This option must be enabled to proceed with the backup.
7. Click **Policy** and choose a policy to specify backup frequency and retention. For more information on policy, see Create or Edit a Standard Policy.
8. Optionally, to configure any of the settings or additional settings, select More Options.
9. Click **Protect**.

You can monitor the status of the backup on the **Activity** page. Also, the **Activity** tab of a specific Azure SQL shows the history of all protection runs, including the one in progress.

Cohesity starts backing up the databases you selected.

Settings

Setting	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.

Setting	Description
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Change Data Capture (CDC) for Backups	Toggle on the **Use CDC During Backups** option to use CDC data for database backups. This option is enabled by default, and the value is set to **Yes**. > **Note:** CDC data is required (mandatory) for database backups. This option must be enabled to proceed and perform backups.

Additional Settings

CDC Data Processing Location	Select the location where CDC data processing is performed: • **SQL DB Server** - CDC data processing is performed on the server, using temporary tables. • **Cohesity Cluster** - CDC data processing is performed on the client (Cohesity cluster nodes).
CDC Batch Count	Enter the number of CDC records processed in batches to prevent consuming the entire temporary database. Set this value based on the server capacity, specific to each instance. The default value is **-1** indicating unlimited (no batching). > **Note:** This option is displayed only if **Change Data Capture (CDC) for Backups** is enabled, and **CDC Data Processing Location** is set as **SQL DB Server**.
Pause Future Runs	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
End Date	If you need to end protection on a specific date, enable this option to select the date.

QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.
Priority	Select a priority for the Object Protection Grou execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Azure SQL Backup Configuration

Backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure SQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure SQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure SQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure SQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Azure SQL Databases

After you protect your Azure SQL databases, you can recover them from Cohesity to their original or a new location.

To recover protected Azure SQL databases:

1. Navigate to **Data Protection > Recoveries > Recover**, and then select **Azure > SQL**.
2. In the **New Recovery** page, use the filter, search box, and views to locate and select the required databases.

    The latest snapshot will be displayed (protection run).

    If you need to recover an earlier snapshot, click the **Edit** icon to open the **Recovery Point** page.

    Click **List** to view the available recovery points by timestamp and click **Select Recovery Point**.
3. Click **Next: Recover Options**.
4. Under **Recover To**, select either of the following:
    - **Original Location** - To recover the Azure SQL database to the Azure SQL instance or Managed Instance from which the database was backed up.

        Recovering to the original location will replace any existing database with the same name in the selected instance. Type **YES** to confirm.
    - **New Location** - Select the source and a target Azure SQL instance or Managed Instance to which the Azure SQL database will be recovered. Toggle on **Overwrite existing database(s)** to allow Cohesity to overwrite the existing databases in the target location. Type **YES** to confirm.
5. Select your **Recovery Options**:
    - **Elastic Pool**. To use an elastic pool, check the **Use Existing Elastic Pool** option and enter the pre-existing **Elastic Pool Name** associated with the server. The default value is **No**.

- **Rename**. Add a Prefix and/or Suffix to the full name of the recovered Azure SQL database.
- **Task Name**. Change the default name of the recovery task.

6. Click **Recover**.

You can monitor the status of the recovery on the **Activity** page.

Cohesity starts recovering the selected Azure SQL databases.

Azure SQL Managed Instance

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Microsoft Azure SQL Managed Instance.

Azure Database for Azure SQL Managed Instance is a fully managed relational database service offered by Microsoft Azure. It enables users to deploy, manage, and scale Azure SQL Managed Instance in the cloud without the need to manage the underlying infrastructure.

> **Important:** The backup and recovery of Azure SQL Managed Instance is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).
> This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

This topic covers the following:

Plan and Prepare for Azure SQL Managed Instance Protection

Before protecting Azure SQL Managed Instances with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure SQL Managed Instances with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Authentication for Azure SQL Managed Instance Protection

You must authenticate to Azure SQL Managed instance with Cohesity cluster for cloud in Azure to perform backup and restore. You can use one of the methods such as Username and Password, User Managed Identity, or System Managed Identity.

- Username and Password:

  With this method, you need to provide the credentials of the admin user of the SQL Managed Instance. This is the user account configured in your Azure environment during SQL Managed Instance creation. If needed, instead of using the admin user credentials, you can create a custom user and use it for authentication during backups. However, for restoring that backup, the admin user credentials are required. So, if you have used custom user credentials for backup authentication, you must change them back to the admin user credentials at the time of restore.
  You can use the following commands to create a custom user on the server/Managed Instance:

  - To create a custom user on the server/Managed Instance:
    ```
    #CREATE LOGIN [newlogin] WITH PASSWORD = 'somepassword;
    ```
  - To create same user under database which should be backed up by the custom user you have created:
    ```
    #CREATE USER [newlogin] FOR LOGIN [newlogin];

    #ALTER ROLE db_owner ADD MEMBER [newlogin];
    ```
  - Run the following command on the master system database:
    ```
    #CREATE USER [newlogin] FOR LOGIN [newlogin];
    ```

- User Managed Identity (UMI):

  It uses a user assigned managed identity to authenticate to the Azure SQL Managed Instance. For this authentication, you must assign the Managed Identity Operator role to the UMI identity. Also, this UMI needs to be from the same subscription as the one in which the SQL Managed instance exists.

  A single identity is shared by all nodes in the cluster. This identity should be assigned as the Microsoft Entra admin of the SQL Managed instance. Ensure all cluster nodes have this UMI set as their identity. To use the UMI method for authentication, you must perform the following activities as prerequisite:

  - Create a user-assigned managed identity that would be used for authentication.
    Manage user-assigned managed identities
  - Assign the Managed Identity Operator role to the UMI identity.
    Assign Azure roles to a managed identity
  - Set the UMI as identity of the cluster nodes.
    Assign a user-assigned managed identity
  - Assign this User Managed identity as Entra ID Admin on the server.
    Set Microsoft Entra admin
  - Assign the UMI to Azure SQL Managed Instance/server under the Identity section > User assigned managed identity on the Azure Portal.

- System Managed Identity (SMI):

It uses a unique identity that is automatically created by Azure when you enable SMI on the node. This identity belongs only to that node and is deleted if the node is deleted. Each node in your cluster will have its own SMI, meaning multiple distinct identities. However, Azure SQL Managed Instance allows only one Microsoft Entra admin per node. This creates a challenge when you have multiple nodes, each with a different SMI. To mitigate this issue, you can do the following:

- Create a Microsoft Entra group.
- Add all the nodes' SMI identities as group members.
- Assign this group as Microsoft Entra Admin.

This way, all nodes can authenticate using their respective SMIs through the group. SMI is not created separately as a resource here. You need to enable it on the cluster nodes from the Azure portal.
To use the SMI method for authentication, you must perform the following activities as prerequisite:

- Enable system-assigned managed identity on all the nodes of the cluster. Enable system-assigned managed identity
- Create a Microsoft Entra basic group and add SMI's of all the nodes created in previous step as members. SMI name is the same as the VM (node) name which can be easily searched for while finding member to include in the group. Create a basic group and add members
- Assign this Microsoft Entra basic group as Entra ID Admin on the server. Set Microsoft Entra admin

Considerations for Azure SQL Managed Instance Protection

You must review the following considerations before you proceed for registration, backup, or recovery:

- Backup and recovery are not supported for the following:
  - GRAPH_TABLE,
  - TEMPORAL_TABLE,
  - LEDGER_TABLE,
  - MEMORY_OPTIMIZED_TABLE
- Incremental backups fall back to full backups in the following cases:
  - Schema changes are detected during backup.
  - The full backup state file from an earlier backup is not available.
  - CDC (Change Data Capture) is enabled or disabled on the database/tables.
  - A previous full backup does not exist.

- If a backup fails due to insufficient tempDB storage on the SQL Managed Instance, you need to select the Client-side processing for CDC. Go to **Protection Group** > **Additional Settings** > **CDC processing location** > select **Client**.
- Restoring a backup taken from a higher database version to a lower version is not supported.
- If public access to the SQL Managed Instance is not available and a private endpoint is used for connectivity, add the private endpoint mapping in Cohesity Host Mapping, so name resolution works over the private network. Go to **Settings** > **Networking** > **Host Mapping** and add the required endpoints.
- An Azure VM, which is used as a Cohesity cluster node, should be in the same Vnet as that of the Azure-managed instance. Alternatively, if the Cohesity cluster node and SQL-managed instance are in different Vnet, then both the Vnets must peer to access the database instance.
- Backup fails when Readlock is placed on the database or resource group.
- Database diagrams are not restored.
- Backup and restore are not supported for the databases, which makes it mandatory to use client certificates for their connection to Cohesity cluster.
- Only default ports are supported.
- Database names containing the characters '#' and '/' are not supported for backup and restore operations. Also, the database name should adhere to the naming conventions suggested by the cloud vendors.
- ";" is not supported in server or database passwords.
- Restore across subscriptions of different tenants is not supported.
- Backup failure with invalid object name error: Subsequent full/incremental backups of an Azure SQL Managed Instance backup may fail with the following error, "Invalid object name 'cdc.<schema_name>_<table_name>_CT'".
    - This issue occurs while executing the backup query script with a transfer object in the query (to transfer table from one schema to another).
    - This may cause a mismatch between the user tables and their associated CDC tables. The mismatch typically arises after the schema changes or when database objects are moved between schemas after CDC is enabled.
      Workaround:
        - Disable CDC on the database before running the backup.
        - Run the following command: # EXEC sys.sp_cdc_disable_db
        - Perform a full backup. The CDC will be automatically enabled during the next backup.

- **Duplicate Records**: Recovered database may contain duplicate records if the source database does not include primary or foreign keys."
- **CDC Job Cleanup**: Full backups can fail during CDC processing if CDC jobs are not cleaned at the SQL MI level. This issue occurs when a database is deleted and then recreated with the same name.

  Workaround:

  You can manually perform the cleanup and disable CDC on database before backup. For more details, see Cleanup job.

  > **Note:** The CDC cleanup job removes all data from the change tracking tables based on the configured retention period.

- **CMK as TDE**: For native backups, Cohesity supports only CMK as TDE on the instance. During recovery, ensure that the same key (and version) is associated with the target server's CMK. Since Azure does not allow export of the keys, ensure that the keys are backed up and associated correctly during recovery.
- **Recovery Database Encryption**: During backup and recovery of Azure SQL MI without CDC, if the source database has encryption disabled but Transparent Data Encryption (TDE) is configured on the SQL server with Service-Managed Keys (SMK) or Customer-Managed Keys (CMK), encryption will be turned on for the recovered database by default.
- **Database Tables**: If the databases contain any of the following types of tables, the backup fails due to CDC limitations:
  - Graph tables
  - Temporal tables
  - Ledger tables (updatable ledger)
  - Memory optimized table (business critical tier only)
  - These types of tables are supported for Azure Managed Instance databases, if the backup uses the native backup database workflow, using TDE enabled by Customer-Managed Keys (CMK) or TDE disabled.
- **Column-level Encryption**: Column-level encryption is supported only for Azure SQL Managed instances where TDE is configured with CMK. Full backups must be performed either in the native BAK format (without CDC data).
- **Public Endpoint Support**: Public endpoint on Azure SQL MI workload is not supported.
- **Incremental Backup Failure**: Incremental backup for Azure SQL MI fails when the **tempdb** storage is exhausted due to a large number of CDC transactions on the database. To avoid this, use the batch count to limit the number of transactions being

processed at a time.

- **CDC Performance**: For optimal CDC performance, ensure the following:
  - Do not use DDL triggers that block schema changes, like CREATE/DROP TABLE, CREATE/DROP FUNCTION.
  - Ensure that the SQL Server Agent service is running.
  - Use an account with db_owner or sysadmin privileges for CDC setup.
  - To avoid capture_instance conflicts, check if the table is already CDC-enabled.
- **Protection to New Environment**: DataProtect can back up and recover databases to new environments, but the identity columns may not retain their original seed values. This can cause new records to receive unexpected identity numbers, especially if rows were deleted before the backup. However, this issue does not affect TDE-enabled Azure SQL Managed Instance databases configured with a CMK and backed up with the native BAK format (without CDC data).
- **SqlPackage Utility**: Consider the following for the SqlPackage utility:
  - SqlPackage may fail to export the following SQL Server objects: Logins, Jobs (SQL Agent Jobs), Linked Servers, Server-level permissions, Server roles, Credentials, Database Mail configuration, Replication settings, SSIS packages, Service Broker objects, CLR assemblies, Data in FILESTREAM/FILETABLE, Encrypted objects (if keys/certs are not exported), Cross-database dependencies, extended stored procedures, and some system objects.
  - SqlPackage primarily exports database-level objects and data, but not server-level or agent-level configurations. For a complete migration, use additional tools or scripts to handle these objects.

Backup Azure SQL Managed Instance

After registration of your Azure tenant as a source, Cohesity uses the cloud credentials to authenticate with Azure.

After authentication, Cohesity cluster auto-discovers all the subscriptions (that the service account can access) and supported workloads such as Azure SQL Managed Instance, MySQL, and so on within those subscriptions.

The discovered workloads appear under **Data Protection** > **Sources**, under the **Azure** section. Click your Azure instance and then click **SQL Managed Instance** tab to view all the discovered objects.

Hierarchy View displays the entity hierarchy of the Azure SQL Managed Instance, allowing you to navigate from the tenant-level down to individual databases.

Azure SQL Managed Instance Entity Hierarchy

The Azure SQL Managed Instance entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Managed Instance	Fully SQL Managed Instance hosted within the Azure cloud.
Database	Fully managed relational database service in Microsoft Azure.

Backup Scope by Azure SQL Managed Instances Hierarchy

Cohesity cluster allows you to backup Azure SQL Managed Instance database(s) at different levels of hierarchy.

You can choose any of the following objects to back up:

- **Tenant** – Protect all Azure SQL Managed Instance databases within the tenant.
- **Subscription** – Protect all Azure SQL Managed Instance databases within the subscription.
- **Resource Group** – Protect all Azure SQL Managed Instance databases within the resource group.
- **Database Server** – Protect all Azure SQL Managed Instance databases within the database server.
- **Database** – Protect individual databases.

Supported Backup Types

Full and incremental backups of Azure SQL Managed Instance are supported.

Protect Azure SQL Managed Instance

To back up a registered Azure SQL Managed Instance, you must:

1. Decide whether to use an existing Protection policy or create a new one with the required parameters.
2. Decide which objects to include in the Protection Group.
3. Set credentials for authentication.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy

Prerequisites to Protect Azure SQL Managed Instance

Configure Authentication for Azure SQL Managed Instance

Before performing a backup of Azure SQL Managed Instances that are discovered in Cohesity cluster, you must configure credentials for these Managed Instances. These credentials are essential because they enable secure authentication and connectivity between Cohesity and your Azure SQL Managed Instance. Without valid credentials, the backup process cannot access your Managed Instance or its databases.

You need to enter the credentials that you have already configured for the Managed Instances in your Azure environment. Once these credentials are validated and the connection is established, Cohesity can access the databases within the instance to perform operations such as dumping, protecting and restoring in case required.

Consider the following points:

1. You can set credentials for the SQL Managed Instance at any object in the entity hierarchy.
    1. Tenant
    2. Subscription
    3. Resource group
    4. Instance level
2. The credentials you set are applied to all child objects within the selected object.
3. You cannot set credentials for individual databases.
4. It is recommended to apply credentials at the Managed Instance level. As if you apply credentials at a higher level in the entity hierarchy (Resource group, Subscription, or Tenant), they will be applied to all child-level Managed Instances. If the credentials do not match any of the Managed Instances at the child-level, the backup of the databases within those Managed Instances will fail.
5. Three authentication methods are supported:
    1. Username and Password
    2. User Managed Identity
    3. System Managed Identity

To set the credentials for Azure SQL Managed Instance:

1. Navigate to **Data Protection > Sources**, find the registered Azure source, and click on it.
2. Click on the **SQL Managed Instance** tab.
3. Click the Hierarchy view icon (if you are in a different view) to display the hierarchy. This view makes it easier to view the entity hierarchy and select the required object.
4. Select the checkbox next to object to which you want to set credentials and click **Select All Child Objects**.
5. Click **Database Credentials**.
6. You can select one of the following options:
    1. **Username and Password** - Enter the SQL Managed Instance admin username and password. The account must have access to all databases within the Azure SQL Managed Instance. Or enter the credentials of the custom user created for backup.
    2. **User Managed Identity (UMI)** - Provide the Managed Identity Client ID and Subscription ID. This uses authentication automatically managed by Azure. The Subscription ID of UMI is optional. If not specified, the UMI is assumed to belong to the same subscription as the database server.

        > **Note:**
        > - User Managed Identity must be attached to the Azure VM on which Cohesity cluster is deployed.
        > - User Managed Identity authentication method does not work across tenants.
        > - Authentication fails if the provided database credentials are incorrect. Ensure that the username, password, or UMI client ID is correct, and that the database user associated with the User Managed Identity is properly configured.

    3. **System Managed Identity**: Select the option.

> **Note:**
> - System Managed Identity must be attached to the Azure VM on which Cohesity cluster is deployed.
> - System Managed Identity authentication method does not work across tenants.

7. Click **Save**.

Protect Azure SQL Managed Instance

Before you proceed to protect an Azure SQL Managed Instance, make sure that the authentication is configured. For more information, see the Configure Authentication for Azure SQL Managed Instance topic.

> **Note:** An object cannot be included in two different protection groups for protection.

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** > **Azure** > **SQL Managed Instance**.
3. On the **New Protection** page,
    1. Click **Add Objects**. From the **Registered Source** field, select the Azure instance you have registered.
    2. Do any of the following:
        1. To apply auto-protection,
            1. To apply auto-protection for the objects (Tenants, Subscriptions, Resource Groups, Managed Instances, and Databases), click the auto-protection icon next to it.
            2. All child objects under the selected object will be auto protected.
            3. Note: Auto-protection is not supported for individual database-level. To apply auto-protection for databases, you must select a parent object in the hierarchy.
            4. To exclude any object from auto-protection, click the auto-protection icon next to that object.
        2. To apply manual protection,

1. To apply manual protection, click the checkboxes of the objects (Tenant, Subscription, Resource Group, Managed Instance, and Database) you want to protect.
2. The child objects under the selected object will be selected.
3. You can select or deselect checkboxes of the objects as required.

3. click **Continue**.

4. On the **New Protection** page,
   1. Enter a name for your Protection Group.
   2. Click **Policy** and select the policy to specify backup frequency and retention. If you do not have a policy, you can create one.
   3. Click **Change Data Capture (CDC) for Backups** and toggle on the **Use CDC During Backups** option to use CDC data for database backups. If CDC data is not used, incremental backups will not be supported.
   4. Click **Staging Container** and select either of the following options to create the staging container:
      - **Auto Create** – Cohesity automatically creates a staging container in your storage account.
      - **Specify Container to Use** – Provide the **Staging Container Name**. Cohesity creates a staging container in your storage account with the provided name. If the container already exists, it is reused.

      > **Note:** A storage account must already be created with the Cohesity pre-defined naming convention: first 16 characters of SHA256 hash of concatenated string of the Azure Subscription ID and Region prefixed with "cohe". For example, *cohe2e5fb564552100a6*.

      This option is available only if **Change Data Capture (CDC) for Backups** is disabled.
5. Optionally, to configure any of the advanced settings, select **More Options**.
6. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Storage Domain**	Select the appropriate Storage Domain. For more information, see Manage Storage Domains
	Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**CDC Data Processing Location**	Select the location where CDC data processing is performed: • **SQL MI Server** - CDC data processing is performed on the server, using temporary tables. • **Cohesity Cluster** - CDC data processing is performed on the client (Cohesity cluster nodes). > **Note:** This option is displayed only if **Change Data Capture (CDC) for Backups** is enabled.
	CDC Batch Count	Enter the number of CDC records processed in batches to prevent consuming the entire temporary database. Set this value based on the server capacity, specific to each instance. The default value is **-1** indicating unlimited (no batching). > **Note:** This option is displayed only if **Change Data Capture (CDC) for Backups** is enabled, and **CDC Data Processing Location** is set as **SQL MI Server**.
	Pause Future Runs	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
	End Date	If you need to end protection on a specific date, enable this option to select the date.

Section	Setting	Description
	QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
	Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the protection takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.
	Priority	Select a priority for the Object Protection Group Execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

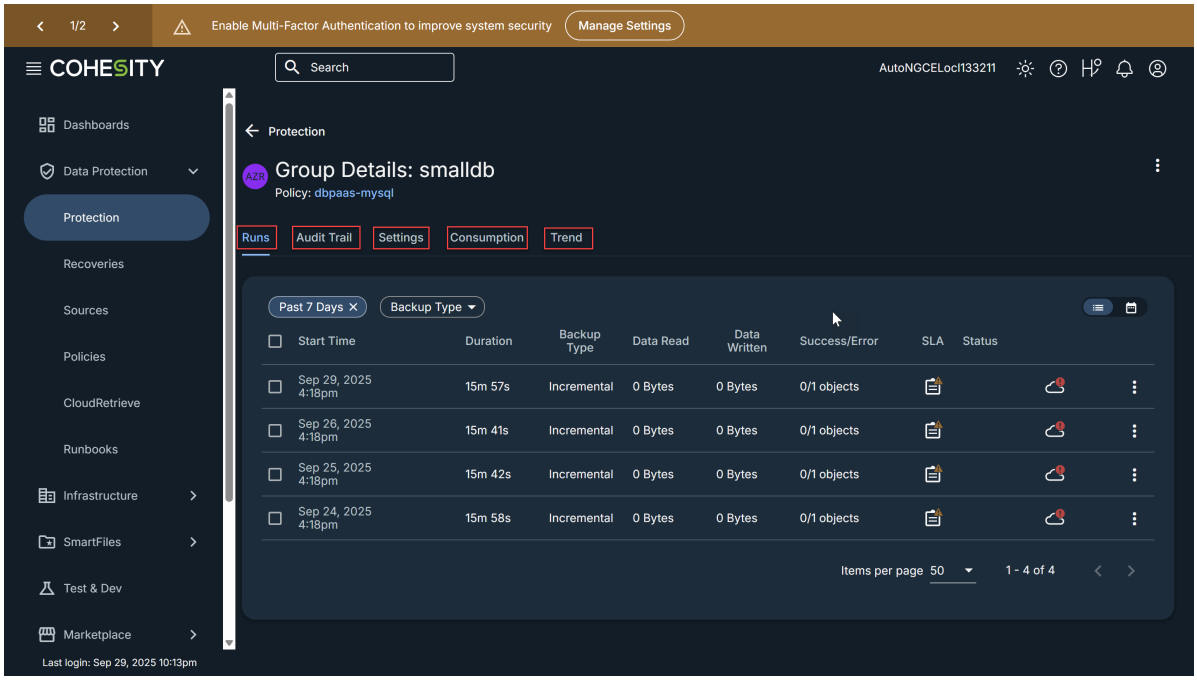

Manage Azure SQL Managed Instance Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure SQL Managed Instance.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.

4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure SQL Managed Instance.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Azure SQL Managed Instance.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.

3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Cloudspin artifacts generated by the Protection Group must be manually deleted from the respective cloud portals. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for Azure SQL Managed Instance.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.

- Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.

5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Azure SQL Managed Instance Databases

After you protect your Azure SQL Managed Instance, you can recover them from Cohesity cluster to their original location or a new location.

> **Note:**
> - If you plan to overwrite an existing database, Cohesity recommends backing up the database first.
> - To restore to a new location, the destination instance credential validation must be done, and the host mapping of the destination instance must be added in **Settings**> **Networking**> **Host Mapping**.

To recover protected Azure SQL Managed Instance databases:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** > **Azure** > **SQL Managed Instance**.
3. On the **New Recovery** page,
    1. Use the filters and search box to locate the required databases. Wildcards are supported.
    2. Select the required database that is to be recovered.
    3. Click **Next: Recovery Option**.
    4. In the **Recover To** section, do one of the following:
        1. To restore to the original location:
            1. Select the **Original Location** option to recover the database to its original location. Consider the following points:
                1. The existing database at that location with the same name will always be overwritten.
                2. You must type **Yes** to confirm and proceed.
                3. To avoid overwriting the database with the same name, enter a new name in the **Rename** field below. The database will be recovered with the new name you have provided.
        2. To restore to a new location:

1. Select the **New Location** option to recover the database to an alternative SQL Managed Instance.
2. Select the registered Azure tenant from the **Source** drop-down and select the Managed Instance to which you want to recover the database.
3. To overwrite an existing database, toggle on the **Overwrite existing database(s)** option and type **Yes** to confirm.
4. If you want to recover the database with a different name, enter the new name in the **Rename** field below.
5. If you enter a new name, the database will be restored at the selected location with the new assigned name.
6. If you do not enter a new name, the existing database with the same name will be overwritten.

5. Select the **Staging Container** by choosing either of the following options to create the staging container:
   - **Auto Create** – Cohesity automatically creates a staging container in your storage account.
   - **Specify Container to Use** – Provide the **Staging Container Name**. Cohesity creates a staging container in your storage account with the provided name. If the container already exists, it is reused.

     > **Note:** A storage account must already be created with the Cohesity pre-defined naming convention: first 16 characters of SHA256 hash of concatenated string of the Azure Subscription ID and Region prefixed with "cohe". For example, *cohe2e5fb564552100a6*.

   This option is applicable only if CDC data is not used for the corresponding backup job. Otherwise, this setting is ignored.

4. Select your **Recovery Options**:
   1. **Rename**. If you have renamed the database, it will be recovered with the new name at either the original or the new location, based on the option you have selected.
   2. **Task Name**. Change the default name of the recovery task.
5. Click **Recover**.

Azure Table Storage

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Microsoft Azure Table Storage. Table Storage is a fully managed relational database service offered by Microsoft Azure. It enables users to deploy, manage, and scale Azure Table Storage in the cloud without the need to manage the underlying infrastructure.

> **Important:** The backup and recovery of Azure Table Storage is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

This topic covers the following:

Plan and Prepare for Azure Table Storage Protection

Before protecting Azure Table Storage with the Cohesity cluster, ensure you have reviewed the considerations described in this topic.

Before You Begin

To protect Azure Table Storage with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Considerations for Azure Table Storage Protection

You must review the following considerations before you proceed for registration, backup, or recovery:

- Restoring a backup taken from a higher database version to a lower version is not supported.
- Only up to 5000 tables can be protected in a single Azure Table Storage protection run.

Backup Table Storage

After registration of your Azure tenant as a source, Cohesity uses the cloud credentials to authenticate with Azure.

After authentication, Cohesity cluster auto-discovers all the subscriptions (that the service account can access) and supported workloads such as Table Storage, MySQL, and so on within those subscriptions.

The discovered workloads appear under **Data Protection** > **Sources**, under the **Azure** section. Click your Azure instance and then click **Table Storage** tab to view all the discovered objects.

Hierarchy View displays the entity hierarchy of the Table Storage, allowing you to navigate from the tenant-level down to individual databases.

Azure Table Storage Entity Hierarchy

The Azure Table Storage entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Storage Accounts	Fully managed Table Storage account hosted within the Azure cloud.
Storage Account Tables	Fully managed relational database service in Microsoft Azure.

Backup Scope by Table Storage Hierarchy

Cohesity cluster allows you to backup Table Storage database(s) at different levels of hierarchy.

You can choose any of the following objects to back up:

- **Tenant** – Protect all Azure Table Storages within the tenant.
- **Subscription** – Protect all Table Storages within the subscription.
- **Resource Group** – Protect all Azure Table Storages within the resource group.
- **Storage Account** – Protect all Table Storage Tables within the database storage accounts.
- **Database** – Protect individual Tables.

Supported Backup Types

Only full backups are supported. Incremental backups are not supported.

Protect Azure Table Storage

To back up a registered Azure Table Storage, you must:

1. Decide whether to use an existing Protection policy or create a new one with the required parameters.
2. Decide which objects to include in the Protection Group.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Protect Azure Table Storage

Before you proceed to protect an Azure Table Storage, make sure that the authentication is configured. For more information, see Plan and Prepare for Azure Table Storage Protection.

> **Note:** An object cannot be included in two different protection groups for protection.

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** > **Azure** > **Table Storage**.
3. On the **New Protection** page,
    1. Click **Add Objects**. From the **Registered Source** field, select the Azure instance you have registered.
    2. Do any of the following:
        1. To apply auto-protection,
            1. To apply auto-protection for the objects (Tenants, Subscriptions, Resource Groups, Managed Instances, and Databases), click the auto-protection icon next to it.
            2. All child objects under the selected object will be auto protected.
            3. Note: Auto-protection is not supported for individual database-level. To apply auto-protection for databases, you must select a parent object in the hierarchy.
            4. To exclude any object from auto-protection, click the auto-

protection icon next to that object.

2. To apply manual protection,
    1. To apply manual protection, click the checkboxes of the objects (Tenant, Subscription, Resource Group, Managed Instance, and Database) you want to protect.
    2. The child objects under the selected object will be selected.
    3. You can select or deselect checkboxes of the objects as required.
3. click **Continue**.

4. On the **New Protection** page,
    1. Enter a name for your Protection Group.
    2. Click **Policy** and select the policy to specify backup frequency and retention. If you do not have a policy, you can create one.
5. Optionally, to configure any of the advanced settings, select **More Options**.
6. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Storage Domain**	Select the appropriate Storage Domain. For more information, see Manage Storage Domains
	Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
	End Date	If you need to end protection on a specific date, enable this option to select the date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
	Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the protection takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.

Section	Setting	Description
	Priority	Select a priority for the Object Protection Group Execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

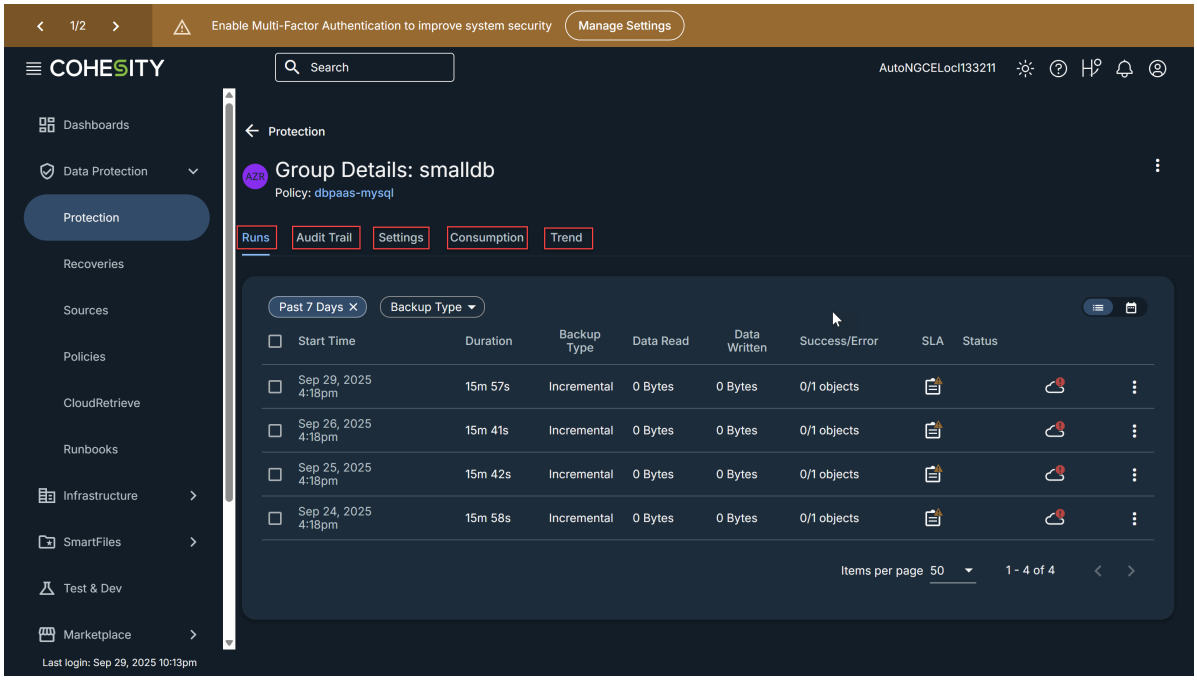

Manage Table Storage Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Table Storage.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Table Storage.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Table Storage.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Cloudspin artifacts generated by the Protection Group must be manually deleted from the respective cloud portals. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for SQL Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Azure Table Storage

After you protect your Azure Table Storage, you can recover them from Cohesity cluster to their original location or a new location.

> **Note:** If you plan to overwrite an existing database, Cohesity recommends backing up the database first.

To recover protected Azure Table Storage databases:

1. Navigate to **Data Protection** > **Recoveries**.

   Select the required database that is to be recovered.
2. Click **Recover** > **Azure** > **Table Storage**.
3. On the **New Recovery** page,
   1. Use the filters and search box to locate the required databases. Wildcards are supported.
   2. Select the required database that is to be recovered.
   3. Click **Next: Recovery Option**.
   4. In the **Recover To** section, do one of the following:

1. To restore to the original location:
    1. Select the **Original Location** option to recover the database to its original location. Consider the following points:
        1. The existing database at that location with the same name will always be overwritten.
        2. You must type **Yes** to confirm and proceed.
        3. To avoid overwriting the database with the same name, enter a new name in the **Rename** field below. The database will then be recovered with the new name you have provided.
2. To restore to a new location:
    1. Select the **New Location** option to recover the database to an alternative Table Storage Account.
    2. Select the registered Azure tenant from the **Source** drop-down and then select the **Storage Account** to which you want to recover the Storage Account Tables.
    3. To overwrite an existing Storage Account Tables, toggle on the **Overwrite existing database(s)** option and type **Yes** to confirm.
    4. If you want to recover the database with a different name, enter the new name in the **Rename** field below.
    5. If you enter a new name, the database will be restored at the selected location with the new assigned name.
    6. If you do not enter a new name, the existing database with the same name will be overwritten.
4. Select your **Recovery Options**:
    1. **Rename**. If you have renamed the database, it will be recovered with the new name at either the original or the new location, based on the option you have selected.
    2. **Task Name**. Change the default name of the recovery task.
5. Click **Recover**.

Azure PostgreSQL

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Microsoft Azure PostgreSQL databases.

Azure Database for Azure PostgreSQL is a fully managed relational database service offered by Microsoft Azure. It enables users to deploy, manage, and scale Azure PostgreSQL databases in the cloud without the need to manage the underlying infrastructure.

> **Important:** The backup and recovery of Azure PostgreSQL is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).

This topic covers the following:

Plan and Prepare for Azure PostgreSQL Protection

Before protecting Azure PostgreSQL with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure PostgreSQL with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Authentication for Azure PostgreSQL Protection

You must authenticate to Azure PostgreSQL database from the Cohesity cluster using one of the following methods:

- Username and Password
- User Managed Identity (UMI)
- System Managed Identity (SMI)

Username and Password

This method uses the credentials of the admin user for the PostgreSQL database. This is the user account configured during server creation.

For backup, you can create a custom user. In the CLI run the following commands:

1. To create login user on the server:

   ```
   #CREATE LOGIN [newlogin] WITH PASSWORD = 'somepassword;
   ```

2. To create same user under database which should be backed up by user **newlogin**:

   ```
   User <database>;
   #CREATE USER [newlogin] FOR LOGIN [newlogin];
   ALTER ROLE db_owner ADD MEMBER [newlogin];
   ```

3. Run the following command on the master system database:

   ```
   CREATE USER [newlogin] FOR LOGIN [newlogin];
   ```

For recovery, the default admin credentials created during server creation are required if the custom user credentials were set for authentication before backup. During recovery, you must change the credentials back to the default admin credentials.

User Managed Identity (UMI)

This method uses a user-assigned managed identity to authenticate to the Azure PostgreSQL database.

Perform the following to use the UMI authentication:

1. Create a user-assigned managed identity. For more information, see Manage user-assigned managed identities using the Azure portal.
2. Assign the **Managed Identity Operator** role to the UMI identity. For more information, see User-assigned managed identity.
3. Set the UMI as the identity for all the cluster nodes. Ensure that the UMI is from the same subscription as the one in which the PostgreSQL database exists. For more information, see Assign a user-assigned managed identity to an existing VM.
4. Assign this User Managed Identity as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.
5. Assign the required user managed identity to Azure PostgreSQLdatabase under the Identity section - User assigned managed identity on the Azure Portal.

System Managed Identity (SMI)

This method uses a unique identity that is automatically created by Azure when you enable SMI on the node. This identity belongs only to the specific node and is deleted if the node is deleted.

Each node in the cluster includes its own SMI (multiple distinct identities). Since the Azure PostgreSQL database allows only one Microsoft Entra admin per node, having multiple nodes, each with a different SMI may cause issues.

To mitigate this issue, you can perform the following:

1. Create a Microsoft Entra group.
2. Add all the nodes' SMI identities as group members.
3. Assign this group as Microsoft Entra Admin.

This ensures that all the nodes can be authenticated using their respective SMIs through the group.

> **Note:** SMIs are not created separately as a resource. You must enable it on the cluster nodes from the Azure portal.

Perform the following to use the SMI authentication:

1. Enable system-assigned managed identity on all the nodes of the cluster. For more information, see Enable system-assigned managed identity on an existing VM.

2. Create a Microsoft Entra basic group and add SMI's of all the nodes created in previous step as members. The SMI name is the same as the VM (node) name, which can be easily searched for while finding members to include in the group. For more information, see Create a basic group and add members.
3. Assign this Microsoft Entra basic group as Entra ID Admin on the server. For more information, see Set Microsoft Entra admin.

Considerations for Azure PostgreSQL Protection

You must review the following considerations before you proceed for registration, backup, or recovery:

- Cohesity cluster does not restore accounts and their grants to the destination server.
- Restoring a backup taken from a higher PostgreSQL database version to a lower version is not supported.
- Restoring security privileges is not supported.
- Use the --no-owner and --no-privileges options during restore.
  After restore, metadata captured at the time of backup is shown as the owner/ACL in the progress log restore activity on the web UI.
- Restore does not fail if the owner or role does not exist on the destination.
- Modify database ownership after restore.
- Azure PostgreSQL database restore from a single server to a flexible server or vice versa is not supported.
- Unsupported characters in database names during restore: `, @, \, [,], !, #, %, ^, ., ,, &, *, (,), <, >, ?, /, |, }, {, ~, :, ', ", ;, +, =, -
- Uppercase names are not supported for new accounts added after PostgreSQL server creation.
- If full backups fail due to temporary objects, delete the temporary objects manually and run the backup again.
- When Azure Managed Identity (UMI or SMI) is configured for authentication, Cohesity cluster performs full backups only. Incremental backups are not supported with MI. Use Username and Password authentication to enable incremental backups.

Considerations for Azure PostgreSQL Incremental Backups

- You can protect Azure PostgreSQL sever assets with a protection plan and policy. You can use full and incremental schedules.
- Always set the **wal_level** server parameter as logical.
- The backup user must be an administrator.
- The backup user must have the CREATEROLE and REPLICATION permissions.
- To backup tables of other users, the tables must have the primary key.

- Intelligent groups are supported.
- Replica server databases are not supported for full and incremental backups.
- Large object data type is not supported for incremental backups.
- Incremental backup with the **bytea_output** server parameter with an escape value is not supported.
- Azure SMI and UMI are not supported for incremental backups.
- The last value of the sequence generator may not be consistent after recovery of an incremental backup.
- To rename a database, perform the following:
  - Note the name of the replication slot created for the database.

    ```
    SELECT slot_name FROM pg_replication_slots WHERE database =
    <database name to be renamed>
    ```

  - Drop the replication slot after renaming the database.

    ```
    SELECT pg_drop_replication_slot('<replication slot name>')
    ```

- The default value of the **max_replication_slots** server parameter is **10**. Since each database uses one replication slot, ensure you set this value to a higher number based on the approximate number of databases to be protected.
- Set the **azure.enable_temp_tablespaces_on_local_ssd** server parameter to **OFF** when the following error occurs:

  "psycopg2.errors.DiskFull: could not write to file "pg_tblspc/16386/PG_16_202307071/pgsql_tmp/pgsql_tmp449554.119": No space left on device."
- Ensure to check the storage space while incremental backups are running.
- Avoid maintenance windows during backup and recovery.

Considerations for Azure PostgreSQL Encryption

- PostgreSQL does not natively provide Transparent Data Encryption (TDE) in standard community deployments. However, PostgreSQL supports application-level and column-level encryption through cryptographic extensions and functions (for example, pgcrypto).
- In Azure PostgreSQL, encryption at rest is typically provided through Azure-managed storage encryption mechanisms.
- During protection, the encrypted column data is preserved in its encrypted form. After recovery, applications can continue to access and decrypt the data using the appropriate encryption keys and application logic.
- Backup and recovery operations do not modify or decrypt application-encrypted column data. You must ensure that the required encryption keys remain accessible after recovery.

Backup Azure PostgreSQL

After registration of your Azure tenant as a source, Cohesity uses the cloud credentials to authenticate with Azure.

After authentication, Cohesity cluster auto-discovers all the subscriptions (that the service account can access) and supported workloads such as PostgreSQL, Azure SQL Managed Instance, MySQL, and so on within those subscriptions.

The discovered workloads appear under **Data Protection** > **Sources**, under the **Azure** section. Click your Azure instance and then click **PostgreSQL** tab to view all the discovered objects.

Hierarchy View displays the entity hierarchy of the Azure PostgreSQL, allowing you to navigate from the tenant-level down to individual databases.

Azure PostgreSQL Entity Hierarchy

The Azure PostgreSQL entity hierarchy follows this structure:

Structure	Description
Tenant	Dedicated instance of the identity service.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and web apps.
Database Servers	Fully managed Azure PostgreSQL server instance hosted within the Azure cloud.
Database	Fully managed relational database service in Microsoft Azure.

Backup Scope by Azure PostgreSQL Hierarchy

Cohesity cluster allows you to backup Azure PostgreSQL database(s) at different levels of hierarchy.

You can choose any of the following objects to back up:

- **Tenant** – Protect all Azure PostgreSQL databases within the tenant.
- **Subscription** – Protect all Azure PostgreSQL databases within the subscription.
- **Resource Group** – Protect all Azure PostgreSQL databases within the resource group.

- **Database Server** – Protect allAzure PostgreSQL databases within the database server.
- **Database** – Protect individual databases.

Supported Backup Types

Cohesity supports full backups and incremental backups for Azure PostgreSQL.

Protect Azure PostgreSQL

To back up a registered Azure PostgreSQL, you must:

1. Decide whether to use an existing Protection policy or create a new one with the required parameters.
2. Decide which objects to include in the Protection Group.
3. Set credentials for authentication.

A Protection Group enables you to select and group the objects that you want to backup. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy

Prerequisites to Protect Azure PostgreSQL

Configure Authentication for Azure PostgreSQL

Before performing a backup of Azure PostgreSQLs that are discovered in Cohesity cluster, you must configure credentials for these Managed Instances. These credentials are essential because they enable secure authentication and connectivity between Cohesity and your Azure PostgreSQL. Without valid credentials, the backup process cannot access your Managed Instance or its databases.

You need to enter the credentials that you have already configured for the Managed Instances in your Azure environment. Once these credentials are validated and the connection is established, Cohesity can access the databases within the instance to perform operations such as dumping, protecting and restoring in case required.

Consider the following points:

1. You can set credentials for the PostgreSQL at any object in the entity hierarchy.
    1. Tenant
    2. Subscription

3. Resource group
4. Instance level

2. The credentials you set are applied to all child objects within the selected object.
3. You cannot set credentials for individual databases.
4. It is recommended to apply credentials at the Managed Instance level. As if you apply credentials at a higher level in the entity hierarchy (Resource group, Subscription, or Tenant), they will be applied to all child-level Managed Instances. If the credentials do not match any of the Managed Instances at the child-level, the backup of the databases within those Managed Instances will fail.

To set the credentials for Azure PostgreSQL:

1. Navigate to **Data Protection > Sources**, find the registered Azure source, and click on it.
2. Click on the **PostgreSQL** tab.
3. Click the Hierarchy view icon (if you are in a different view) to display the hierarchy. This view makes it easier to view the entity hierarchy and select the required object.
4. Select the checkbox next to object to which you want to set credentials and click **Select All Child Objects**.
5. Click **Database Credentials**.
6. You can select one of the following options:
    1. **Username and Password** - Enter the PostgreSQL admin username and password. The account must have access to all databases within the Azure PostgreSQL. Or enter the credentials of the custom user created for backup.
    2. **Managed Identity** - Authenticate using a User Assigned (UMI) or System Assigned (SMI) Managed Identity configured on the Azure source. Only full backup is supported with Managed Identity; incremental backup fails with Managed Identity.
7. Click **Save**.

Protect Azure PostgreSQL

Before you proceed to protect an Azure PostgreSQL, make sure that the authentication is configured.

> **Note:** An object cannot be included in two different protection groups for protection.

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** > **Azure** > **PostgreSQL**.
3. On the **New Protection** page,
    1. Click **Add Objects**. From the **Registered Source** field, select the Azure instance you have registered.
    2. Do any of the following:
        1. To apply auto-protection,
            1. To apply auto-protection for the objects (Tenants, Subscriptions, Resource Groups, Servers, and Databases), click the auto-protection icon next to it.
            2. All child objects under the selected object will be auto protected.
            3. Note: Auto-protection is not supported for individual database-level. To apply auto-protection for databases, you must select a parent object in the hierarchy.
            4. To exclude any object from auto-protection, click the auto-protection icon next to that object.
        2. To apply manual protection,
            1. To apply manual protection, click the checkboxes of the objects (Tenant, Subscription, Resource Group, Servers, and Database) you want to protect.
            2. The child objects under the selected object will be selected.
            3. You can select or deselect checkboxes of the objects as required.
    3. Click **Continue**.
4. On the **New Protection** page,
    1. Enter a name for your Protection Group.
    2. Click **Policy** and select the policy to specify backup frequency and retention. If you do not have a policy, you can Create or Edit a Standard Policy.
5. Optionally, to configure any of the advanced settings, select **More Options**.
6. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Storage Domain**	Select the appropriate Storage Domain. For more information, see Manage Storage Domains
	Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group to run. An incremental backup captures only data changes since the last successful backup.

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Toggle on this option to stop protection runs from executing. Once you enable this option, no protection runs will be scheduled.
	End Date	If you need to end protection on a specific date, enable this option to select the date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. Cohesity recommends specifying Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Object. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Object. Only specify this policy if you need fast ingest speed for a small number of Objects. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the current running task.
	Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when an Object completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Object fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time specified in the SLA field. Emails are sent when warning alerts are created.

Section	Setting	Description
	Priority	Select a priority for the Object Protection Group Execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	Enter a brief description of the object.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

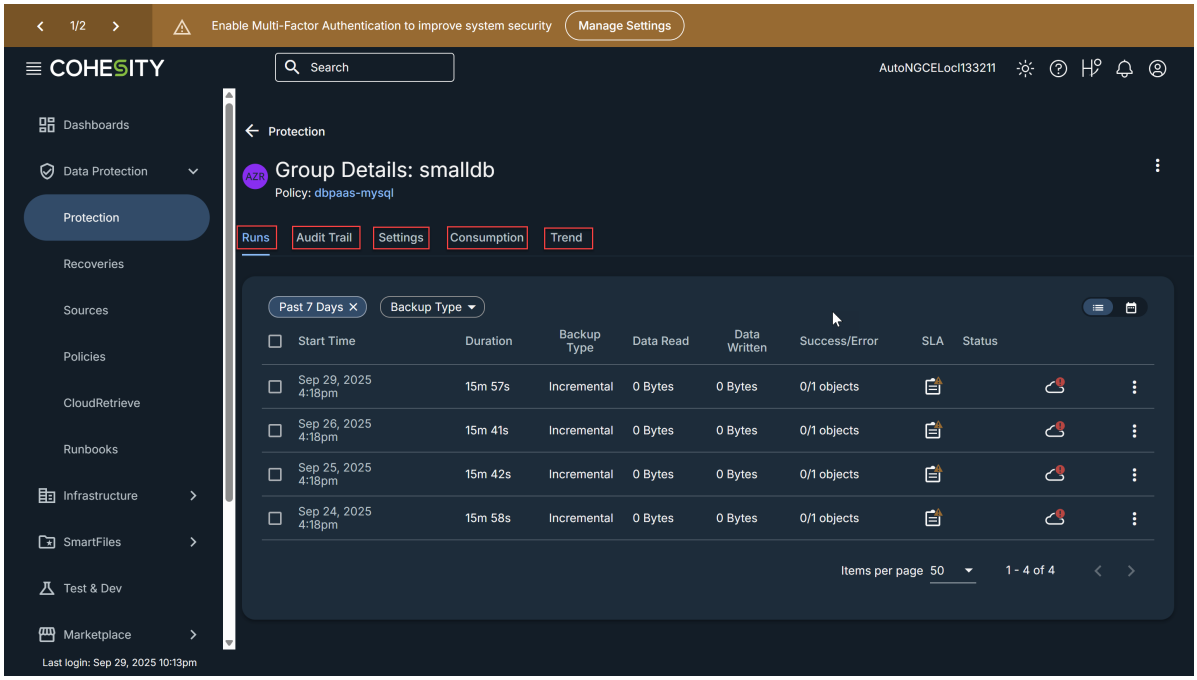

Manage PostgreSQL Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to the Protection Group

You can add new objects to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add objects to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Incremental Backup

You can instantly start a protection run of the Protection Group that you have created to back up Azure PostgreSQL. This allows you to back up selected objects in the Protection

Group.

Perform the following steps to start an on-demand protection run:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, click the actions menu (⋮) next to the required Protection Group and click **Run Now**.
3. In the **Run Now** window, click on **Backup** and perform the following:
    1. From the **Backup** drop-down list, select any one of the following:
        1. All objects in the Protection Group.
        2. All objects that succeeded last run.
        3. Selected objects in the protection run.
    2. In the **Backup Type** drop-down list, select **Incremental**.
4. Optionally, you can click the + icon and perform replication, archival, and vault for the protection group.
5. Click **Run**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Cloudspin artifacts generated by the Protection Group must be manually deleted from the respective cloud portals. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for SQL Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
    - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Azure PostgreSQL Databases

After you protect your Azure PostgreSQL, you can recover them from Cohesity cluster to their original location or a new location.

> **Note:**
> - If you plan to overwrite an existing database, Cohesity recommends backing up the database first.
> - To restore to a new location, the destination instance credential validation must be done, and the host mapping of the destination instance must be added in **Settings**>**Networking**> **Host Mapping**.

To recover protected Azure PostgreSQL databases:

1. Navigate to **Data Protection** > **Recoveries**.

   Select the required database that is to be recovered.
2. Click **Recover** > **Azure** > **PostgreSQL**.
3. On the **New Recovery** page,
    1. Use the filters and search box to locate the required databases. Wildcards are supported.
    2. Select the required database that is to be recovered.
    3. Click **Next: Recovery Option**.
    4. In the **Recover To** section, do one of the following:
        1. To restore to the original location:
            1. Select the **Original Location** option to recover the database to its original location. Consider the following points:
                1. The existing database at that location with the same name will always be overwritten.
                2. You must type **Yes** to confirm and proceed.
                3. To avoid overwriting the database with the same name, enter a new name in the **Rename** field below. The database will be recovered with the new name you have provided.
        2. To restore to a new location:
            1. Select the **New Location** option to recover the database to an alternative PostgreSQL.

2. Select the registered Azure tenant from the **Source** drop-down and then select the PostgreSQL server to which you want to recover the database.
3. To overwrite an existing database, toggle on the **Overwrite existing database(s)** option and type **Yes** to confirm.
4. If you want to recover the database with a different name, enter a new name in the **Rename** field below.
5. If you enter a new name, the database will be restored at the selected location with the new name assigned.
6. If you do not enter a new name, the existing database with the same name will be overwritten.

4. Select your **Recovery Options**:
    1. **Refresh Materialized Views**. Toggle On Refresh materialized views automatically after recovery to automatically refresh the PostgreSQL materialized views post recovery to ensure they reflect the latest recovered data.

       Materialized views in PostgreSQL are database objects that physically store the pre-computed results of a query on the disk. Cohesity automatically runs a refresh script after recovery to refresh materialized views and ensure data consistency.

       If this option is toggled Off, the recovery time can be reduced, but materialized views can remain outdated until manually refreshed. To manually refresh, download the script from the recovery task details page, and then run it on the recovered PostgreSQL database.
    2. **Rename**. If you have renamed the database, it will be recovered with the new name at either the original or the new location, based on the option you have selected.
    3. **Task Name**. Change the default name of the recovery task.
5. Click **Recover**.

Refresh Azure PostgreSQL Materialized Views

When recovering Azure PostgreSQL databases, materialized views must be refreshed to reflect the recovered data. This topic explains the available options to refresh the materialized views and recommendations to choose the right approach for your environment.

Materialized views in PostgreSQL are database objects that physically store the pre-computed results of a query on the disk. Unlike regular views, they need to be explicitly refreshed to reflect changes in their underlying base tables.

After a database recovery operation, materialized views contain stale data from the backup point. The views must be rebuilt to work with the newly recovered data. Without refreshing, queries against materialized views return outdated or incorrect results.

You can refresh the materialized views using one of the following options:

- Automatic Refresh Materialized Views During Recovery
- Manual Refresh Materialized Views After Recover

The following is a comparison between the refresh options:

Factor	Automatic Refresh	Manual Refresh
Database Size	Small to medium	Large
Refresh Duration	Size	Large
Recovery Time Objective	Flexible	Strict
Permissions Available	Yes	Limited
Scheduling Flexibility	Low	High
Administrative Control	Minimal	Greater

Automatic Refresh Materialized Views During Recovery

Cohesity supports performing automatic materialized view refresh using the recovery options in the Cohesity UI.

You can use this option in the following scenarios:

- Smaller databases where refresh completes quickly
- Immediate data consistency is critical
- Database credentials have the required permissions to refresh materialized views
- Longer recovery duration is acceptable

Prerequisites

Before enabling automatic refresh, verify that the Azure PostgreSQL database user has the required permissions using the following commands:

- Check schema usage:

```
SELECT has_schema_privilege('refresh_user', 'schema_name', 'USAGE');
```

- Check table select permissions:

```
SELECT has_table_privilege('refresh_user', 'table_name', 'SELECT');
```

- Check materialized view permissions:

```
SELECT has_table_privilege('refresh_user', 'materialized_view_name',
'INSERT, UPDATE, DELETE');
```

Procedure

To automatically refresh materialized views:

1. During the recovery workflow, navigate to the **Refresh Materialized Views** option in the **Recovery Options** section.
2. Toggle On **Refresh materialized views automatically after recovery**.
3. Complete the database recovery.

Manual Refresh Materialized Views After Recovery

Cohesity supports manually performing materialized view refresh using the script available in the Cohesity UI.

You can use this option in the following scenarios:

- Large databases with materialized views require significant time to refresh
- DataProtect source credentials do not have sufficient permissions to refresh materialized views
- Recovery Time Objective (RTO) requires faster database availability
- Materialized view refresh must be scheduled during off-peak hours

The following are the benefits of manual post-recovery refresh:

- **Faster Recovery Time**: Your database becomes available immediately after recovery.
- **Controlled Timing**: Schedule the refresh during off-peak hours.
- **Better Resource Management**: Avoid resource contention during critical recovery operations.
- **Flexibility**: Database administrators can monitor and manage the refresh process independently.
- **Permission Flexibility**: Use a different user with appropriate permissions for refresh operations.

Prerequisites

Before refreshing the materialized view manually, verify that the Azure PostgreSQL database user has the required permissions using the following commands:

- Check schema usage:

```
SELECT has_schema_privilege('refresh_user', 'schema_name', 'USAGE');
```

- Check table select permissions:

```
SELECT has_table_privilege('refresh_user', 'table_name', 'SELECT');
```

- Check materialized view permissions:

```
SELECT has_table_privilege('refresh_user', 'materialized_view_name',
'INSERT, UPDATE, DELETE');
```

Procedure

To manually refresh materialized views after recovery:

1. Disable Automatic Refresh:
    1. During recovery, disable the Refresh Materialized Views option.
    2. Complete the database recovery.
    3. Download the refresh script from the Cohesity UI.
    4. Run the script manually after the recovery.
2. Download the refresh script:
    1. Navigate to Recoveries.
    2. Select the PostgreSQL recovery task.
    3. Open the Options tab.
    4. Locate the Materialized View Refresh section.
3. Download the refresh script.
    1. Run the Refresh Script:
    2. Execute the downloaded script on the recovered database
    3. The script automatically determines materialized view dependencies.
    4. Materialized views are refreshed in dependency order.

Best Practices

- **Assess your data size.** For databases with materialized views exceeding several GB, consider manual refresh.
- **Verify permissions.** Before using automatic refresh, ensure dataprotect source credentials have required permissions.
- **Save the script.** Download and save the refresh script from the recovery details page for future use.

- **Test in non-production.** Validate your chosen approach in test/dev environments first.
- **Schedule appropriately.** For manual refresh, run during maintenance windows or low-traffic periods.
- **Monitor resources.** Materialized views refresh can be resource-intensive; ensure adequate system resources are available.
- **Regular maintenance.** Consider implementing periodic refresh schedules (using tools like pg_cron) independent of recovery operations.
- **Document your views.** Maintain documentation of materialized views and their dependencies.

Azure Blob Storage and Azure Data Lake Storage

Azure Blob Storage and Azure Data Lake Storage are types of storage accounts in Microsoft Azure, designed for storing unstructured data. Within these storage accounts, you can create containers (for Blob Storage) or file systems (for Data Lake Storage) to organize and manage your data. Azure Data Lake Storage (ADLS) offers a hierarchical namespace, enabling efficient file system operations and enhanced file-level security, while Azure Blob Storage utilizes a flat namespace for simpler object storage. Cohesity seamlessly integrates with Azure, providing a straightforward, fast, and cost-effective solution for backup, recovery, and data management across both Azure Blob Storage and Azure Data Lake Storage accounts.

> **Important:** The backup and recovery of Azure Blob Storage and Azure Data Lake Storage is supported only on Cohesity Cloud Edition on Azure (formerly known as Azure NGCE cluster).

Plan and Prepare for Azure Blob and Data Lake Storage Protection

Before protecting Azure Blob and Data Lake Storage with the Cohesity cluster, ensure you have met the prerequisites and reviewed the considerations described in this topic.

Before You Begin

To protect Azure Blob and Data Lake Storage with the Cohesity cluster, you must register your Azure as a source in the Cohesity cluster and perform backup and recovery operations. For more information, see Register an Azure Source at Tenant Level.

Prerequisites for Registering Azure Source

Before registering Azure with the Cohesity cluster, you must perform the following steps:

1. Register an application with Microsoft Entra ID and create a service principal. For information, see the Azure documentation.
2. Create an application secret key for setting up authentication for the service principal. For information, see the Azure documentation.
3. Create a custom role at subscription level with the required permissions for backup and recovery.
4. For information about creating a custom role, see the Azure documentation.
5. Assign the custom role to the Microsoft Entra ID application created in step 1.

The Entra ID and application secret key are required when you register the Azure source with the Cohesity cluster.

Considerations for Azure Storage Protection

Review and understand the following considerations before you protect Azure Blob Storage and Data Lake Storage:

- The Microsoft Entra ID you create must have access to the Storage Container you want to protect.
- Protection of Page blobs and blobs in the archive tier is not supported.
- Both full and incremental backups are supported.
- You cannot recover Blob data to a Data Lake container, or Data Lake data to a Blob container

Backup Azure Blob and Data Lake Storage

After registering your Azure source, Cohesity uses the cloud credentials provided to authenticate with Azure.

Once authenticated, Cohesity automatically discovers all the subscriptions that the service account can access, as well as all supported workloads such as Azure Blob and Data Lake Storage within those subscriptions. The discovered workloads appear under **Data Protection** > **Sources**. You can select the **Blob and Data Lake Storage** tab to view all the Storage Containers discovered.

The Hierarchy view displays the entity hierarchy of Azure, allowing you to navigate from the tenant-level down to individual Storage Container.

Azure Storage Container Entity Hierarchy

The **Blob and DataLake Storage** entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
Tenant	The top-level Azure Active Directory boundary that contains subscriptions.
Subscription	Logical container for managing billing, resources, and permissions for services.
Resource Group	Logical container for organizing and managing related Azure resources, such as virtual machines, storage accounts, and storage services.
Storage Account	The Azure Blob Storage or Azure Data Lake Storage having storage container to protect.
Storage Container	The namespace within Azure Blob Storage or Data Lake Storage. Containers store data objects such as blobs.

Backup Scope by Azure Storage Hierarchy

The Cohesity cluster allows you to backup Azure Storage Container at different levels of the hierarchy.

You can choose to back up any of these levels in hierarchy:

- **Tenant level**– Protect all Azure Storage Accounts (Blob/ADLS) across all levels.
- **Subscription level** – Protect all Azure Blob or Data Lake Storage Accounts within the Subscription.
- **Resource Group level**– Protect all Azure Blob Storage or ADLS Storage Accounts within the subscription.
- **Storage Account level**– Protect all containers within a selected Azure Blob or ADLS Storage Account.
- **Storage Container** – Protect individual containers within an Azure Blob or ADLS Storage Account.

Supported Backup Types

Cohesity supports both full and incremental backups for Azure Blob and Data Lake Storage.

Protect Azure Blob or Data Lake Storage Containers

In Cohesity, the containers of Azure Blob and Data Lake Storage Accounts are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

To protect a Storage Container:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **Azure** > **Blob and Data Lake Storage**.
3. In the **New Protection** page, click **Add Objects** and select your registered source or click **Register Source** to add a new one.
4. Select the Storage Container using one of the following views or option:
    - **List View**: Use this option to select the Storage Container for protection from a flat, searchable list.
    - **Hierarchy View**: Use this option to select Storage Container for protection at various hierarchy levels.

      Optionally,

        - You can click the auto-protect icon at the tenant level, subscription level, and Storage account level to automatically protect newly added Storage Container.
        - You can exclude the Storage Container from auto-protection within an tenant, Subscription, resource group, or Storage Account by clicking the auto-protect icon ( ) next to the Storage Container. The Storage Container excluded from auto-protection are displayed with the exclude ( ) icon.
    - **Tags**: Use this option to select the Storage Container for protection based on tags. Select the checkbox of a tag, and then select **Select All Child Objects** to include all containers with that tag for protection.

      Optionally:

        - You can configure auto-protection based on a tag. You can select the tag checkbox and then choose **Auto Protect This Tag** to automatically protect all newly added Storage Containers with this tag.
        - You can exclude the Storage Containers from auto-protection within a tag by clicking the auto-protect icon ( ) next to the Storage Container. The Storage Containers excluded from auto-protection are displayed with the exclude ( ) icon.
        - You can exclude Storage Containers with a specific tag from auto-protection by clicking the exclude icon (⊘). The tag excluded from auto-protection are displayed with the exclude icon.
5. Click **Continue**.
6. Provide a name for the **Protection Group**. The name can contain alphanumeric characters, underscores, dashes, and periods.

7. From the **Policy** drop-down, create or select and existing policy.
8. Select **More Options** to change or configure any of the advanced settings.
9. Click **Protect**.

Advanced Settings

Section	Setting	Description
Settings	**Start Time**	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
	SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

Section	Setting	Description
Additional Settings	**Pause Future Runs**	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
	End Date	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
	QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: This policy achieves similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. The data is written to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. For best performance, we recommend the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
	Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Section	Setting	Description
	Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: • High-priority protection runs • Medium-priority protection runs • Low-priority protection runs
	Description	**Optional**. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Storage Container Backup Configuration

In the Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

Depending on the Protection Type you chose while creating the Protection Group, you can add new Storage Containers to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Storage Container.

3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, click the plus icon (**+**) to enter multiple objects that you plan to back up. Click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Storage Container.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Storage Container.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the product UI.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Storage Container.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - **Delete Protection Group Only**
    - **Delete Protection Group and Snapshots**

Recover Azure Blob or Data Lake Storage Account Containers

You can restore the Storage Containers of Azure Blob or Data Lake Storage Account either to its original location or to a different location.

> **Note:** We recommend that you review the Plan and Prepare for Azure Blob and Data Lake Storage Protection.

To recover a Storage Container:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Azure** > **Blob and DataLake Storage**.
3. Search for the Storage Container by entering characters of the Storage Container name, or by specifying the wildcard character *.
4. From the list of search results, select the Storage Container to recover. By default, the latest snapshot of the Storage Container is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

   From the selected lists, mouse over the snapshot, and then click the edit icon. The **Edit recovery point for *<container_name>*** page appears. By default, the option to select the snapshot is displayed in the **List** mode. You can also choose to view the option in the **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the **Edit recovery point for *<container_name>*** page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
     b. Select the snapshots to be recovered.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
     b. Select the snapshot to be recovered by clicking the corresponding blue dot.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
5. Click **Next: Recover Options**.

6. In the **Recover To** section, select any one of the following options based on where you want to restore the container and then provide the required information:
   - **Original Location**: Restore into the original container.
     a. Enable **Overwrite existing container(s)** to recover the Storage Container by overwriting the original container. Type **Yes** to confirm.

        If the source and target objects are identical, then the target container will not be updated even if the **Overwrite existing container(s)** option is enabled.

     b. Optionally, in the **Prefix or file path (optional)** field, add a prefix or file path to the files of the containers that are recovered.
   - **New Location**: Restore into another container within the same or a different Azure source.
     a. From the **Source** drop-down, select a registered Azure source and from the **Container** drop-down, select the target Storage Container to which you want to restore the object from the source containers.

        If the source and target objects are identical, then the target container will not be updated even if the **Overwrite existing container(s)** option is enabled.

     b. To overwrite any existing Storage Container with the same name, enable **Overwrite existing container(s)** and type **Yes** to confirm.

        If the source and target objects are identical, then the target container will not be updated even if the **Overwrite existing container(s)** option is enabled.

     c. Optionally, in the **Prefix or file path (optional)** field, add a prefix or file path to the files of the containers that are recovered.

   > **Note:** If **Overwrite existing container(s)** is enabled, any existing object in the target container with the same name as a recovered object is overwritten. Cohesity does not delete or replace the Azure Storage Container itself.

7. In the **Recovery Options** section, configure the following options:
   a. **Granular recovery**: To perform granular recovery the Storage Container, toggle on **Enable to recover specific prefixes or file paths** and provide the prefix of the files to be recovered.

      For example, **Finance/statement/document**.

      You can click the add icon to provide multiple prefixes.

The maximum number of prefixes that you can provide per recovery task is 100.

b. **Continue on Error**: Enable this option if you want to continue the recovery even if one of the files in the container encounters an error. By default, this option is disabled and the recovery operation will fail if one of the files encounters an error.

c. **Task Name**: Update the default name of the recovery task.

8. Click **Recover**.

GCP

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud Services, such as virtual machines, databases and so on.

This chapter includes the following topics:

Register a Google Cloud Source

To protect your GCP workload, you need to register your Google Cloud environment as a source on the Cohesity cluster. After registering your Google Cloud environment as a source, you can:

- Update the Google Cloud source configuration.
- Refresh the source details.
- Unregister the Google Cloud source from the Cohesity cluster.

Google Cloud Source Registration Requirements

The following table lists the requirements that are needed only for source registration.

Requirement	Details
Service Account	Create a dedicated service account in Google Cloud for Cohesity. Assign the required permissions to the service account.
Service Account Key	Generate a JSON key for the service account to allow authentication from outside Google Cloud.
Project ID	The Google Cloud Project ID of the project associated with the service account.

Requirement	Details
Virtual Private Cloud (VPC)	The name of the Virtual Private Cloud network to connect the Cohesity cluster with Google Cloud.
Subnet	The name of the subnet within the VPC for network connectivity.
Network Connectivity	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	Allowlist the required fully qualified domain names (FQDNs) to allow Cohesity to communicate with supported services during source registration operations. **Note:** This FQDNs are also required during backup and recovery.

Register Google Cloud Source

Before registering Google Cloud as a source, make sure you have completed all prerequisite tasks and reviewed the considerations.

To register Google Cloud source:

1. Navigate to **Data Protection > Source**.
2. Click **Register** at the top-right of the page and then select **Cloud Platforms > Google Cloud Platform**.
3. In the **Register GCP Source** dialog box, configure the following and click **Continue**.

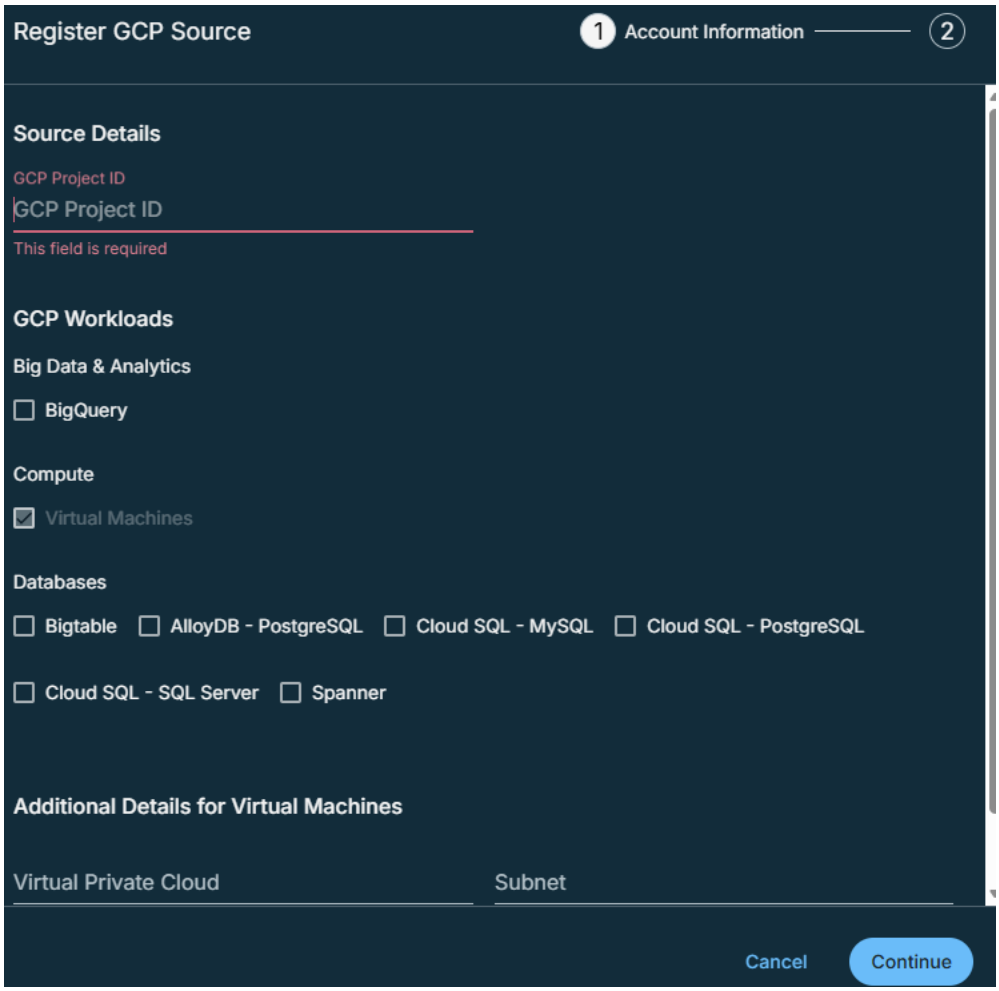

Field	Description
GCP Project ID	The Google Cloud Project ID of the project associated with the service account.
GCP Workloads	Select the required workload.
Virtual Private Cloud	The name of the VPC network that connects the Cohesity cluster to Google Cloud.
Subnet	The subnet within the VPC for network connectivity.

4. In the Authentication section of the dialog box, choose any one of the following authentication types:

- **Upload JSON Key File** - Upload the service account JSON key file. You can either drag and drop the file or click to upload the JSON file.
- **Enter Email and Private Key** - Manually enter the service account email address and private key associated with the service account.

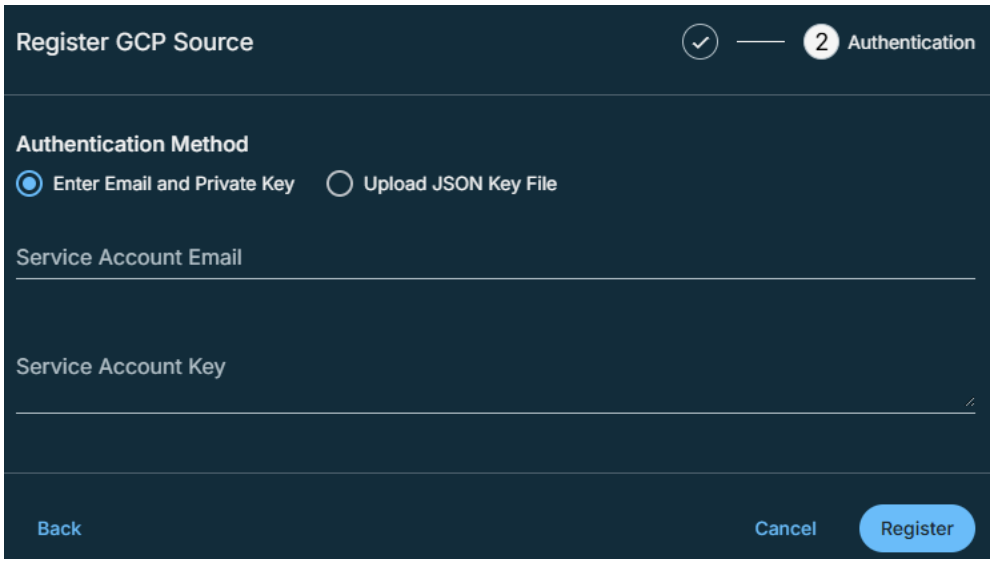

5. Click **Register**.

Port Requirements

Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data for backup and recovery from Google Cloud Source using TCP/IP protocol.

Ports	Worloads
443	• MySQL • Firestore • Spanner • Bigquery • SQL Server • AlloyDB • PostgreSQL • Bigtable
3306	• MySQL

Ports	Worloads
9010	• Spanner
1433	• SQL Server
5432	• AlloyDB • PostgreSQL

Required Permissions in GCP for Workloads

The following lists of workloads specifies the permissions and privileges that must be included in a custom role. Assign this role to the service account to enable GCP workload protection.

- Required Permissions for GCP VM Data Protection
- Required Permissions for BigQuery Data Protection
- Required Roles for Spanner Data Protection
- Required Roles for Firestore Data Protection
- Required Permissions for Cloud SQL for MySQL Data Protection
- Required Permissions for Cloud SQL for PostgreSQL Data Protection
- Required Permissions for AlloyDB for PostgreSQL Data Protection
- Required Permissions for Cloud SQL for SQL Server Data Protection
- Roles for Bigtable Data Protection

Required Permissions for GCP VM Data Protection

The following tables list the GCP permissions required for cloud services.

GCP Native Backup and Recovery

Following are the permission required for GCP native backup and recovery:

```
compute.addresses.createInternal
compute.addresses.deleteInternal
compute.disks.create
compute.disks.createSnapshot
compute.disks.delete
compute.disks.get
compute.disks.list
compute.disks.setLabels
compute.disks.use
compute.images.useReadOnly
```

```
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.get
compute.instances.list
compute.instances.setDeletionProtection
compute.instances.setLabels
compute.instances.setMetadata
compute.instances.setTags
compute.instances.setServiceAccount
compute.instances.stop
compute.instances.update
compute.networks.get
compute.networks.list
compute.projects.get
compute.regions.list
compute.snapshots.create
compute.snapshots.delete
compute.snapshots.get
compute.snapshots.setLabels
compute.snapshots.useReadOnly
compute.subnetworks.list
compute.subnetworks.use
compute.subnetworks.useExternalIp
compute.zones.list
compute.zoneOperations.get
iam.serviceAccounts.actAs
resourcemanager.projects.get
vpcaccess.connectors.list
vpcaccess.locations.list
```

Cloud Archival

Action	Permissions
Cloud Archive	`storage.buckets.get` `storage.objects.create`
Cloud Recovery	`storage.buckets.get` `storage.objects.create`

Cloud Edition

Action	Permissions
Cloud Edition deployment using Helios	cloudfunctions.functions.call cloudfunctions.functions.create cloudfunctions.functions.delete cloudfunctions.functions.get cloudfunctions.functions.list cloudfunctions.locations.list compute.addresses.createInternal compute.addresses.deleteInternal compute.addresses.delete compute.addresses.get compute.addresses.list compute.addresses.useInternal compute.disks.create compute.disks.delete compute.disks.get compute.disks.list compute.disks.use compute.disks.useReadOnly compute.images.get compute.images.useReadOnly compute.instances.attachDisk compute.instances.create compute.instances.delete compute.instances.get compute.instances.list compute.instances.setDeletionProtection compute.instances.setLabels compute.instances.setMetadata compute.instances.setTags compute.instances.start compute.instances.stop compute.instances.update compute.networks.get compute.networks.list compute.projects.get compute.regions.get compute.regions.list compute.subnetworks.get compute.subnetworks.list

Action	Permissions
	compute.subnetworks.use compute.subnetworks.useExternalIp compute.zones.get compute.zones.list iam.serviceAccounts.actAs resourcemanager.projects.get resourcemanager.projects.getIamPolicy storage.buckets.get storage.objects.create storage.objects.delete storage.objects.get storage.objects.list storage.objects.update vpcaccess.connectors.get vpcaccess.connectors.list vpcaccess.locations.list vpcaccess.operations.get vpcaccess.operations.list

Action	Permissions

Cloud Edition deployment using the setup utility on the Control VM.	`compute.addresses.createInternal`
	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.list`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setServiceAccount`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.instances.update`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions
Add or remove a node using the setup utility on the Control VM.	compute.addresses.delete compute.addresses.get compute.addresses.list compute.addresses.useInternal compute.disks.create compute.disks.delete compute.disks.get compute.disks.list compute.disks.use compute.disks.useReadOnly compute.images.get compute.images.useReadOnly compute.instances.attachDisk compute.instances.create compute.instances.delete compute.instances.get compute.instances.list compute.instances.setDeletionProtection compute.instances.setMetadata compute.instances.setTags compute.instances.start compute.instances.stop compute.networks.get storage.buckets.get storage.objects.create storage.objects.delete storage.objects.get storage.objects.list storage.objects.update compute.subnetworks.get compute.subnetworks.use

Action	Permissions

Add or remove a node using an existing node in the cluster.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.list`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions

Power on the cluster using the setup utility on the Control VM.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions

Power off the cluster using the setup utility on the Control VM.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions

Destroy the cluster using the setup utility on the Control VM.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.list`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Shared VPC

During source registration, if you are registering a shared VPC, enter the VPC and the subnet name associated with the VPC in the following format:

`<host-project-id>/<vpc-name>` where `<host-project-id>` is the host project in Shared VPC.

`<host-project-id>/<subnet-name>` where `<host-project-id>` is the host project in Shared VPC.

The registered service account must have the Compute Network User IAM role in the host project. For more information, see Register or Edit a GCP Source.

Required Permissions for BigQuery Data Protection

The following table lists the permissions and privileges that must be included in a custom role. Assign this role to the service account to enable BigQuery data protection.

Type	Permission	Purpose
Google Cloud Storage	`storage.buckets.get`	Retrieve the metadata of a specific Cloud Storage bucket
	`storage.buckets.getIamPolicy`	Read the Identity and Access Management (IAM) policy of a specific Cloud Storage bucket.
	`storage.buckets.setIamPolicy`	Ability to read the Identity and Access Management (IAM) policy (ACLs) for a specific Cloud Storage bucket.
	`storage.objects.create`	Create new objects (files) within a specified Cloud Storage bucket
	`storage.objects.delete`	Ability to delete objects (files) from a Cloud Storage bucket.
	`storage.objects.get`	Ability to read the data and metadata of an object stored in a Cloud Storage bucket.
	`storage.objects.list`	List the objects (files) within a specified Cloud Storage bucket.

Type	Permission	Purpose
BigQuery	`bigquery.config.get`	View project-level BigQuery configuration settings.

Type	Permission	Purpose
	`bigquery.datasets.create`	Create new datasets within a project.

Type	Permission	Purpose
	`bigquery.datasets.delete`	Delete existing datasets.

Type	Permission	Purpose
	bigquery.datasets.get	Retrieve metadata for a specific dataset.
	bigquery.jobs.create	Submit new query or load jobs.
	bigquery.jobs.get	View metadata and status of a specific job.
	bigquery.jobs.list	List jobs submitted by the current user.
	bigquery.jobs.listAll	List all jobs in the project, including those submitted by other users.
	bigquery.jobs.update	Update job metadata (for example, labels).
	bigquery.routines.get	Retrieve metadata for a specific routine (for example, UDF or stored procedure).
	bigquery.routines.list	List routines in a dataset.
	bigquery.tables.create	Create new tables in a dataset.
	bigquery.tables.createSnapshot	Create a snapshot of an existing table for backup or archival.
	bigquery.tables.delete	Delete tables from a dataset.
	bigquery.tables.deleteSnapshot	Delete a snapshot table.
	bigquery.tables.export	Export table data to external storage (for example, Cloud Storage).
	bigquery.tables.get	Retrieve metadata for a specific table.
	bigquery.tables.getData	Read data from a table (used in queries or API calls).
	bigquery.tables.list	List all tables in a dataset.
	bigquery.tables.update	Update table metadata (for example, schema, description, labels).
	bigquery.tables.updateData	Modify table data (for example, insert, update, delete rows).

Required Roles for Spanner Data Protection

The following table lists the roles that must be assigned to the service account to enable Spanner data protection.

Role Name	Purpose/Scope
Cloud Spanner Admin	Manage Spanner resources
Dataflow Admin	Manage Dataflow jobs for export/import
Storage Admin	Access and manage GCP buckets
Log Viewer	View logs for troubleshooting
DataBoost Role	Use DataBoost for performance
Cloud Spanner Reader	Read database data
Cloud Spanner Database Admin	Administer database operations

> **Note:**
> - Assign these roles to the service account at the appropriate scope (project, database, or database) to ensure successful backup and restore operations.
> - For DataBoost-enabled exports, the additional DataBoost role is required.

Required Roles for Firestore Data Protection

The following table lists the roles that must be assigned to the service account to enable Firestore data protection.

Role Name	Purpose/Scope
Storage Admin	Access and manage GCS buckets.
Firebase Admin	List, Create, Delete, Export and Import Firestore database.

Required Permissions for Cloud SQL for MySQL Data Protection

The following table lists the permissions and privileges that must be included in a custom role. Assign this role to the service account to enable Cloud SQL for MySQL data protection.

Permission	Purpose
`cloudsql.instances.list`	Allows listing the Cloud SQL instances in the project.
`cloudsql.databases.list`	Allows to retrieve the list of databases.
`cloudsql.databases.get`	Allows to retrieve details of a specific database.
`cloudsql.instances.get`	Allows to retrieve details of a specific Cloud SQL instance.

Required Permissions for Cloud SQL for PostgreSQL Data Protection

The following table lists the permissions and privileges that must be included in a custom role. Assign this role to the service account to enable Cloud SQL for PostgreSQL data protection.

Permission	Purpose
`cloudsql.instances.list`	Allows listing the Cloud SQL instances in the project.
`cloudsql.databases.list`	Allows to retrieve the list of databases.
`cloudsql.databases.get`	Allows to retrieve details of a specific database.
`cloudsql.instances.get`	Allows to retrieve details of a specific Cloud SQL instance.

Required Permissions for AlloyDB for PostgreSQL Data Protection

The following table lists the permissions and privileges that must be included in a custom role. Assign this role to the service account to enable AlloyDB for PostgreSQL data protection.

Permission	Purpose
`alloydb.clusters.list`	Allows you to view a list of all AlloyDB clusters in your project.
`alloydb.clusters.get`	Enables you to retrieve detailed information about a specific AlloyDB cluster.
`alloydb.instances.get`	Grants the ability to list all AlloyDB instances within a cluster.
`alloydb.instances.list`	Allows you to access detailed information about a specific AlloyDB instance.

Required Permissions for Cloud SQL for SQL Server Data Protection

The following table lists the permissions and privileges that must be included in a custom role. Assign this role to the service account to enable Cloud SQL for SQL Server data protection.

Type	Permission	Purpose
Google Cloud Storage	`storage.buckets.create`	Allows creating new Google Cloud Storage (GCS) buckets within the project.
	`storage.buckets.get`	Ability to retrieve the metadata of a specific Cloud Storage bucket
	`storage.buckets.getIamPolicy`	Ability to read the Identity and Access Management (IAM) policy of a specific Cloud Storage bucket.
	`storage.buckets.setIamPolicy`	Ability to set the Identity and Access Management (IAM) policy (ACLs) for a specific Cloud Storage bucket.
	`storage.objects.create`	Ability to create new objects (files) within a specified Cloud Storage bucket
	`storage.objects.delete`	Ability to delete objects (files) from a Cloud Storage bucket.
	`storage.objects.get`	Ability to read the data and metadata of an object stored in a Cloud Storage bucket.
	`storage.objects.list`	Ability to list the objects (files) within a specified Cloud Storage bucket.
Cloud SQL for SQL Server	`cloudsql.databases.get`	Ability to retrieve metadata for a specific database within a Cloud SQL instance.
	`cloudsql.databases.list`	Allows listing all databases inside a Cloud SQL instance.
	`cloudsql.instances.export`	Allows initiating an export operation on a Cloud SQL instance.
	`cloudsql.instances.get`	Allows reading configuration details for a Cloud SQL instance, such as version, region, and settings.
	`cloudsql.instances.import`	Allows starting an import operation on a Cloud SQL instance.
	`cloudsql.instances.list`	Allows listing all Cloud SQL instances in the project.

Roles for Bigtable Data Protection

The following table lists the common predefined. Assign the role to the service account to enable Bigtable data protection.

> **Note:** See Google Bigtable IAM roles for full permission mapping.

Role	High level coverage of the role
`roles/bigtable.admin`	This is the admin role which includes create/update/delete tables and read/write rows, plus instance/cluster admin access.
`roles/bigtable.user`	This is the non-admin role with read/write access to the table data and some metadata.
`roles/dataflow.developer` (or `roles/dataflow.admin`)	This role is specific during backup and restore operations. This Dataflow job submission/management permission includes dataflow.jobs and related.
`roles/dataflow.worker`	This role is about the permissions for Dataflow workers to execute the pipeline and access resources.
`roles/storage.admin`	This role provides full administrative access to Cloud Storage. It allows creating, updating, and deleting buckets and objects, as well as managing bucket-level settings and permissions.

Create an IAM Custom Role in GCP

To enable GCP workload data protection, you must create an IAM custom role in Google Cloud and assign it to the service account used by the Cohesity cluster.

For detailed instructions, refer to the following topics in the Google Cloud documentation:

Before you begin creating IAM custom role, ensure that you have created a dedicated service account in Google Cloud for the Cohesity cluster.

> **Warning:**
> This task is NOT applicable for GCP Firestore and Spanner.

- Create a custom IAM role.
- Assign a role to a service account.

Manage GCP Source Registration

After registering Google Cloud as a source in the Cohesity cluster, you may have updated your environment in one or more of the following ways:

- Changed the project associated with the service account.
- Updated VPC or subnet details.
- Generated a new JSON key file for the service account

You can update the Google Cloud registration details provided during the source registration with the latest Google Cloud configuration details.

To edit the Google Cloud source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click action menu (⋮) next to the Google Cloud source and select **Edit**.
3. In the **Edit GCP Source** page, update the required configurations.
4. Click **Update**.

Refresh Google Cloud Source

You can refresh the Google Cloud source configuration and fetch the latest changes on your Google Cloud environment.

To refresh the Google Cloud source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click action menu (⋮) next to the Google Cloud source and select **Refresh**.

Unregister Google Cloud Source

If you plan to stop taking the backup of your Google Cloud source, you can unregister the Google Cloud source from the Cohesity cluster.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the Google Cloud source.

To unregister the Google Cloud source:

1. Navigate to **Data Protection > Sources**.
2. Click action menu (⋮) next to the Google Cloud source and select **Unregister**.

GCP Virtual Machines

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud Services, such as virtual machines.

This chapter includes the following topics:

Plan and Prepare for GCP Protection

Review the following prerequisites and considerations of GCP.

Prerequisites

- To add a new node to a Cohesity cluster, both the node and Cohesity cluster must have Cohesity 6.0 or later installed. For more information, see Manage Nodes.
- Make sure `googleapis.com` is reachable through firewall port 443 when registering the GCP Source.
- You must allowlist the following source FQDNs for registering and backing up the GCP source:
  - https://accounts.google.com/
  - https://www.googleapis.com/
  - https://storage.googleapis.com
  - https://vpcaccess.googleapis.com/
  - https://cloudfunctions.googleapis.com/
  - https://cloudresourcemanager.googleapis.com/

Considerations

Review the following considerations of GCP native snapshotting:

- App consistent backup is not supported.
- Recovering files/folders to source is not supported. File download is supported.
- Cloning is not supported.
- For incremental snapshots, native backup uses source side dedup and not Google API as Google does not have any APIs for incremental backup. Hence full and incremental

backups after the first run takes about the same time to complete.

- For backup and recovery of GCP VMs, the service account, the VPC and Subnet where the fleet instances are launched should be on the same project.
- You can perform files and folder recovery to an GCP VM Linux instance only if the GCP VM Linux instance is of x86_64 architecture.
- File-level recovery is not supported for GCP VMs with XFSv5 file system.
- You cannot exclude disk at the object-level for auto-protected VMs. To exclude disk for auto-protect VMs, you must use the **Volume Exclusion Settings** option in the Additional Settings.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.

Create an IAM Custom Role for GCP

You need to create IAM custom roles when registering GCP as a source on the Cohesity cluster.

Prerequisites

Before you begin creating IAM custom roles, ensure the following prerequisites are met:

- Enable IAM API. For more information, see Enable GCP APIs.
- Create service accounts. For more information, see Creating and managing service accounts.

Perform the following steps to create an IAM custom role:

1. In the Google Cloud Platform Console, navigate to **IAM & Admin**.
2. From the drop-down located at the top of the page, select the project in which you want to create the custom role.

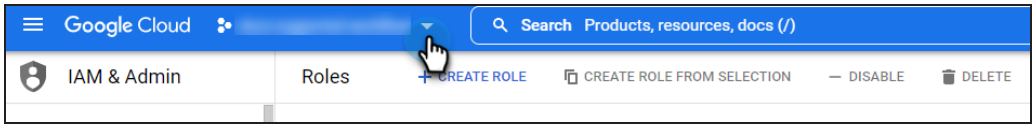

3. Click **Create Role**.
4. In the **Title** field, enter a title for the role.
5. In the **Description** field, enter a brief description.
6. In the **ID** field, enter the ID.
7. From the **Role launch stage** drop-down, select the role launch stage for the role.

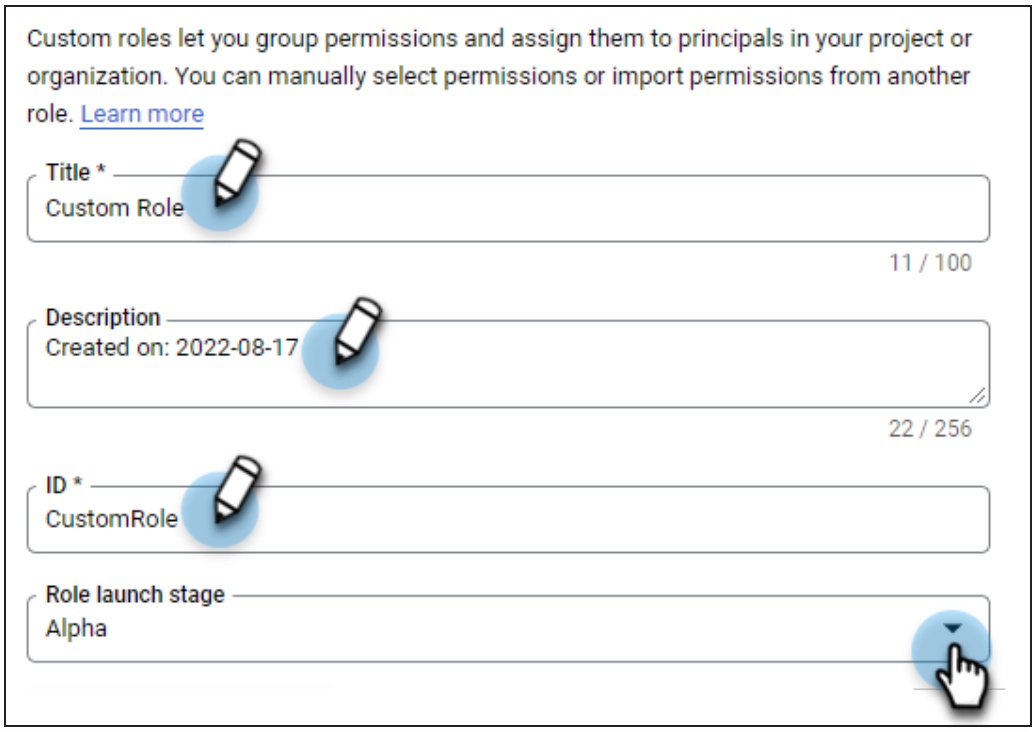

8. Click **Add Permissions**. The **Add Permissions** page is displayed.
9. Select the permission you want to include in the role from the list or filter the permissions by role. For GCP permissions required for cloud services, see GCP Permissions.
10. Click **Add** to add the permissions to the role.

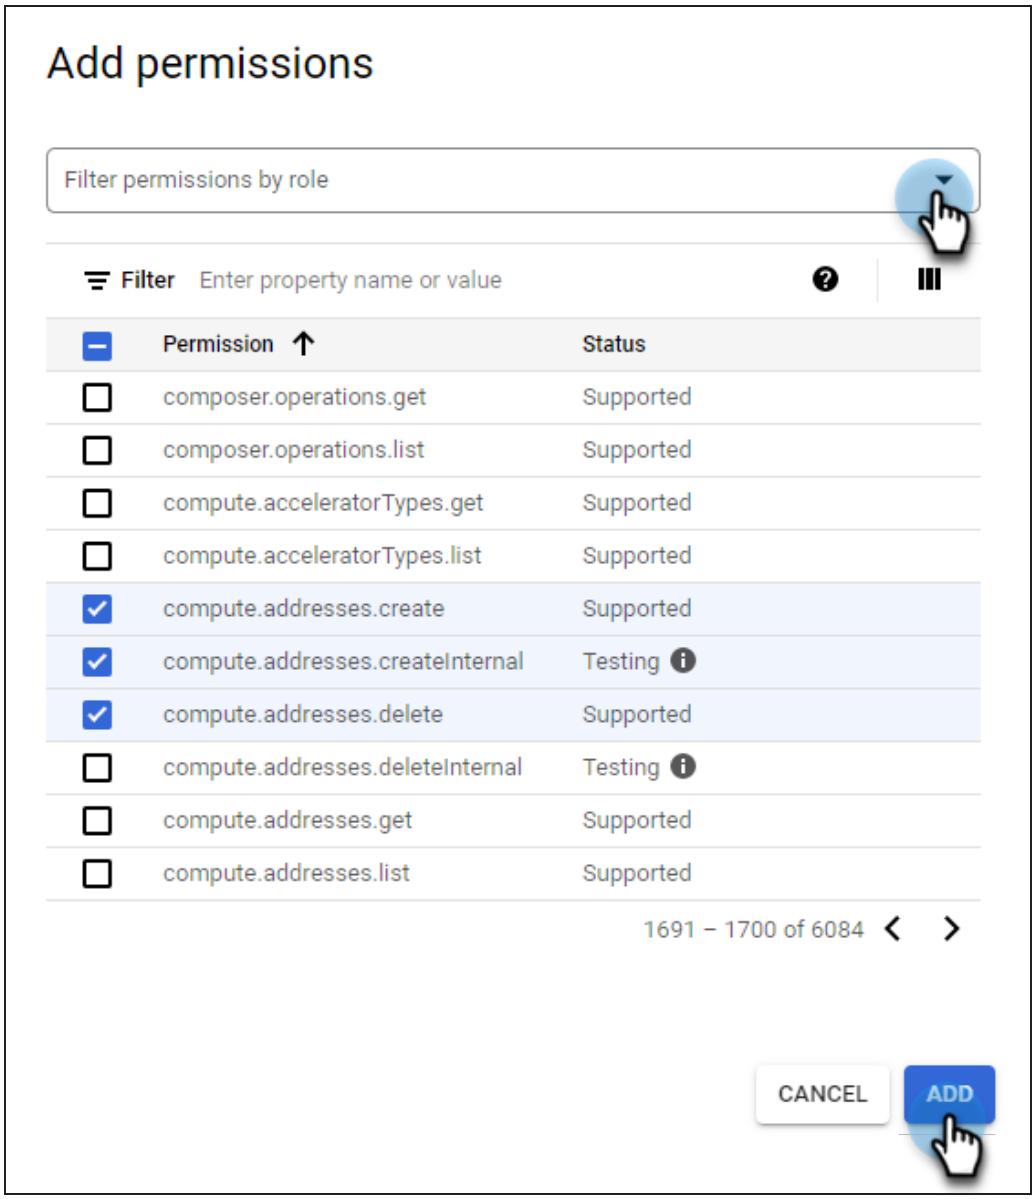

Enable GCP APIs

The following GCP APIs must be enabled for GCP native backup support:

- Compute Engine API
- Cloud Resource Manager API

To enable the APIs:

1. In the Google Cloud Platform Console, in the project you are using for GCP, go to **APIs & Services > Library**.
2. Search for one of the APIs listed above and then click the API tile as shown below.

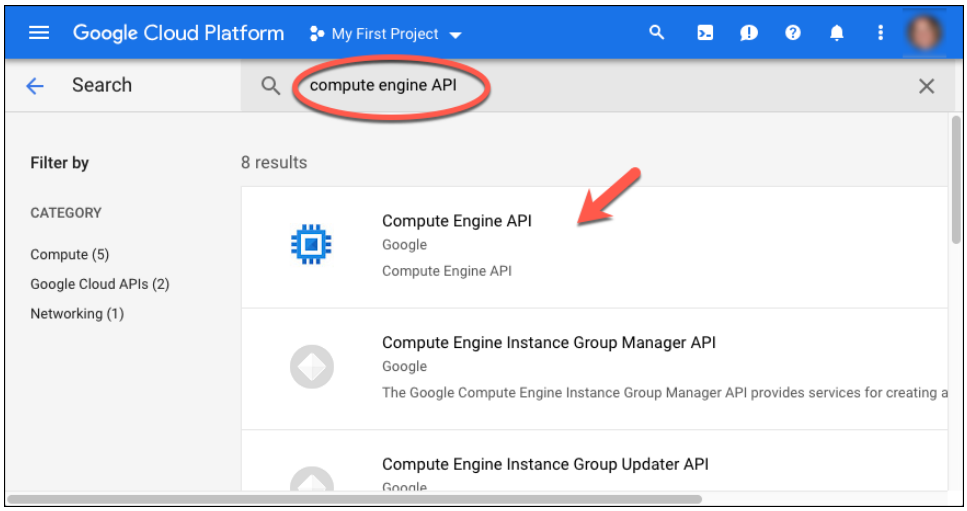

3. If the tile indicates **API enabled**, perform the previous step for the next API.

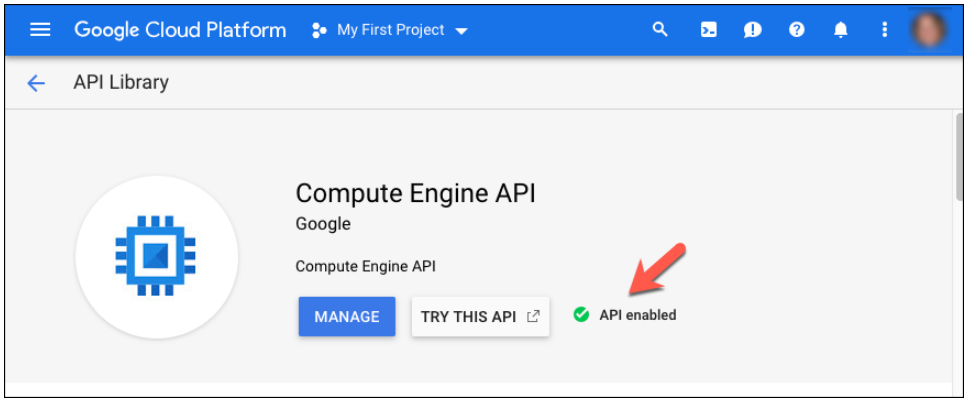

If the tile displays the **Enable** button, click it and respond to any billing prompts as needed.

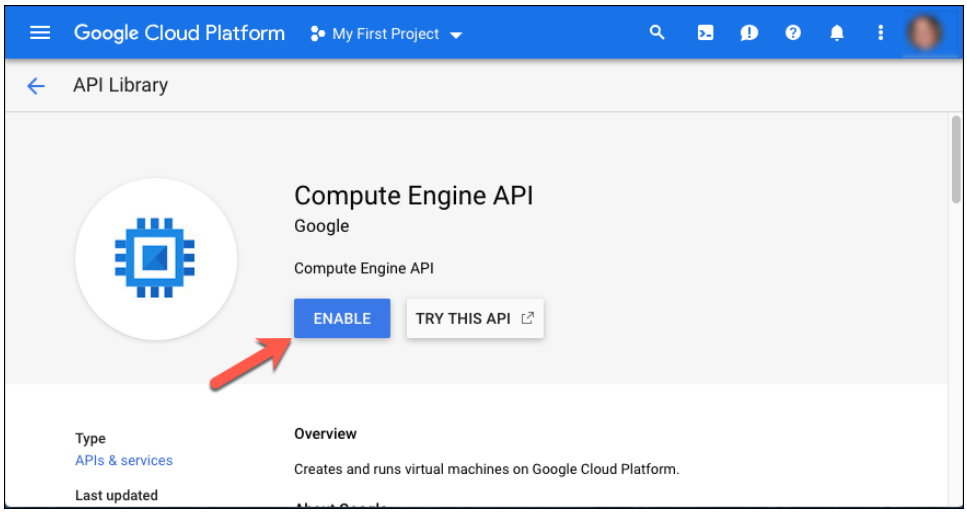

4. Repeat the steps above for the other APIs.

Add Project IDs to Allowlist

Add the following project IDs to the allowlist on the GCP Organization Policies page if the GCP organization has a policy that restricts the use of images belonging to external projects. You must add the project IDs to the allowlist for creating Cohesity cluster and backup and recovery of native VMs.

Project IDs	Description
cloud-edition-exports	Required for creating Cohesity cluster.
cohesity-public	Required for backup and recovery of native VMs.

Register or Edit a GCP Source

Before you can register the GCP source:

- Create an IAM Custom Role for GCP
- Enable GCP APIs
- Review and understand the Plan and Prepare for GCP Protection.

After creating the IAM user and enabling the APIs, you can register the source:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Select **Register > Virtual Machines**.
3. From the **Select Hypervisor Source Type** drop-down, select **Cloud > GCP: IAM User**.
4. Either enable the **Use JSON Key File** option and upload the **Service Account JSON Key File** or disable the **Use JSON Key File** option and provide the following details:

Field Name	Description
Client Email Address	The client email address for the GCP service account associated with the project. You can find the client email address on the IAM page in the Google Cloud Platform Console. The service account must be assigned the Editor role in GCP. Example: `servaccnt@project-name.iam.gserviceaccount.com` **Tip:** The portion of the email address to the left of the at symbol @ will be used as the Source name. For the example above, the Source name will be `servaccnt`. **Note:** A single service account can be used to protect multiple GCP projects. When a service account is registered on the Cohesity cluster as a source, all GCP projects in which the service account has Editor access are made available for protection.
Client Private Key	The private key for the service account. The key could have been downloaded as a JSON file and retained when the service account was created, or the key could have been generated and downloaded after the account was created. Paste the entire private key into this field, from `-----BEGIN PRIVATE KEY-----` to `-----END PRIVATE KEY-----`.

Field Name	Description
Fleet Settings	Specify the VPC and subnet in which the GCP fleet instance should be launched. You can specify the fleet instance location twice as primary and secondary so that if the fleet instance location specified in the primary is unavailable, the location specified in the secondary will be considered. Following are the fleet instance location you can choose: **Same as Source VM**: (Default behavior) The fleet instance will be launched in the same VPC and subnet of the VMs that are to be protected. **Note:** It is recommended to launch the fleet instance in the same VPC and subnet as source VMs to avoid cross-region charges for data read operations. • **Specify ProjectID, VPC, and Subnets**: The fleet instance will be launched in one of the specified subnets, depending on the IP availability in the subnet. If the GCP source is not registered with the Cohesity cluster, the cluster cannot detect the Project and the associated VPC and subnets. To configure the source for launching the fleet instance, you must first register the source and then update the registered source to specify the Project IDs, VPCs, and subnets. For more information, see Configure Project IDs, VPCs and Subnet for Fleet Instance. **Note:** If you have not selected the **Fleet Settings** option, then the GC fleet instance is launched in the same VPC and subnet of the VMs that are to be protected.

5. Click **Register**.

   The following figure is the screenshot of the GCP registeration page:

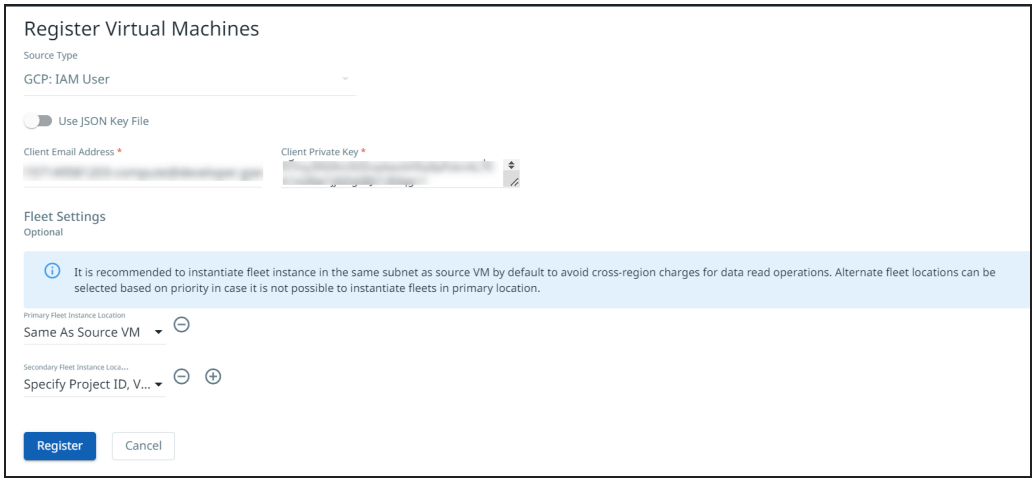

Customer-Managed Key Restoration

Cohesity now supports backing up encrypted GCP VMs. The source VM encrypted using Customer-Managed Key encryption can now be recovered back to the same location. The recovered VM is encrypted with the same key as the original source VM.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Configure Project IDs, VPCs and Subnet for Fleet Instance

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

To configure a new VPC and subnet for the GCP fleet instance:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Navigate the hierarchy to the registered GCP source for which you want to configure a new VPC and subnet for the fleet instance.
3. Click the actions menu next to the GCP source and select **Edit**.

   The Edit Source page appears.
4. From the **Primary Fleet Instance Location** drop-down list, select **Specify ProjectID, VPC, and Subnet** and provide the following information:
   1. **Project ID**: Select the Google Cloud Platform (GCP) project which will be used by the Cohesity cluster to launch the temporary proxy VM (fleet instance) used for backup and recovery operations.

      Example: `my-company-project`
   2. **VPC**: Select the Virtual Private Cloud network name associated with the project. The Cohesity cluster will use this VPC to launch the temporary proxy VM used during backup and recovery operations.
   3. **Subnet**: A subnet name associated with the VPC. The Cohesity cluster will use this subnet to launch the temporary proxy VM used during backup and recovery operations. The subnet belongs to a specific region.

      Example: `mycompany-google-subnet-us-west1`
5. (*Optional*) Click the add icon to specify multiple numbers of fleet instance locations. The maximum number of fleet instance locations you can specify is 20. If the subnet you specified in the first entry is unavailable or the IP in the subnet is exhausted, then based on the order you specified the instance location, the fleet instance will be launched in the next subnet and so on, until Cohesity finds a subnet that is available. If none of the subnets are available, then the fleet instance will either be launched in the secondary location (if present) or the backup and recovery of GCP VMs will fail.

6. Click **Save**.

   The GCP fleet instance will be launched in the specified subnet when you backup the VMs of the GCP cloud service to the Cohesity cluster.

   The following animated image shows how to configure Project ID, VPC, and subnet to launch the fleet instance:

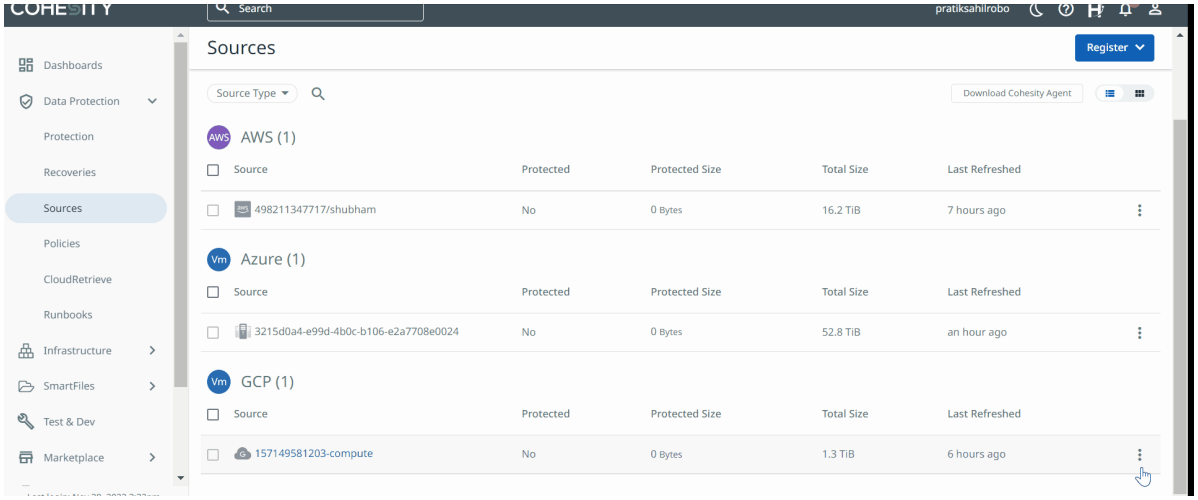

Unregister GCP Source

If you plan to stop taking backup of the GCP VMs, you can unregister the GCP source from the Cohesity cluster. Before you unregister, you must delete all the Protection Groups associated with the GCP source.

To unregister a GCP source on the Cohesity cluster:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the GCP source and select **Unregister**.
3. Click **Unregister**.

> **Note:** When you unregister a source, all the objects in the source will be unprotected. Any new protection after the re-registration of the source will trigger a full backup and not an incremental backup.

Protect GCP VMs

In Cohesity, GCP VMs with their shielded VM properties are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Create or Edit a Standard Policy.

> **Note:** Backup and recovery of GCP VMs requires a proxy VM (fleet instance) to be launched on the cloud source. For enhanced security, after creating the proxy VM on the cloud source, Cohesity automatically deploys an X.509 certificate. Contact Cohesity Support if you want to launch the proxy VM with its Shielded VM properties enabled.

To protect GCP VMs:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **Virtual Machines**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the GCP source you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it.
4. From the objects list, select the VMs you want to protect.

   Optionally, you can:

   - Click the auto protect ( ) icon to enable **Auto Protect** at various hierarchy levels (Source, Project, Region, Availability Zone). The auto-protect option enables you to automatically protect the new VMs that are added.
   - Click the **Tags** icon on the right and select a tag to:
     - Auto Protect ( ) the VMs having this tag.
     - Exclude ( ) the auto protection of VMs having this tag.

       If a VM has multiple tags with Auto Protect and Exclude applied, exclusion takes precedence. These excluded VMs are also excluded even if a parent object is auto protected.

     The **Tags** button will not be displayed if the Cohesity cluster detects that no sources have tags assigned to them.
   - Exclude any disks of the VM from protection. To exclude the disks:
     a. Click the edit icon displayed next to the VM of which you want to exclude the disk.
     b. Unselect the disks which you want to exclude from protection.

> **Note:**
> - You cannot exclude the root disk from protection.
> - You cannot exclude disk at the object-level for auto-protected VMs. To exclude disk for auto-protect VMs, you must use the **Volume Exclusion Settings** option in the Additional Settings.

5. Click **Continue**.
6. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
8. If you want to change or configure any of the additional setings, then select **More Options** and perform the below steps or else, click **Protect**.
9. Optionally, from the **Policy** drop-down list, you can change the policy you selected earlier or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
10. In the **Storage Domain** field, select the appropriate storage domain.
11. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
12. In the **SLA** field, select how long the administrator expects a Protection Group run to take.
    1. **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
    2. **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run.
13. If you need to change any of the advanced settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
14. Click **Protect**.

    The protection of GCP VMs is now active and running.

Additional Settings

Additional Settings	Description
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Time	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Indexing	This option is toggled on by default to index the data of the Protection Group. Indexing improves files and folders search. Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.

Additional Settings	Description
Exclusions	Enable **Exclude Disks** to specify the disks to exclude for all VMs in this Protection Group. You can use the following parameters to form a boolean expression for excluding disks: • **Type**: The type of persistent disk - Standard, Balanced, SSD, Extreme. • **Name**: Name of the disk. • **Label**: Label name tagged with the disk. For example, if you have provided the following boolean expression for exclusion: (`'Type' = Balanced AND 'Name'=gcp AND 'label' = region-us-east) OR ('label'=region-us-west)`. Then the Cohesity cluster will exclude the disks of the VMs that are of type, `Balanced persistent disks` with "`gcp`" as the disk name, and with label name as "`region-us-east`" or disks with "`region-us-west`" as the label name. **Note:** Exclusion of disks enabled at the object level will take precedence over the exclusion configured here. Enable **Exclude VMs with no disk** to exclude the VM with no disk from protection. 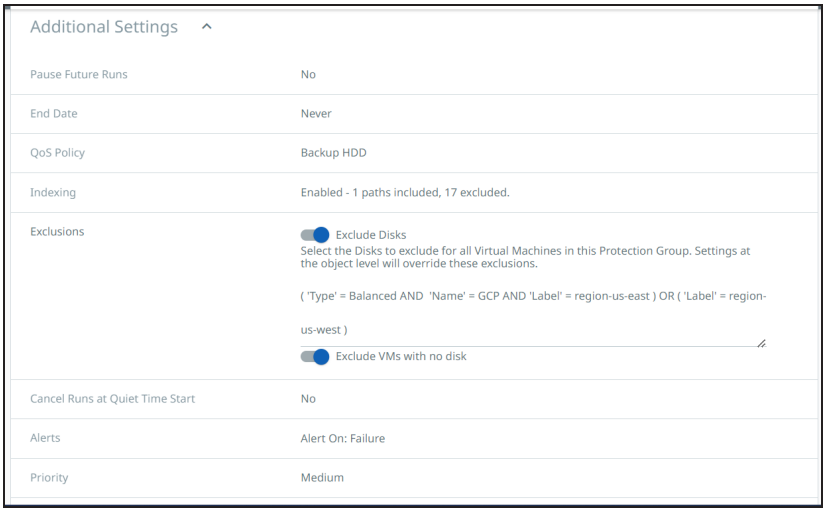
Cancel Run at Quite Start Time	*Available only if the selected policy has at least one quiet time period.* Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. If the option is toggled off, then after a protection run starts, it continues to execute even when a quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period.

Additional Settings	Description
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
Description	Enter a brief description about the Protection Group.

Manage GCP Recovery

This chapter includes the following topics:

Recover GCP VMs

You can recover GCP VMs from a snapshot created by a Protection Group.

1. In the cluster UI, select **Data Protection > Recoveries**.
2. Click **Recover** > **Virtual Machines** > **VMs** located at the top right of the page.
3. Search for a protected object (VM) or Protection Group that contains the snapshots to recover from. Enter characters of the VM name or Protection Group name and a list of items displays. You can optionally specify the wildcard character *. For more information, see Use Wildcard * in Search for Recovery. You can also optionally narrow the search results by specifying filter criteria, for example, you can filter the search results by a specific Protection Group. Click the **Add Filters** icon, specify the filter and click **Add**.

4. Select a Protection Group or a VM.

> **Note:**
> - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
> - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the Storage Domain and Source filters will automatically get locked.

5. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a...	Then...
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the **Protection** Run drop-down menu, select the snapshot from which you want to recover the GCP VM. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a...	Then...
VM	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. b. Select the snapshots to be recovered. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.

   From the **Recover To** options, select one of the following:

   - **Original Location**: Select this option to recover the VM their original location
   - **New Location**: Select this option to recover the VM to an alternate project, region and availability zone. Provide the following information:
     - **Source**: Select a registered Source to recover the VMs to.
     - **Project**: Select a Google Cloud Platform (GCP) project to recover the VMs to.
     - **Region**: Select a region in the project to recover the VMs to.
     - **Availability Zone**: Select an availability zone in GCP to recover the VMs.
     - **VPC**: Select a Virtual Private Cloud (VPC) to store the recovered VMs.
7. Under **Recovery Options**, edit the default settings. Click the **edit** icon to edit the settings:
   a. **Rename**: Add a Prefix and/or Suffix to the recovered GCP VMs.
   b. **Power State**: To ensure that the restored GCP VMs are functioning properly, the VMs are powered on by default after the restoration process. However, you can enable the **Power Off** option if you prefer to keep the VMs powered off once Cohesity verifies their functionality. If you are recovering VMs from a Storage Domain with Cloud Tier enabled, Cohesity recommends keeping the VMs powered off.
   c. **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
   d. **Task Name**: Change the default name of the recovery task.
8. Click **Recover**.

Cohesity begins to restore the selected GCP VMs.

Recover Files or Folders for GCP to the New Server

Prerequisites

- You must install the Cohesity agent on the target VM.
- Firewall port 50051 must be open on the target VM. It is required for transferring the files from the Cohesity cluster to the target VM. For more information, see Manage Firewall Ports.
- The target VM must be reachable with private IP.

> **Note:** You can perform files and folder recovery to an GCP VM Linux instance only if the GCP VM Linux instance is of x86_64 architecture.

Recover Files and Folders to New Server

Perform the following steps to recover files and folders to the new server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on Cohesity.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on Cohesity.

If you select..	Then..
Files and Folders	a. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. b. From the list of search results, select the required files or folders to recover. c. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: i. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. ii. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. iii. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. d. Go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display

If you select..	Then..
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. By default, the snapshot is recovered from the local cluster. 8. Go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Server Type	Volume Type	

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next: Recover Options**: If you select this option, then continue with the step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task completes, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
     Example:
     - The original name is "a:b" then the archived name is "a_COLON_b"
     - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

     If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
     Example:

- The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **New Server** to recover the files or folders to a new server.
6. **Source**: From the **Source** drop-down menu, select a registered server to which you want to recover the files and folders. You can also register a new source by clicking **Register Source**.
7. **Target**: From the Target (for VM and physical server) or Volume (for NAS) drop-down menu, select a VM or volume to which the files and files are to be recovered.

   > **Note:** For recovering files and folders of a VMware VM, ensure that VMware Tools are installed and running on the target VM. The target list displayed includes only VMs with VMware Tools installed and running.

8. **Recover to**: In the **Recover to** field, change the default location to which the files or folders are recovered. By default, the files and folders are recovered to the /**tmp**/ directory for VMs.

   For NAS, an example path on an alternate location could be /tmp/recover. You can recover encrypted files/folders to their original location only.

   Example for Windows alternate location:

   - D:\cohesity_restore\1.txt
   - D:\cohesity_restore\bar

   User Account Requirements for VM Protection may need to be met to recover files to a Windows VM.

   If recovering to a standalone Hyper-V host or cluster, you must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.
9. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

- C:\restore\1.txt
- C:\restore\cohesity<taskid>_0\1.txt
- C:\restore\bar
- C:\restore\cohesity<taskid>_1\bar

If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

- **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only files, then files will receive the new timestamps.
- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

10. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

BigQuery

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including BigQuery.

Cohesity supports both dataset and table level protection.

> **Note:** Table-level protection is an Early Access feature. Contact your Cohesity account team to enable the feature.

This chapter includes the following topics:

Plan and Prepare for BigQuery Data Protection

To protect BigQuery datasets and tables with the Cohesity cluster, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and

recovery operations. The following table lists the requirements are grouped into those needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Network Connectivity	backup/restore	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	backup/restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/`
Staging Bucket for Backup and Restore	Backup/Restore	You must have a staging bucket in Google Cloud Storage to enable backup and restore operations for BigQuery objects. • **Auto Create**: Cohesity can automatically create a staging bucket when you configure backup and recovery for BigQuery objects. • **Manual Create**: You can manually create one or more staging buckets in Google Cloud Storage. Use different staging buckets for protecting different objects or for recovery operations. Each bucket can be in the same or a different region as the object. > **Note:** Using a bucket in a different region than the object may incur additional Google Cloud Storage costs.

Create an IAM Custom Role in GCP

To enable BigQuery data protection, you must create an IAM custom role in Google Cloud and assign it to the service account used by the Cohesity cluster.

For detailed instructions, refer to the following topics in the Google Cloud documentation:

- Create a custom IAM role
- Assign a role to a service account

> **Note:** Before you begin creating IAM custom role, ensure that you have created a dedicated service account in Google Cloud for the Cohesity cluster.

Create Staging Bucket on Google Cloud Storage

To enable BigQuery data protection with the Cohesity cluster, you need a staging bucket in Google Cloud Storage. The staging bucket is used to temporarily store exported BigQuery objects during backup and recovery operations.

You can choose one of the following options when setting up the staging bucket:

- **Auto Create**: Cohesity can automatically create a staging bucket for you as part of the backup configuration workflow in the Cohesity cluster.
- **Create Manually**: Manually create and specify an existing Google Cloud Storage bucket to be used for staging.

> **Note:** Create one or more staging buckets in Google Cloud Storage. You can use different staging buckets for protecting different objects or for recovery operations. Each bucket can be in the same or a different region. Note that using a different region may incur additional Google Cloud Storage costs.

Manually Create a Staging Bucket

To manually create a staging bucket:

1. In the **Google Cloud** console, go to the Cloud Storage **Buckets** page.
2. Click **Create**.
3. On the **Create a bucket** page, provide the following information. After completing each section, click **Continue** to proceed:
    1. **Get started**: Enter a unique name for the staging bucket. Follow naming guidelines to ensure compatibility.
    2. **Choose where to store your data**: Select either **Region** or **Multi-region** as the location type.
    3. **Choose how to store your data**: Select **Set a default class**, then choose **Standard** as the default storage class.
    4. **Choose how to control access to objects**:
        - Enable **Enforce public access prevention**.
        - Select **Uniform as the access control model**.
    5. **Choose how to protect object data**:

- Under **Data protection**, enable **Soft delete policy (For data recovery)** and select **Use default retention duration**.
- Under **Data encryption**, select **Google-managed encryption key** as the encryption method.

6. Click **Create**.

Considerations for BigQuery Data Protection

Review and understand the following limitations and known issues before you protect BigQuery objects:

- **User permissions are not restored** — Google Cloud project users and their permissions are excluded from restores. The restored dataset owner will be the GCP service principal configured in the Cohesity cluster when you registered the Google Cloud as a source.
- **Daily export limit** — Google Cloud imposes a maximum data export limit of 50 TB per project, per day.
- **Tags are not restored** — Dataset and table tags are not included in backups or restores.
- **Unsupported data types** — A record with the data type: DATETIME in the tables is not supported for backup. This also includes DATETIME fields nested inside a RECORD-type field.

  The RANGE data type is also not supported.
- **Linked datasets** — Datasets linked from other projects are not discovered, backed up, or restored.
- **External datasets** — Datasets that reference external data sources outside BigQuery (for example, data in Cloud Storage or Spanner) are not discovered, backed up, or restored.
- **Schema requirements** — Tables without a defined schema or without any columns cannot be backed up. Tables must have a defined schema to be backed up. Empty tables are supported for protection and recovery if the table schema is present.
- **Access controls not supported** — Tables that use column-level access control or row-level security cannot be backed up.
- **Backup type limitation** — Only full backups are supported; incremental backups are not available for BigQuery objects.
- **Cross-region charges** — Cross-region backup and recovery operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.
- **Unsupported table types** — Views, materialized views, and table snapshots are not supported for backup and restore in this feature. External tables are supported.

Backup BigQuery Objects

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, it automatically discovers all projects that the service account can access, as well as all supported workloads such as BigQuery datasets and tables within those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab. Select the **BigQuery** tab to view all the discovered datasets or tables.

The **BigQuery** tab displays the entity hierarchy for the BigQuery environment, allowing you to navigate from the project level down to individual tables.

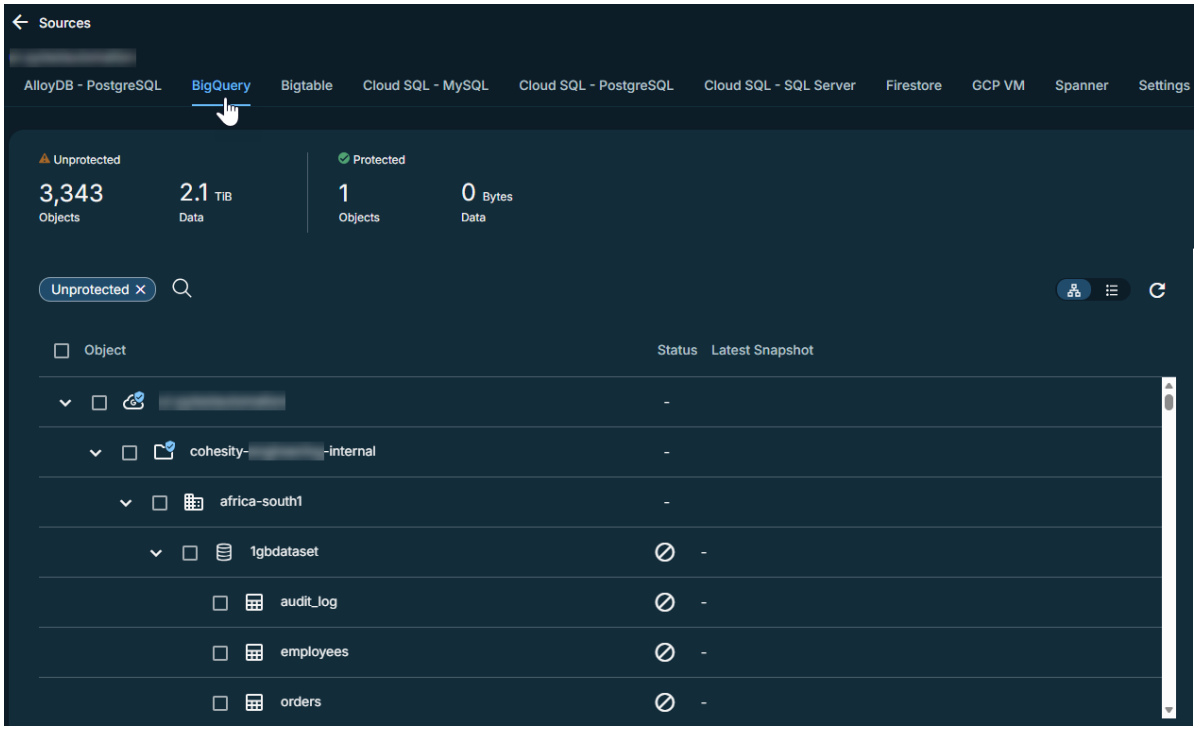

BigQuery Entity Hierarchy

The BigQuery entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
Service Account	The dedicated service account in Google Cloud assigned to Cohesity cluster for BigQuery protection.

Structure	Description
Project ID	The Google Cloud Project ID of the project (s) accessible with the service account.
Region	The Google Cloud region where the datasets reside.
Dataset	The dataset that can be selected for backup and recovery.
Table	Individual BigQuery tables displayed under a dataset.

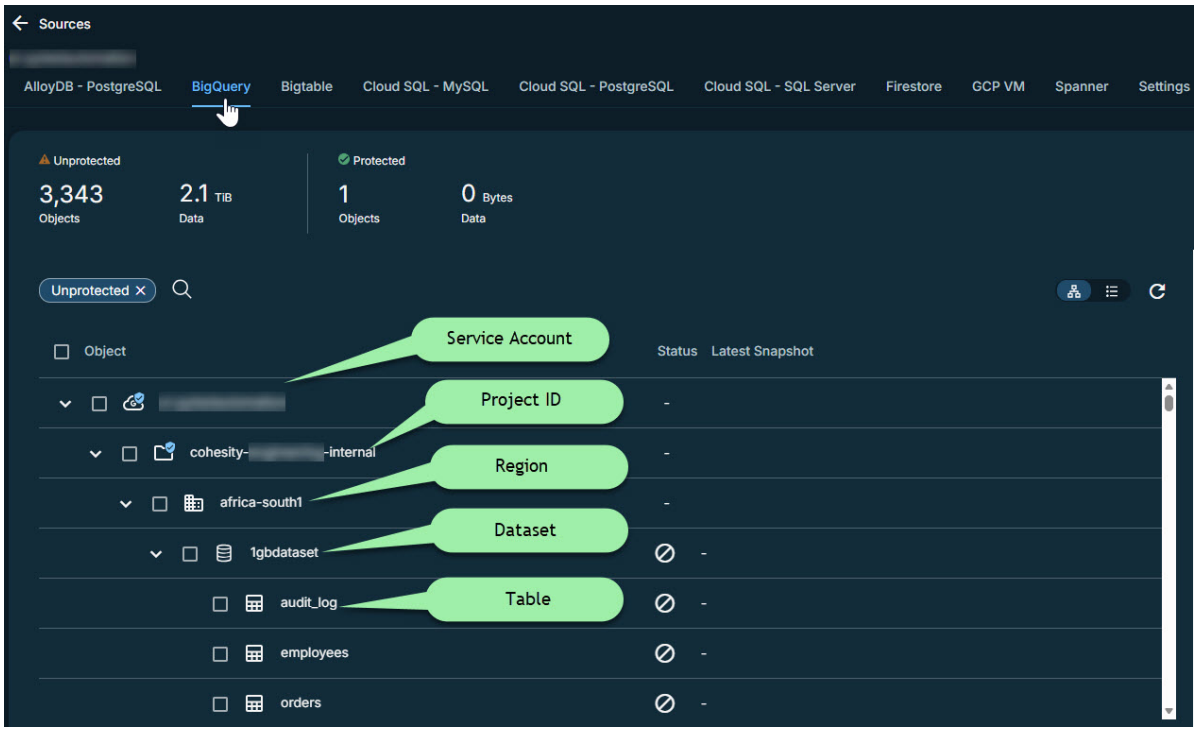

Backup Scope by BigQuery Hierarchy

Cohesity cluster allows you to back up BigQuery datasets at different levels of the hierarchy. You can choose to back up:

- **Individual dataset** – Back up a single dataset.
- Individual tabless - Backup up one or more selected tables within a dataset.

> **Note:** Table-level jobs run one parallel subtask per table; partial success is possible.

- **Multiple datasets within a region** – Back up selected datasets from the same region.
- **Region level** – Back up all datasets within a specific region.
- **Multiple datasets across regions** – Back up selected datasets from different regions.
- **Project level** – Back up all datasets in all regions within the specified Project ID.

BigQuery Backup Process

When you run a protection job in the Cohesity cluster to back up BigQuery datasets, the datasets specified in the Protection Group are first exported from BigQuery to a staging bucket in Google Cloud Storage. The Cohesity cluster then copies the data from the staging bucket to the cluster. After the backup completes successfully, the contents of the staging bucket are deleted to free up storage.

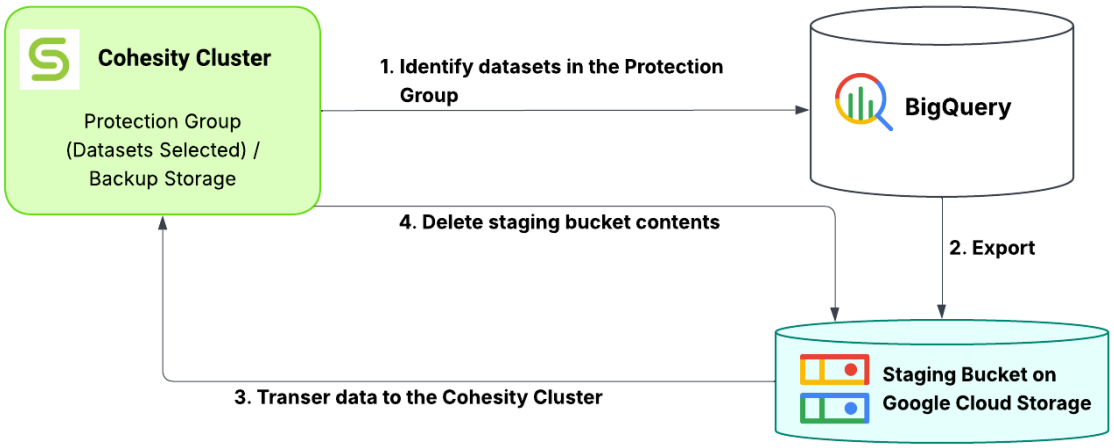

Supported Backup Types

Cohesity supports only full backups for BigQuery datasets or tables. Incremental backups are not supported.

Protect BigQuery

In the Cohesity cluster, Google Cloud BigQuery datasets and tables are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

> **Note:** Cohesity supports only full backups for BigQuery datasets and tables. Incremental backups are not supported.

Prerequisites

Before you protect BigQuery objects, ensure that a staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the backup configuration.

- If you choose to **Auto Create** the staging bucket: Cohesity automatically creates a staging bucket in the same region as your dataset(s) or tables when you start a backup operation.
- If you choose to **specify an existing bucket**: You can specify an existing bucket located in the same region as the datasets or tables being protected. Using a bucket in a different region is supported but may incur additional Google Cloud Storage charges.

> **Note:** After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

Protect BigQuery Objects

To protect BigQuery objects:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **BigQuery**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose objects (datasets or tables) using one of the following views:

> **Note:**
> - If you select a dataset, Cohesity protects the dataset as a single unit and does not create table-level restore points.
> - If you select individual tables, each selected table is protected independently.
> - If all tables in a dataset are selected, or if Auto Protect is enabled at the dataset level, Cohesity protects the dataset as a single unit and does not create table-level restore points.

   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the source, project, or region level.
   - **List View**: Select datasets or tables from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your objects. The bucket follows the naming format `cohe-<project-id >-<region-name>`.

     Where:
     - `project-id` is the Google Cloud project ID.
     - `region-name` is the Google Cloud region where the object reside.
   - **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for backups.

   After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.
7. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** This option is not applicable for BigQuery data protection. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

8. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

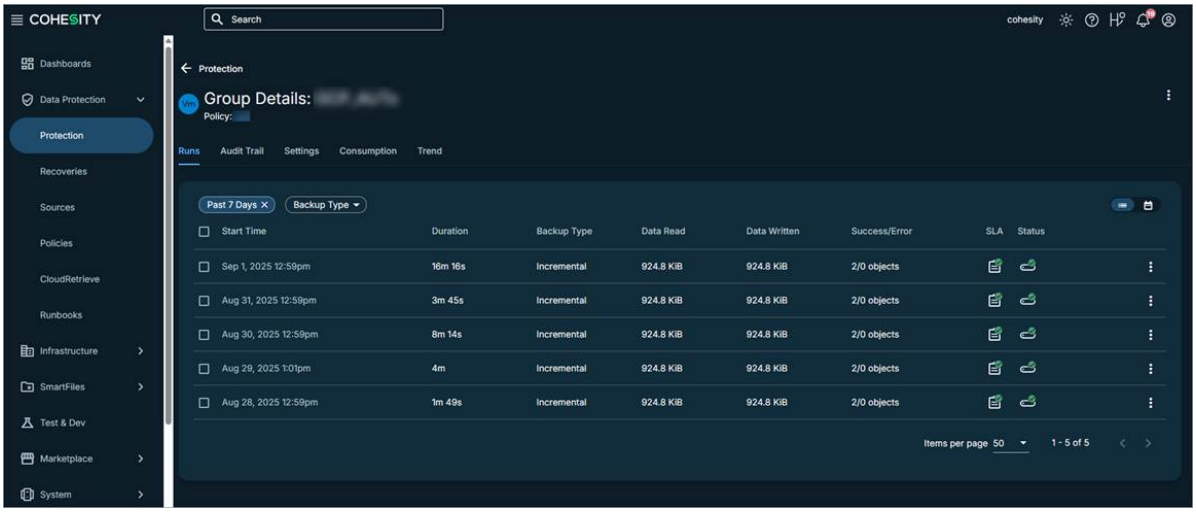

Manage BigQuery Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for BigQuery.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.

5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for BigQuery.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for BigQuery.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for BigQuery.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover BigQuery Objects

You can restore BigQuery datasets or tables either to their original location or to a new location. When you restore a dataset to a new location, you can restore it to a different region. When you restore a table to a new location, you can restore it to a different dataset.

Prerequisites

Before you restore BigQuery objects, ensure that a staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the restore process.

- If you choose to **Auto Create** the staging bucket: Cohesity automatically creates a staging bucket in the same region as your objects when you start the recovery.
- If you choose to **specify an existing bucket**: You can specify an existing bucket located in the same region as the recovery target. Using a bucket in a different region is supported but may incur additional Google Cloud Storage charges.

> **Note:** After each recovery job run, Cohesity deletes only the temporary data stored in the staging bucket. The staging bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.

Recover BigQuery Objects

To recover BigQuery objects:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **BigQuery**.
3. On the BigQuery recovery page, select one of the following options:
    - Dataset
    - Table
4. Search for the object by entering characters of the dataset name, or by specifying the wildcard character *****.
5. From the list of search results, select the database to recover.
    - By default, the latest snapshot of the database is selected for recovery.
    - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <database_name>** page appears.
    - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

To select a snapshot from the **List** mode:

a. On the **List** tab, snapshots available in the past month are displayed by default.

b. To view snapshots from a different range, select a date range from the **Date** drop-down list.

c. Select the snapshot(s) to recover.

d. By default, the snapshot is recovered from the local Cohesity cluster.

e. Click **Select Recovery Point** and then go to step 5.

To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

6. Click **Next: Recover Options**.
7. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the object to the same project from which it was backed up. **Note:** Restoring to the original location replaces any existing dataset / table with the same name.
New Location	Restore the object to an alternate location in a different GCP source. Select a registered GCP source, project, and region. To overwrite an existing object with the same name, enable **Overwrite existing dataset(s) / table(s)** and type **YES**. **Note:** Restoring to a new location also replaces any existing dataset / tables with the same name in the selected region.

8. In **Restore Staging Bucket** section, choose one of the following options:

- **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your object for the recovery operation. The bucket follows the naming format: `cohe-<project-id>-<region-name>`.

  Where:

  - `project-id` is the Google Cloud project ID.
  - `region-name` is the Google Cloud region where the datasets reside.
- **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for recovery.

Cohesity deletes only the temporary data stored in the staging bucket during each recovery operation. The bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.

9. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the dataset/table if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

10. Click **Recover**.

Bigtable

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including Bigtable.

> **Important:**
> The backup and recovery of Google Cloud NoSQL database for Bigtable is supported only on Cohesity Cloud Edition on GCP (formerly known as GCP NGCE cluster). This is an Early Access feature, contact Cohesity to enable the feature.

This chapter includes the following topics:

Plan and Prepare for Bigtable Data Protection

To protect Bigtable with the Cohesity cluster, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and recovery operations. The following table lists the requirements and are grouped as needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Allowlisted Domains	Source registration, backup/restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/` • `https://bigtable.googleapis.com/` • `https://bigtableadmin.googleapis.com/` • `https://dataflow.googleapis.com/`
Bigtable users	Backup/Restore	Bigtable access is controlled through IAM (service account). The Cohesity service account must have the required Bigtable permissions on the project/instance/table.
Dataflow	Backup operations	Dataflow is a Google Cloud service for running data processing pipelines. Cohesity requires Individual Permission for Bigtable Data Protection to submit/monitor jobs used for Bigtable export/import.

Individual Permission for Bigtable Data Protection

The following table lists the permissions and privileges that must be included in a custom role. Assign this role to the service account to enable Bigtable data protection.

Permission	Purpose
`bigtable.clusters.list`	Allows you to view a list of all Bigtable clusters in your project.
`bigtable.clusters.get`	Enables you to retrieve detailed information about a specific Bigtable cluster.
`bigtable.instances.get`	Allows you to access detailed information about a specific Bigtable instance.

Permission	Purpose
`bigtable.instances.list`	Grants the ability to list all Bigtable instances within a cluster.
`bigtable.tables.list`	Allows you to access detailed information about specific Bigtable tables.
`bigtable.tables.get`	Grants the ability to list all Bigtable tables within a cluster.
`bigtable.tables.create`	Grants the ability to create tables within a Bigtable cluster.
`bigtable.tables.mutateRows`	Allows you to restore/import operations that write rows to Bigtable.
`bigtable.tables.update`	Allow you to recreate the table settings if your workflow changes table configuration.
`bigtable.tables.delete`	Allow you to restore the workflows that drop-and-recreate tables.
`bigtable.appProfiles.get`	Retrieves metadata for a specified Bigtable app profile ID provided during registration.
`dataflow.jobs.create`	Allows you to submit Dataflow jobs.
`dataflow.jobs.get`	Allows you to monitor the Dataflow job progress and display errors, if any.

Create Staging Bucket on Google Cloud Storage

To enable Bigtable data protection with the Cohesity cluster, you need a staging bucket in Google Cloud Storage. The staging bucket is used to temporarily store exported Bigtable during backup and recovery operations.

You can choose one of the following options when setting up the staging bucket:

- **Auto Create**: Cohesity can automatically create a staging bucket for you as part of the backup configuration workflow in the Cohesity cluster.
- **Create Manually**: Manually create and specify an existing Google Cloud Storage bucket to be used for staging.

> **Note:** Create one or more staging buckets in Google Cloud Storage. You can use different staging buckets for protecting different datasets or for recovery operations. Each bucket can be in the same or a different region as the dataset. Note that using a different region may incur additional Google Cloud Storage costs.

Manually Create a Staging Bucket

To manually create a staging bucket:

1. In the **Google Cloud** console, go to the Cloud Storage **Buckets** page.
2. Click **Create**.
3. On the **Create a bucket** page, provide the following information. After completing each section, click **Continue** to proceed:
    1. **Get started**: Enter a unique name for the staging bucket. Follow naming guidelines to ensure compatibility.
    2. **Choose where to store your data**: Select either **Region** or **Multi-region** as the location type.
    3. **Choose how to store your data**: Select **Set a default class**, then choose **Standard** as the default storage class.
    4. **Choose how to control access to objects**:
        - Enable **Enforce public access prevention**.
        - Select **Uniform as the access control model**.
    5. **Choose how to protect object data**:
        - Under **Data protection**, enable **Soft delete policy (For data recovery)** and select **Use default retention duration**.
        - Under **Data encryption**, select **Google-managed encryption key** as the encryption method.
    6. Click **Create**.

Considerations for Bigtable Data Protection

Review and understand the following limitations and known issues before you protect Bigtable:

- **Cluster location** : Before you can use a regional endpoint, you must create an instance. By default, single cluster is created.
- **Application profile**: After you set a regional endpoint, Cohesity recommends to use an app profile that's configured to route requests to a cluster in the endpoint's region.

- **Replication**: To avoid replicating data to another region, create clusters only in the intended region.
- **Availability**: A regional endpoint prevents automatic failovers as it confines all requests to a single region.
  For more information, see Bigtable service level agreement (SLA).
- **Cross-region**: Cross-region backup and restore operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.

Backup Google Cloud Bigtable

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, Cohesity scans the Google Cloud projects accessible to the service account and discovers Cloud Bigtable resources in those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab.

Select the **Bigtable** tab to view all the tables discovered.

The Bigtable tab displays the entity hierarchy, allowing you to navigate from the service account level down to individual tables.

Bigtable Entity Hierarchy

The Bigtable entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
Service Account	The dedicated service account in Google Cloud assigned to Cohesity cluster for Bigtable protection.
Project ID	The Google Cloud Project ID of the project (s) accessible with the service account.
Table	The individual Google Cloud for Bigtable that can be backed up or restored.

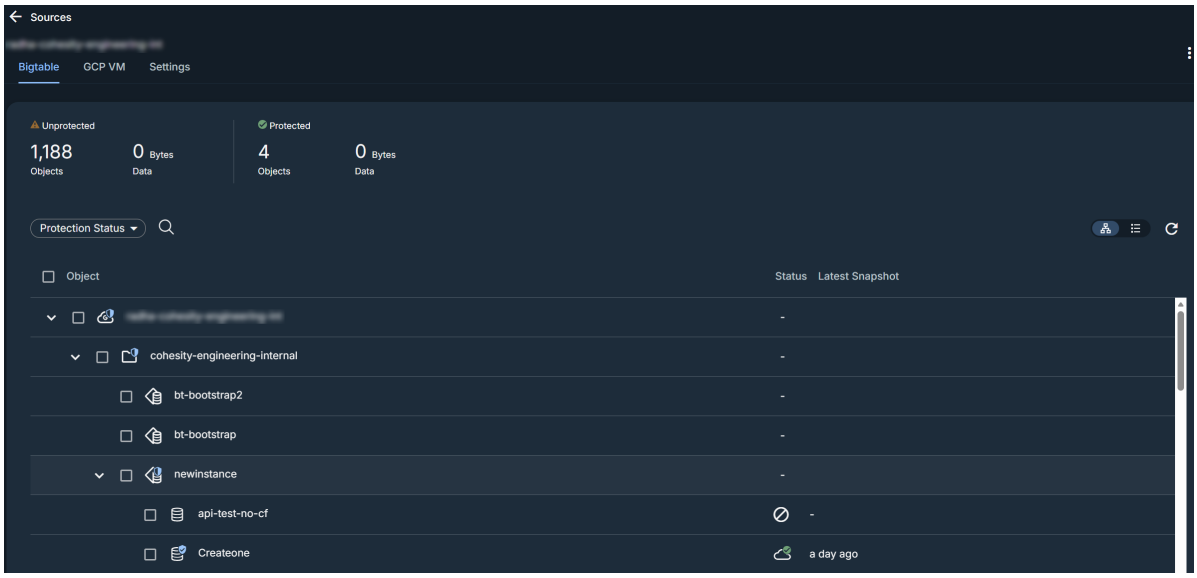

Backup Scope by Bigtable Hierarchy

Cohesity cluster allows you to back up Bigtable at different levels of the hierarchy. You can choose to back up:

- **Individual tables** – Back up a single or multiple tables.
- **Multiple tables across regions**: Back up selected tables from different regions.
- **Project level**: Back up all instances in all regions within the specified Project ID.

Bigtable Backup Process

When you run a protection job for Bigtable, the Cohesity cluster performs a full backup specified in the Protection Group. The Cohesity cluster connects to the selected instance using the specified credentials during source registration. Cohesity triggers a Dataflow job to read data from Bigtable tables and transfers to staging bucket.

Further, the data is transferred from staging bucket to the Cohesity cluster.

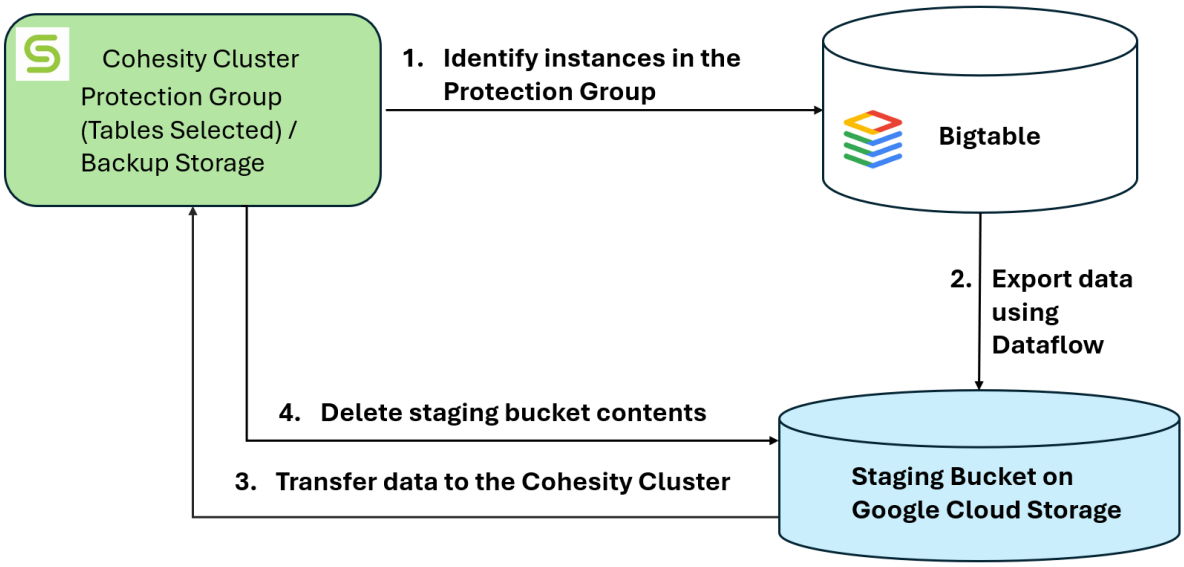

Supported Backup Types

Cohesity supports only full backups for Bigtable datasets. Incremental backups are not supported.

Protect Bigtable

In the Cohesity cluster, Bigtable is protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

> **Note:** Cohesity supports only full backups for Bigtable. Incremental backups are not supported.

Prerequisites

Before you protect Bigtable, ensure that a staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the backup configuration.

- If you choose to **Auto Create** the staging bucket: Cohesity automatically creates a staging bucket in the same region as your dataset(s) when you start a backup operation.
- If you choose to **specify an existing bucket**: You can specify an existing bucket located in the same region as the datasets being protected. Using a bucket in a different region is supported but may incur additional Google Cloud Storage charges.

> **Note:** After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

Protect Bigtable

To protect Bigtable:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **Bigtable**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose datasets using one of the following views:
   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the instance level.
   - **List View**: Select datasets from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your dataset(s). The bucket follows the naming format `cohe-<project-id >-<region-name>`.

     Where:
     - `project-id` is the Google Cloud project ID.
     - `region-name` is the Google Cloud region where the datasets reside.
   - **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for backups.

   After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.

7. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** This option is not applicable for Bigtable data protection. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

8. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Max Dataflow Workers	Specify the maximum number of Dataflow worker instances that a Dataflow job can use. Google Cloud Dataflow auto-scales the number of worker instances based on worker CPU utilization. The Max Dataflow Workers value acts as an upper limit where Dataflow service will not allocate more worker instances than the specified value, even if additional capacity could improve job performance.

Additional Settings	Description
Dataflow Prime	Dataflow Prime is the Dataflow service feature which enhances the backup process by dynamically allocating resources such as CPU. Disable by default. If enabled, the resources are allocated dynamically. The Dataflow Prime allows GCP to allocate more CPUs or reduce them based on the job size (backup or restore). Note: Support and pricing are governed by Google Cloud Dataflow. Enabling attracts additional cost.
Application Profile	Optional. An application profile defines how Bigtable handles requests from a Cohesity, including how requests are routed to Bigtable clusters. Before registering a Bigtable workload, ensure that a valid application profile already exists in Google Cloud for the target Bigtable instance. Cohesity requires the Application Profile ID to be explicitly provided during registration. Cohesity does not list, discover, or automatically select Bigtable application profiles, including the default profile. For more information on Application Profile, see App profile overview.
Pause Future Runs	Enable Pause Future Run only if you do not want to schedule future runs of the Protection Group.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. • **Backup Auto**: The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the Backup SSD policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

Manage Bigtable Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Bigtable.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.

4. On the **Edit Protection** page, click in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Bigtable.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Bigtable.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.

3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Bigtable.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Bigtable

You can restore BigQuery datasets either to their original GCP source or to a different GCP source and region.

Prerequisites

Before you restore Bigtable, ensure that a staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the restore process.

Recover Bigtable

To recover Bigtable:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **Bigtable**.
3. Search for the dataset by entering characters of the dataset name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
    - By default, the latest snapshot of the database is selected for recovery.
    - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <table_name>** page appears.
    - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

      To select a snapshot from the **List** mode:

        a. On the **List** tab, snapshots available in the past month are displayed by default.
        b. To view snapshots from a different range, select a date range from the **Date** drop-down list.
        c. Select the snapshot(s) to recover.
        d. By default, the snapshot is recovered from the local Cohesity cluster.
        e. Click **Select Recovery Point** and then go to step 5.

      To select a snapshot from the **Timeline** mode:

        a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
        b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the dataset to the same project from which it was backed up. **Note:** Restoring to the original location replaces any existing dataset with the same name.
New Location	Restore the dataset to an alternate location in a different GCP source. Select a registered GCP source, project, and region. To overwrite an existing dataset with the same name, enable **Overwrite existing dataset(s)** and type **YES**. **Note:** Restoring to a new location also replaces any existing dataset with the same name in the selected region.

7. In **Restore Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your dataset(s) for the recovery operation. The bucket follows the naming format: `cohe-<project-id>-<region-name>`.

     Where:

     - `project-id` is the Google Cloud project ID.
     - `region-name` is the Google Cloud region where the datasets reside.
   - **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for recovery.

   Cohesity deletes only the temporary data stored in the staging bucket during each recovery operation. The bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.
8. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the dataset if you want to rename it after recovery.
Max Dataflow Workers	Specify the maximum number of Dataflow worker instances that a Dataflow job can use. Google Cloud Dataflow auto-scales the number of worker instances based on worker CPU utilization. The Max Dataflow Workers value acts as an upper limit where Dataflow service will not allocate more worker instances than the specified value, even if additional capacity could improve job performance.
Dataflow Prime	Dataflow Prime is the Dataflow service feature which enhances the backup process by dynamically allocating resources such as CPU. Disable by default. If enabled, the resources are allocated dynamically. The Dataflow Prime allows GCP to allocate more CPUs or reduce them based on the job size (backup or restore). **Note:** Support and pricing are governed by Google Cloud Dataflow. Enabling attracts additional cost.
Task Name	Update the default name of the recovery task.

9. Click **Recover**.

Spanner

Cohesity Platform natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including Spanner.

> **Important:** The backup and recovery of Google Cloud Spanner is supported only on Cohesity Cloud Edition on GCP.

This chapter includes the following topics:

Plan and Prepare for Spanner Data Protection

To protect Spanner databases with Cohesity, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and recovery operations. The following table lists the requirements and are grouped into those needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Network Connectivity	Backup/Restore	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	Backup/Restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/` • `https://spanner.googleapis.com` • `https://dataflow.googleapis.com` • `https://oauth2.googleapis.com`
Staging Bucket for Backup and Restore	Backup/Restore	You must have a staging bucket in Google Cloud Storage to enable backup and restore operations for Spanner databases. • **Auto Create**: Cohesity can automatically create a staging bucket when you configure protection for a Spanner database. • **Manual Create**: You can manually create one or more staging buckets in Google Cloud Storage. Use different staging buckets for protecting different databases or for recovery operations. Each bucket can be in the same or a different region as the database. **Note:** Using a bucket in a different region than the database may incur additional Google Cloud Storage costs.

Create Staging Bucket on Google Cloud Storage

To enable Spanner data protection with the Cohesity cluster, you need a staging bucket in Google Cloud Storage. The staging bucket is used to temporarily store exported Spanner data during backup and recovery operations.

You can choose one of the following options when setting up the staging bucket:

- **Auto Create**: Cohesity can automatically create a staging bucket for you as part of the backup configuration workflow in the Cohesity cluster.

- **Create Manually**: Manually create and specify an existing Google Cloud Storage bucket to be used for staging.

> **Note:**
> - You can create one or more staging buckets in Google Cloud Storage. Different staging buckets can be used for protecting different databases or for recovery operations.
> - A staging bucket can be in the same or a different region as the database. However, using a bucket in a different region may incur additional Google Cloud Storage costs.
> - By default, backup and restore operations fail if the staging bucket is not created in the default leader region of the database configuration. Contact Cohesity Support if you want to enable cross-region transfers.

Manually Create a Staging Bucket

To manually create a staging bucket:

1. In the **Google Cloud** console, go to the Cloud Storage **Buckets** page.
2. Click **Create**.
3. On the **Create a bucket** page, provide the following information. After completing each section, click **Continue** to proceed:
    1. **Get started**: Enter a unique name for the staging bucket. Follow naming guidelines to ensure compatibility.
    2. **Choose where to store your data**: Select **Region** as the location type.

        > **Note:** Do not select Multi-region. Multi-regional buckets are not supported for Spanner backup or recovery operations. If you need to use an existing multi-regional bucket, contact Cohesity Support for guidance.

    3. **Choose how to store your data**: Select **Set a default class**, then choose **Standard** as the default storage class.
    4. **Choose how to control access to objects**:
        - Enable **Enforce public access prevention**.
        - Select **Uniform as the access control model**.
    5. **Choose how to protect object data**:

- Under **Data protection**, enable **Soft delete policy (For data recovery)** and select **Use default retention duration**.

  > **Note:** To avoid unnecessary storage charges, disable the Soft delete policy on staging buckets used for backup and restore operations. Cohesity automatically deletes all staged data after each operation. If soft delete is enabled, this temporary data will be retained for the full soft-delete retention period resulting in additional costs.

- Under **Data encryption**, select **Google-managed encryption key** as the encryption method.

6. Click **Create**.

Considerations for Spanner Data Protection

Review and understand the following limitations and known issues before you protect Spanner databases with Cohesity .

- **Full Backups Only**: Only full backups are supported. Incremental backups and point-in-time recovery (PITR) are not supported.
- **Database-Level Operations**: Backups and recovery operations are performed at the database level. Table-level operations are not supported.
- **Data Format**: Backup data is exported in Avro format. Some Spanner-specific data types or extensions may require adjustments before being restored or imported into other SQL databases.
- **Edition Compatibility**: Backups and restores are supported across all Spanner editions (Standard, Enterprise, and Enterprise+). Cross-edition recovery is allowed but may be limited to features or datatypes common to both editions. For best results, restore to the same edition whenever possible.
- **Database-Level Recovery**: Recovery of complete Spanner instances is not supported. You can restore only individual databases.
- **Overwrite Behavior**: Restoring over an existing database with same name in an instance requires deletion of existing conflicting database. If drop protection is enabled , deletion will fail, and manual intervention is required.
- **Performance Optimization**: Improve backup performance by increasing the number of Dataflow workers and enabling Boost. For large databases, horizontal scaling (adding more workers) is typically more effective than vertical scaling.
- **Plan for Cloud Dataflow quotas**: Cohesity uses Cloud Dataflow jobs during backup and recovery. Google Cloud projects run up to 25 concurrent Dataflow jobs by default. If your environment triggers more than 25 concurrent backup/recovery operations

and they start Dataflow jobs, some jobs can be rejected and the corresponding operations can fail. To avoid this, reduce concurrent operations or request a higher Dataflow job quota in Google Cloud.
For more information, see Quotas and limits.

- **Cross-region charges**: Cross-region backup and recovery operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.

Backup Spanner Databases

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, it automatically discovers all projects that the service account can access, as well as all supported workloads such as Spanner databases within those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab. Select the **Spanner** tab to view all the discovered databases.

The **Spanner** tab displays the entity hierarchy for the Spanner environment, allowing you to navigate from the project level down to individual databases.

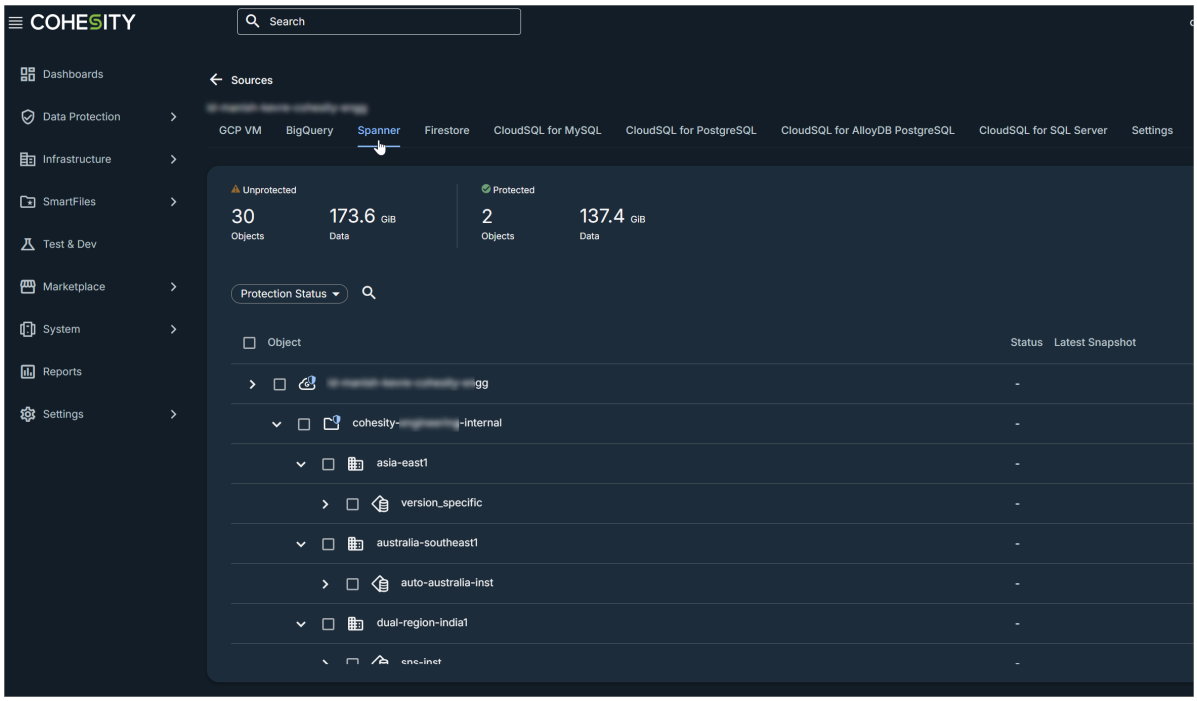

Spanner Entity Hierarchy

The following hierarchy describes how Spanner entities are organized and displayed in the Cohesity cluster:

Structure	Description
Service Account	The Google Cloud service account configured for Cohesity to perform Spanner backup and recovery operations.
Project ID	The Google Cloud Project ID that the service account has permissions to access.
Region	The Google Cloud region or multi-region configuration where the Spanner instance is hosted.
Instance	The Spanner instance available for backup and recovery.
Dataset	The individual Spanner database that can be backed up or restored. Cohesity supports both Spanner-supported dialects: Google Standard SQL and PostgreSQL.

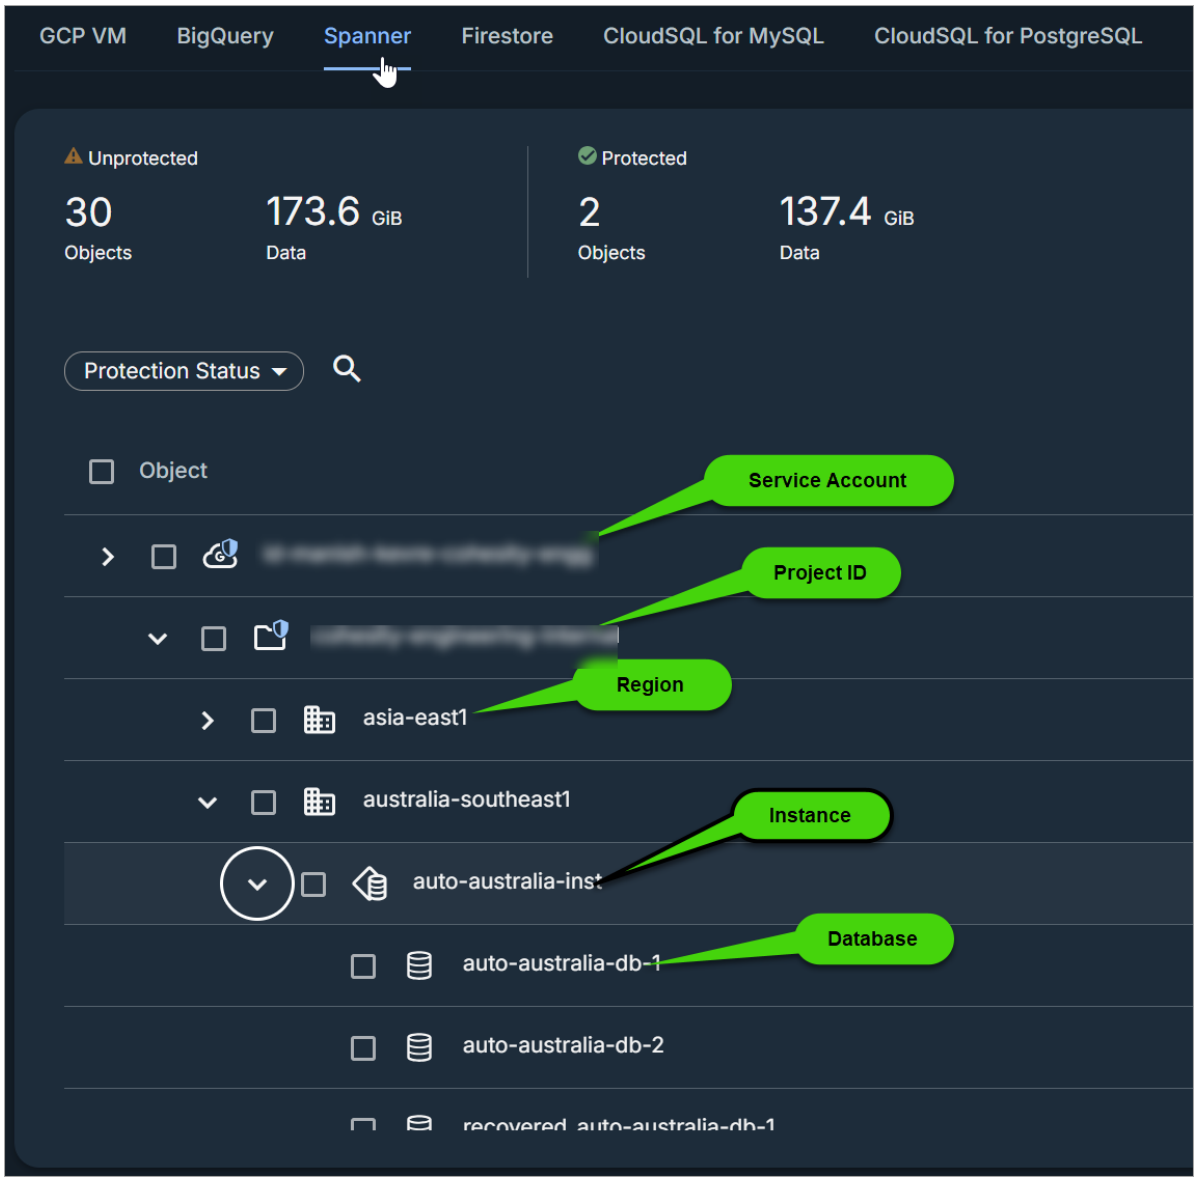

Backup Scope by Spanner Hierarchy

The Cohesity cluster allows you to back up Spanner databases at different levels within the hierarchy. You can choose to back up:

- **Individual Database**: Back up a single selected database.
- **Multiple Databases within an Instance**: Back up selected databases within the same Spanner instance.
- **Multiple databases across instances**: Back up selected databases from different instances, as long as all instances are located in the same region or in the **default leader region**.
- **Instance Level**: Back up all databases within the selected Spanner instance.

Spanner Backup Process

When you run a protection job in the Cohesity cluster, one or more selected Spanner databases are exported to a Google Cloud Storage staging bucket. The staging bucket temporarily stores the exported data before it is ingested into the Cohesity cluster. The Cohesity cluster verifies that the selected bucket meets configuration requirements and ensures that the backup runs efficiently.

During the process, the Spanner databases are securely exported to the staging bucket as full backups. After the export completes, the Cohesity cluster automatically copies the data from the staging bucket to the cluster for retention. Once ingestion is successful, the Cohesity cluster applies the defined retention policies. The contents of the staging bucket are then automatically deleted to free up storage and reduce costs.

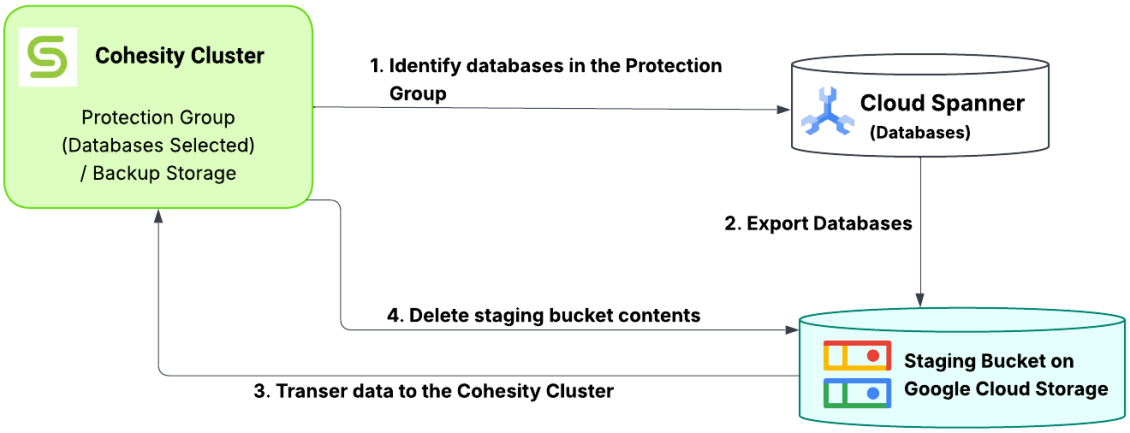

Supported Backup Types

Cohesity supports only full backups for Spanner databases. Incremental backups are not supported.

Protect Spanner

In the Cohesity cluster, Spanner database are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

> **Note:** Cohesity supports only full backups for Spanner databases. Incremental backups are not supported.

Prerequisites

Before you protect Spanner databases, ensure that a staging bucket exists in Google Cloud Storage (GCS) for storing exported data during backups. You can:

- **Auto Create**: Allow Cohesity to automatically create a temporary staging bucket during protection configuration.
  - The staging bucket is created in the same region as the database(s).
  - For multi-region deployments, the bucket is created in the **default leader region**.
- **Manually Create**: Create one or more staging buckets in GCS. You can use different buckets for protecting different databases or for recovery operations.

> **Note:**
> - Staging buckets created in a different region than the Spanner database will cause backup and restore operations to fail by default.
> - Cohesity supports using a staging bucket only in the default leader region of the database.
> - If you need to use a staging bucket in a different region, contact Cohesity Support to enable cross-region transfers.
>
> For more information, refer to Dual-region configurations or Multi-region configurations sections on Google Cloud Spanner documentation.

Protect Spanner Databases

To protect Spanner databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **Spanner**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose databases using one of the following views:
   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the source, project , region, and instance level.
   - **List View**: Select databases from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. In the **Staging Bucket** section, choose one of the following options:

- **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your database(s). For multi-region deployments, the bucket is created in the default leader region. The bucket follows the naming format `cohe-<project-id >-<region-name>`.

  Where:

  - `project-id` is the Google Cloud project ID.
  - `region-name` is the Google Cloud region where the database **default leader region** resides.
- **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for backups.

6. In the Network Configurations section, provide the following Google Cloud VPC details:
   - **VPC Network**: Enter the name of the Google Cloud VPC network to be used for backup operations. Ensure that the network has connectivity to both the Spanner instance and the Google Cloud Storage staging bucket. You can also specify default if the default VPC is configured in your project.
   - **VPC Subnet**: Specify the subnet within the selected VPC network that must be used for all backup activities. Subnets in Google Cloud are region-specific, so ensure that the subnet you select corresponds to the region of your Spanner instance. If your project uses default networking, you can also specify default as the subnet.
7. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.
8. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.

Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** This option is not applicable for Spanner data protection. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

9. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Max Dataflow Workers	Specify the maximum number of Dataflow workers that can be created during backup operations. This value sets the upper limit for autoscaling. Dataflow increases the number of workers based on workload, but it will not create more workers than the limit you set. A higher limit can improve performance but may increase cost and requires sufficient GCP quotas. **Default**: 10 **Recommendations**: • Keep 10 for most environments. • Increase for large databases or tighter backup windows. • Decrease if you need to limit resource usage or have quota constraints.
Data Boost	Enabled by default, Data Boost is a fully managed, serverless service provided by Google cloud that provides independent compute resources. For more information, see Data Boost on Google Cloud Documentation.
Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.

Additional Settings	Description
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

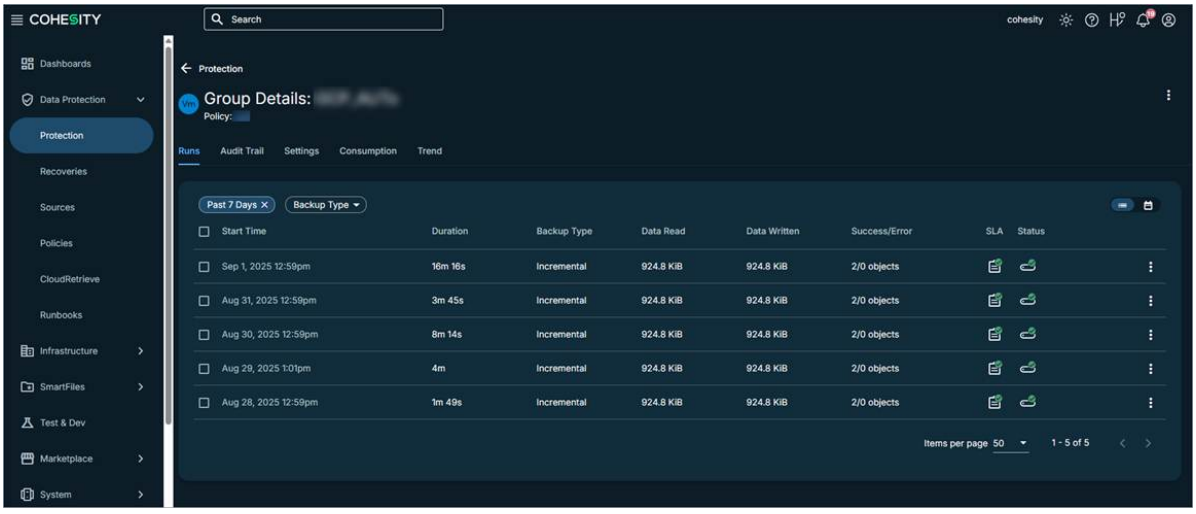

Manage Spanner Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Spanner .
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Spanner .
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Spanner .
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Spanner .
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
    - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Spanner Databases

You can recover Spanner databases either to their original Google Cloud source or to a different source, instance and region.

Spanner Recovery Process

When a recovery job runs on the Cohesity cluster, the selected Spanner database backups are exported to a Google Cloud Storage staging bucket. The Cohesity cluster then recreates the database in the target environment and transfers the exported data to it.

After the data transfer and database creation are complete, the Cohesity cluster triggers a Spanner import job. Google Cloud Spanner automatically ingests the database data from the staging bucket into the Spanner deployment.

Once the import is successful, the staged backup data is automatically deleted from the bucket to free up storage and minimize costs.

> **Note:**
> Under normal conditions, staged backup data in the staging bucket is automatically cleaned up after the import completes. However, in certain failure scenarios, the cleanup may not occur. In such cases, you should manually verify and delete any remaining staged data.
>
> Cleanup may fail in scenarios such as:
>
> - **Import job submission failure during recovery**. Staged data may remain if the import job could not be triggered.
> - **Duplicate subtask–related errors during backup or restore**. These errors may interrupt cleanup, leaving temporary data in the staging bucket.

Prerequisites

Before you recover Spanner databases, ensure that a staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the restore process.

- If you choose to **Auto Create** the staging bucket: Cohesity automatically creates a staging bucket in the same region as your database(s) when you start the recovery. For multi-region deployments, the bucket is created in the **default leader region**.
- The recovery operations may take longer if the target instance has lower compute capacity or is already running active workloads. Smaller instances may process data more slowly during recovery.
- If you choose to **specify an existing bucket**: You can specify an existing bucket located in the same region as the recovery target. Using a staging bucket in a different region than the target instance's default leader region is not supported by default. Contact Cohesity Support if cross-region usage is required.

> **Note:** After each recovery job run, Cohesity deletes only the temporary data stored in the staging bucket. The staging bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges .

Recover Spanner Databases

To recover Spanner databases:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **Spanner**.

3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
   - By default, the latest snapshot of the database is selected for recovery.
   - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <database_name>** page appears.
   - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

     To select a snapshot from the **List** mode:

     a. On the **List** tab, snapshots available in the past month are displayed by default.
     b. To view snapshots from a different range, select a date range from the **Date** drop-down list.
     c. Select the snapshot(s) to recover.
     d. By default, the snapshot is recovered from the local Cohesity cluster.
     e. Click **Select Recovery Point** and then go to step 5.

     To select a snapshot from the **Timeline** mode:

     a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
     b. Select the snapshot to be recovered by clicking the corresponding blue dot.
     c. By default, the snapshot is recovered from the local Cohesity cluster.
     d. Click **Select Recovery Point** and then go to step 5.
5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the database to the same project from which it was backed up. **Note:** • Restoring to the original location replaces any existing database with the same name. • If you need to restore to the same instance without removing the existing database, use the **Rename** option in the Recovery Options section to specify a new database name.
New Location	Restore the database to an alternate location in a different GCP source. Select a registered GCP source and Spanner instance. To overwrite an existing database with the same name, enable **Overwrite existing database(s)** and type **YES**. **Note:** Restoring to a new location also replaces any existing database with the same name in the selected region.

7. In **Restore Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your database(s) for the recovery operation. For multi-region deployments, the bucket is created in the **default leader region**. The bucket follows the naming format: `cohe-<project-id>-<region-name>`.

     Where:
     - `project-id` is the Google Cloud project ID.
     - `region-name` is the Google Cloud region where the databases reside.
   - **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for recovery.

   Cohesity deletes only the temporary data stored in the staging bucket during each recovery operation. The bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.
8. In the **Network Configurations** section, provide the following Google Cloud VPC details:
   - **VPC Network**: Enter the name of the Google Cloud VPC network to be used for the restore operation. Ensure that the network has connectivity to both the target Spanner instance and the Google Cloud Storage staging bucket.

- **VPC Subnet**: Specify the subnet associated with the region of the target Spanner instance configuration within the selected VPC network. This subnet will be used for the current restore operation.

9. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to recover it to the same instance without deleting an existing database with the same name.
Max Dataflow Workers	Specify the maximum number of Dataflow workers to use for recovery operations. More workers improve performance but may increase costs and require sufficient GCP quotas. **Default**: 10 **Recommendations**: • Keep 10 for most environments. • Increase for large databases or tighter backup windows. • Decrease if you need to limit resource usage or have quota constraints.
Task Name	Update the default name of the recovery task.

10. Click **Recover**.

Firestore

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including Firestore.

> **Important:** The backup and recovery of Google Cloud Firestore is supported only on Cohesity Cloud Edition on GCP.

This chapter includes the following topics:

Plan and Prepare for Firestore Data Protection

To protect Firestore databases with the Cohesity cluster, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and recovery operations. The following table lists the requirements are grouped into those needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Network Connectivity	Backup/Restore	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	Backup/Restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/`
Staging Bucket for Backup and Restore	Backup/Restore	You must have a staging bucket in Google Cloud Storage to enable backup and restore operations for Firestore databases. • **Auto Create**: Cohesity can automatically create a staging bucket when you configure backup and recovery for Firestore databases. • **Manual Create**: You can manually create one or more staging buckets in Google Cloud Storage. Use different staging buckets for protecting different databases or for recovery operations. Each bucket can be in the same or a different region as the database. **Note:** Using a bucket in a different region than the database may incur additional Google Cloud Storage costs.

Create Staging Bucket on Google Cloud Storage

To enable Firestore data protection with the Cohesity cluster, you need a staging bucket in Google Cloud Storage. The staging bucket is used to temporarily store exported Firestore databases during backup and recovery operations.

You can choose one of the following options when setting up the staging bucket:

- **Auto Create**: Cohesity can automatically create a staging bucket for you as part of the backup configuration workflow in the Cohesity cluster.
- **Create Manually**: Manually create and specify an existing Google Cloud Storage bucket to be used for staging.

> **Note:** Create one or more staging buckets in Google Cloud Storage. You can use different staging buckets for protecting different databases or for recovery operations. Each bucket can be in the same or a different region as the database. Note that using a different region may incur additional Google Cloud Storage costs.

Manually Create a Staging Bucket

To manually create a staging bucket:

1. In the **Google Cloud** console, go to the Cloud Storage **Buckets** page.
2. Click **Create**.
3. On the **Create a bucket** page, provide the following information. After completing each section, click **Continue** to proceed:
    1. **Get started**: Enter a unique name for the staging bucket. Follow naming guidelines to ensure compatibility.
    2. **Choose where to store your data**: Select either **Region** or **Multi-region** as the location type.
    3. **Choose how to store your data**: Select **Set a default class**, then choose **Standard** as the default storage class.
    4. **Choose how to control access to objects**:
        - Enable **Enforce public access prevention**.
        - Select **Uniform as the access control model**.
    5. **Choose how to protect object data**:
        - Under **Data protection**, enable **Soft delete policy (For data recovery)** and select **Use default retention duration**.
        - Under **Data encryption**, select **Google-managed encryption key** as the encryption method.
    6. Click **Create**.

Considerations for Firestore Data Protection

Review and understand the following limitations and known issues before you protect Google Cloud Firestore databases.

- **Database-Level Protection Only**: Backups and restores are performed at the database level. Individual collections cannot be selected for backup or restore.
- **Full Backups Only**: Incremental and log-based backups are not supported. Each backup creates a full copy of the database.

- **No Snapshot Isolation**: Backup consistency is based on the snapshot time. Updates made during export may not be included.
- **Point-in-Time Recovery (PITR)**: PITR is optional and available for up to seven days. It incurs additional Google Cloud costs.
- **Quota Limits**: Firestore allows a maximum of 20 export/import requests per minute and up to 50 concurrent operations per project. Exceeding limits can result in failures.
- **Manual Service Account Configuration**: Service accounts must be registered and assigned required IAM roles for Firestore and Cloud Storage operations.
- **No Transaction or Log Backups**: Firestore does not support log-based backups.
- **Overwrite Behavior**: Restoring a database with the overwrite option deletes all collections before import. If the restore fails, data loss can occur.
- **No Maintenance Mode**: Firestore does not support disabling access during recovery. Pause application writes manually to maintain consistency.
- **Cross-region charges**: Cross-region backup and recovery operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.

Backup Firestore Databases

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, it automatically discovers all projects that the service account can access, as well as all supported workloads such as Firestore databases within those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab. Select the **Firestore** tab to view all the discovered databases.

The **Firestore** tab displays the entity hierarchy for the Firestore environment, allowing you to navigate from the project level down to individual databases.

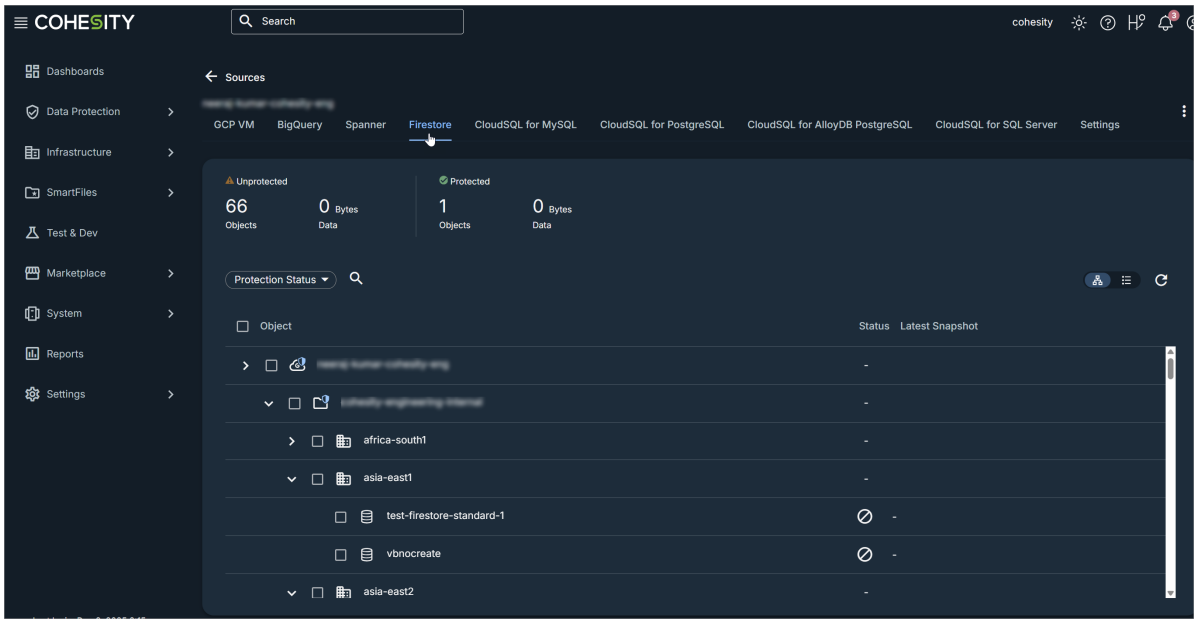

Firestore Entity Hierarchy

The Firestore entity hierarchy in the Cohesity cluster follows this structure:

Structure	Description
Service Account	The dedicated service account in Google Cloud assigned to Cohesity cluster for Firestore protection.
Project ID	The Google Cloud Project ID of the project (s) accessible with the service account.
Region	The Google Cloud region where the databases reside.
Database	The database that can be selected for backup and recovery.

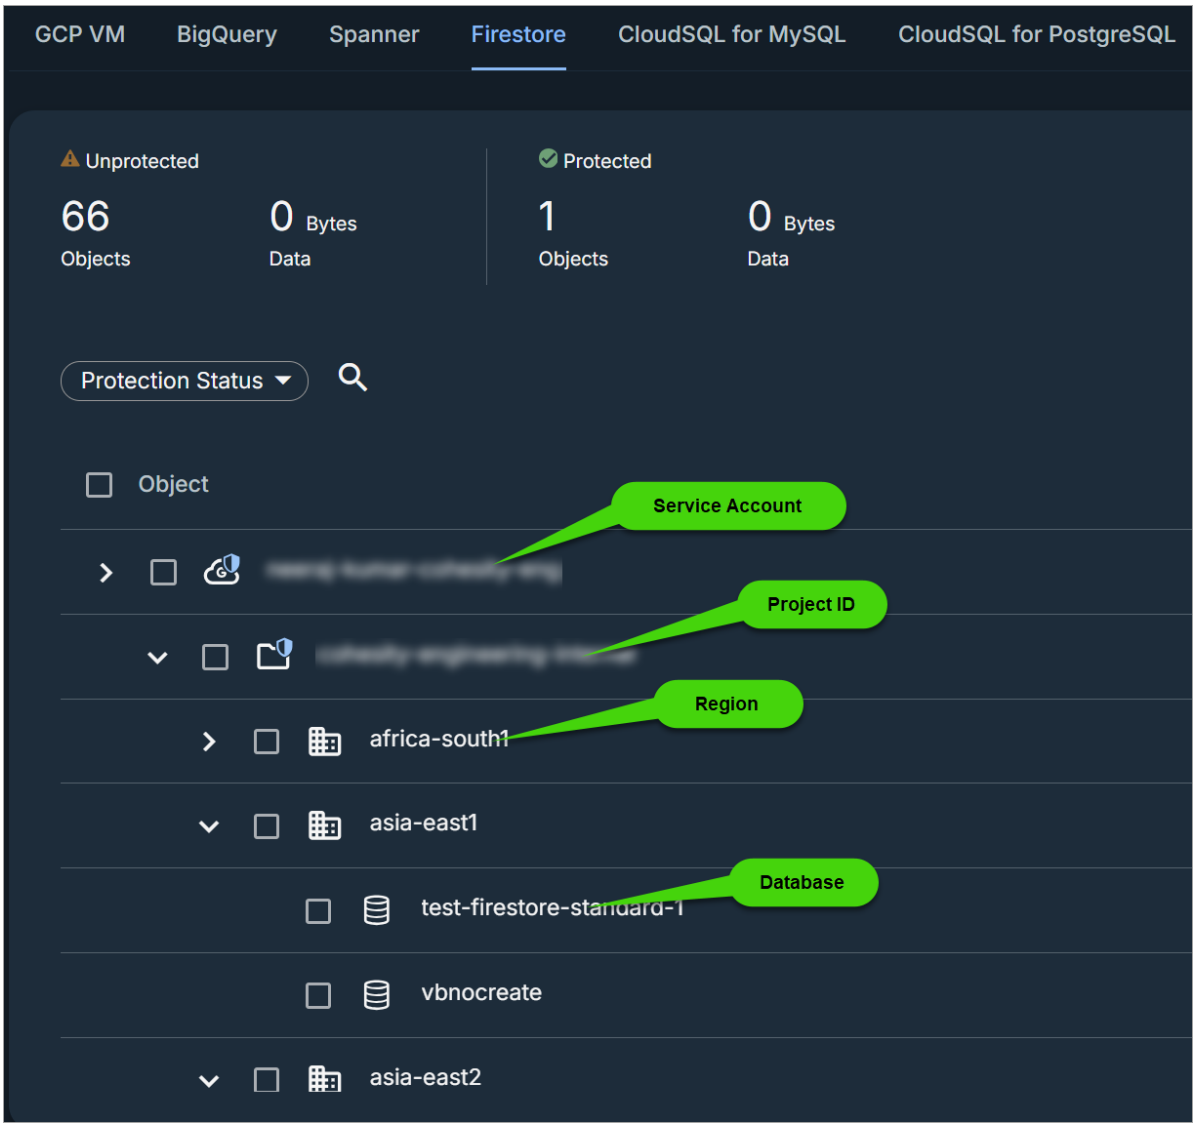

Backup Scope by Firestore Hierarchy

Cohesity cluster allows you to back up Firestore databases at different levels of the hierarchy. You can choose to back up:

- **Individual database** – Back up a single database.
- **Multiple databases within a region** – Back up selected databases from the same region.
- **Region level** – Back up all databases within a specific region.
- **Multiple databases across regions** – Back up selected databases from different regions.
- **Project level** – Back up all databases in all regions within the specified Project ID.

Firestore Backup Process

When you run a protection job in the Cohesity cluster, the selected Firestore database(s) and their collections are exported to a Google Cloud Storage (GCS) staging bucket. The staging bucket temporarily stores the exported data before it is ingested into the Cohesity cluster.

After the export completes, the Cohesity cluster copies the data from the staging bucket into the cluster for retention and applies the defined retention policies.

Once ingestion is complete, Cohesity automatically deletes only the temporary data from the staging bucket to free up storage and reduce Google Cloud Storage costs.

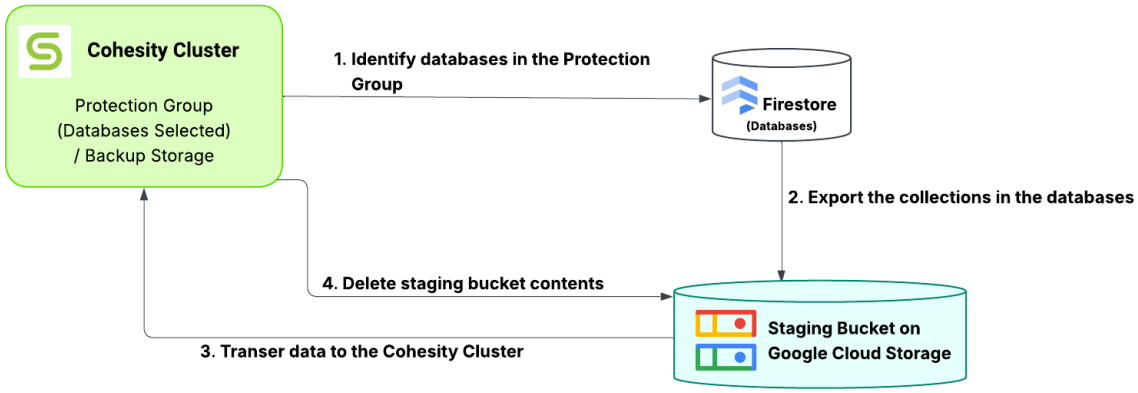

Supported Backup Types

Cohesity supports only full backups for Firestore databases. Incremental backups are not supported.

Protect Firestore

In the Cohesity cluster, Google Cloud Firestore databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

> **Note:** Cohesity supports only full backups for Firestore databases. Incremental backups are not supported.

Prerequisites

Before you protect Firestore databases, ensure that a staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the backup configuration.

- If you choose to **Auto Create** the staging bucket: Cohesity automatically creates a staging bucket in the same region as your database(s) when you start a backup operation.
- If you choose to **specify an existing bucket**: You can specify an existing bucket located in the same region as the databases being protected. Using a bucket in a different region is supported but may incur additional Google Cloud Storage charges.

> **Note:** After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

Protect Firestore Databases

To protect Firestore databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **Firestore**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose databases using one of the following views:
   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the source, project, or region level.
   - **List View**: Select databases from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your database(s). The bucket follows the naming format `cohe-<project-id >-<region-name>`.

     Where:
     - `project-id` is the Google Cloud project ID.
     - `region-name` is the Google Cloud region where the databases reside.
   - **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for backups.

   After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.
7. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** This option is not applicable for Firestore data protection. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

8. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

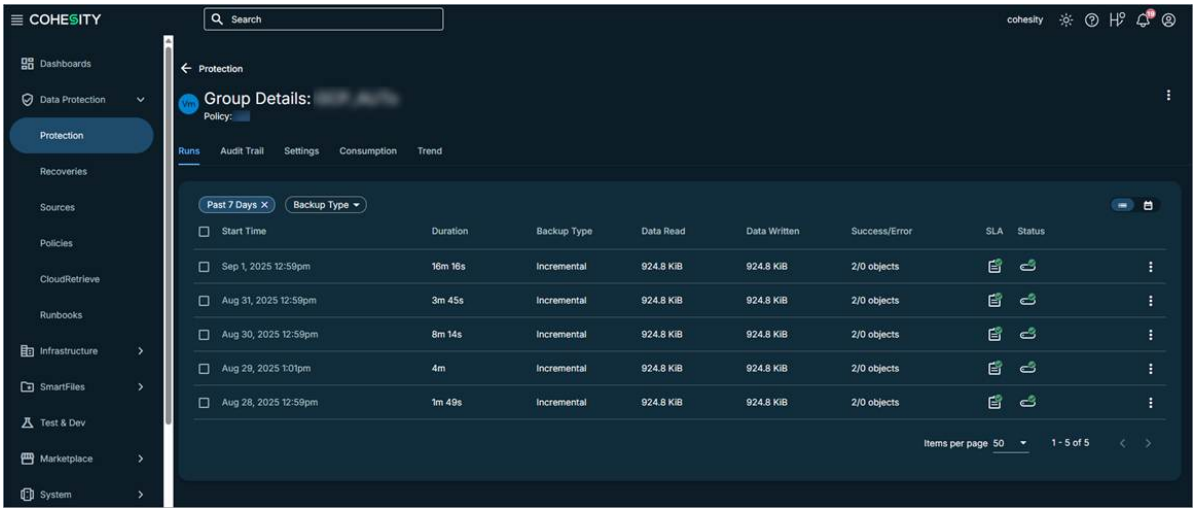

Manage Firestore Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Firestore.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.

5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Firestore.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Firestore.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Firestore.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Firestore Databases

You can restore Firestore databases either to their original GCP source or to a different GCP source and region.

Prerequisites

Before you restore Firestore databases, ensure that a staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the restore process.

- **If you choose to Auto Create the staging bucket**: Cohesity automatically creates a staging bucket in the same region as your database(s) when you start the recovery.
- **If you choose to specify an existing bucket**: You can specify an existing bucket located in the same region as the recovery target. Using a bucket in a different region is supported but may incur additional Google Cloud Storage charges.

> **Note:** After each recovery job run, Cohesity deletes only the temporary data stored in the staging bucket. The staging bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.

Recover Firestore Databases

To recover Firestore databases:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **Firestore**.
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
    - By default, the latest snapshot of the database is selected for recovery.
    - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <database_name>** page appears.
    - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

      To select a snapshot from the **List** mode:

        a. On the **List** tab, snapshots available in the past month are displayed by default.

b. To view snapshots from a different range, select a date range from the **Date** drop-down list.

c. Select the snapshot(s) to recover.

d. By default, the snapshot is recovered from the local Cohesity cluster.

e. Click **Select Recovery Point** and then go to step 5.

To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the database to the same project from which it was backed up. **Note:** Restoring to the original location replaces any existing database with the same name. Type **YES** to confirm.
New Location	Restore the database to an alternate location in a different GCP source. Select a registered GCP source, project, and region. To overwrite an existing database with the same name, enable **Overwrite existing database(s)** and type **YES**. **Note:** Restoring to a new location also replaces any existing database with the same name in the selected region.

7. In **Restore Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your database(s) for the recovery operation. The bucket follows the naming format: `cohe-<project-id>-<region-name>`.

     Where:

- `project-id` is the Google Cloud project ID.
- `region-name` is the Google Cloud region where the databases reside.

- **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for recovery.

Cohesity deletes only the temporary data stored in the staging bucket during each recovery operation. The bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.

8. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

9. Click **Recover**.

Cloud SQL for MySQL

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including Cloud SQL for MySQL.

> **Important:** The backup and recovery of Google Cloud Cloud SQL for MySQL is supported only on Cohesity Cloud Edition on GCP (formerly known as GCP NGCE cluster).

This chapter includes the following topics:

Plan and Prepare for Cloud SQL for MySQL Data Protection

To protect Cloud SQL for MySQL databases with the Cohesity cluster, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and recovery operations. The following table lists the requirements are grouped into those needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Network Connectivity	Source registration, backup/restore	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	Source registration, backup/restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/`

Create an IAM Custom Role in GCP

To enable Cloud SQL for MySQL data protection, you must create an IAM custom role in Google Cloud and assign it to the service account used by the Cohesity cluster.

For detailed instructions, refer to the following topics in the Google Cloud documentation:

- Create a custom IAM role
- Assign a role to a service account

> **Note:** Before you begin creating IAM custom role, ensure that you have created a dedicated service account in Google Cloud for the Cohesity cluster.

Create a Custom Database User

Cohesity requires MySQL credentials to connecto to the database and retrieve the details. These privileges are futher user to perform the backup and recovery operations. You can use MySQL instance admin credentials or creates a dedicated database user with the required privileges and provide those credentials to the Cohesity cluster.

To create a custom database user:

1. Log into the MySQL instance using database administrative privileges.
2. Create a database user by running the following command.

   ```
   CREATE USER <USERNAME> IDENTIFIED BY <PASSWORD>;
   ```

3. Assign the following permissions to the database user.

```
GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, PROCESS,
REFERENCES, INDEX, ALTER, SHOW DATABSE , LOCK TABLES, CREATE VIEW,

SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, SHOW ROUTINE, EVENT, TRIGGER,
SET_USER_ID, REPLICATION CLIENT ON *.* TO <USERNAME>;

FLUSH PRIVILEGES;
```

Considerations for Cloud SQL for MySQL Data Protection

Review and understand the following limitations and known issues before you protect Cloud SQL for MySQL databases:

- **Full Database Backups Only**: Incremental backups are not supported. Each backup captures the complete contents of the selected database.
- **Database-Level Restores**: You can restore individual databases and choose to rename or overwrite them. Restoring an entire MySQL server is not supported.
- **Direct Database Connection**: Cohesity connects directly to the MySQL server using the database credentials you provide. These credentials are required for performing backup and recovery operations, as well as for retrieving the size of the databases on the server.
- **Version Downgrade Not Supported**: Restoring a backup taken from a higher Cloud SQL for MySQL database version to a lower version is not supported.
- **Cross-region charges**: Cross-region backup and recovery operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.

Backup Cloud SQL for MySQL Databases

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, it automatically discovers all projects that the service account can access, as well as all supported workloads such as Cloud SQL for MySQL databases within those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab. Select the **Cloud SQL for MySQL** tab to view all the discovered databases.

The **Cloud SQL for MySQL** tab displays the entity hierarchy for the Cloud SQL for MySQL environment, allowing you to navigate from the project level down to individual databases.

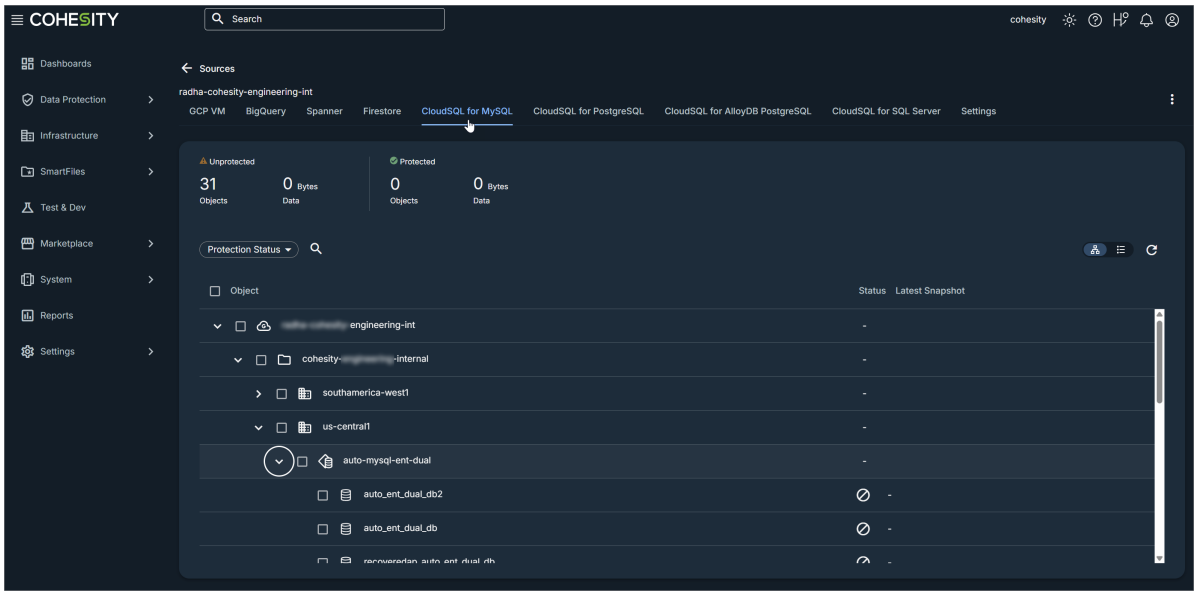

Cloud SQL for MySQL Entity Hierarchy

The following hierarchy describes how Cloud SQL for MySQL entities are organized and displayed in the Cohesity cluster:

Structure	Description
Service Account	The Google Cloud service account configured for the Cohesity cluster to perform Cloud SQL for MySQL backup and recovery operations.
Project ID	The Google Cloud Project ID that the service account has permissions to access.
Region	The Google Cloud region configuration where the MySQL server is hosted.
Database Server	The MySQL server available for backup and recovery.
Database	The individual Cloud SQL for MySQL database that can be backed up or restored.

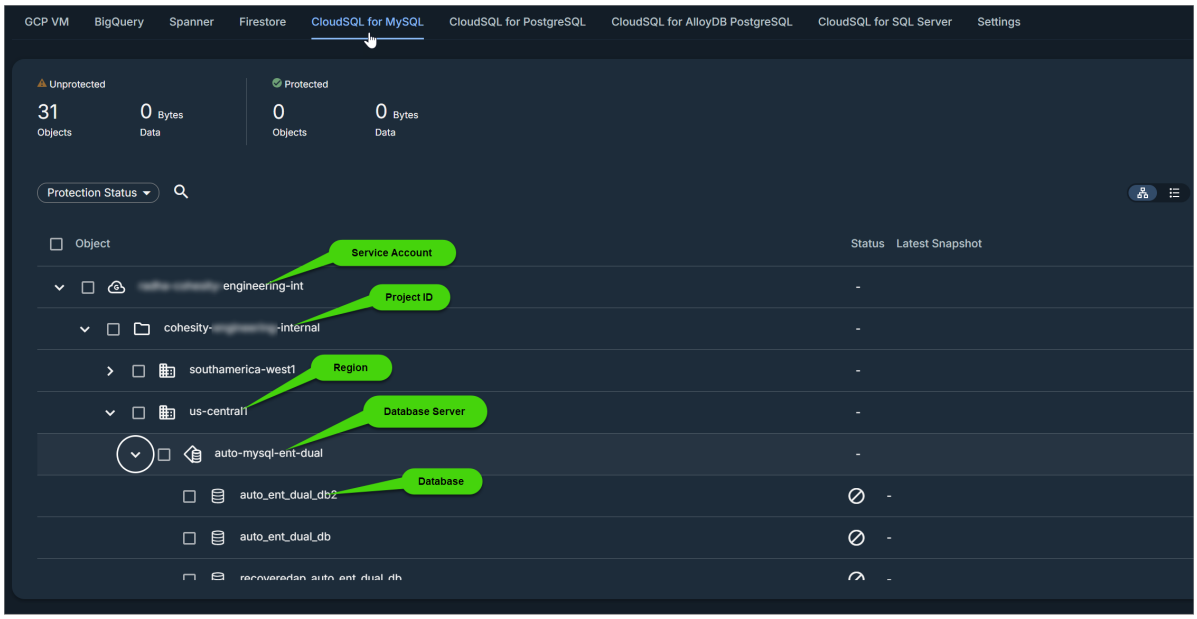

Backup Scope by Cloud SQL for MySQL Hierarchy

Cohesity cluster allows you to back up Cloud SQL for MySQL databases at different levels of the hierarchy. You can choose to back up:

- **Individual database**: Back up a single database.
- **Multiple databases within a region**: Back up selected databases from the same region.
- **Region level**: Back up all databases within a specific region.
- **Multiple databases across regions**: Back up selected databases from different regions.
- **Project level**: Back up all databases in all regions within the specified Project ID.

Cloud SQL for MySQL Backup Process

When you run a protection job for Cloud SQL for MySQL, the Cohesity cluster performs a full backup of the databases specified in the Protection Group. The Cohesity cluster connects to the selected Cloud SQL instance using the database credentials you provided during source registration and uses standard MySQL backup utilities, such as `mysql` and `mysqldump`, to export the complete database contents to the cluster.

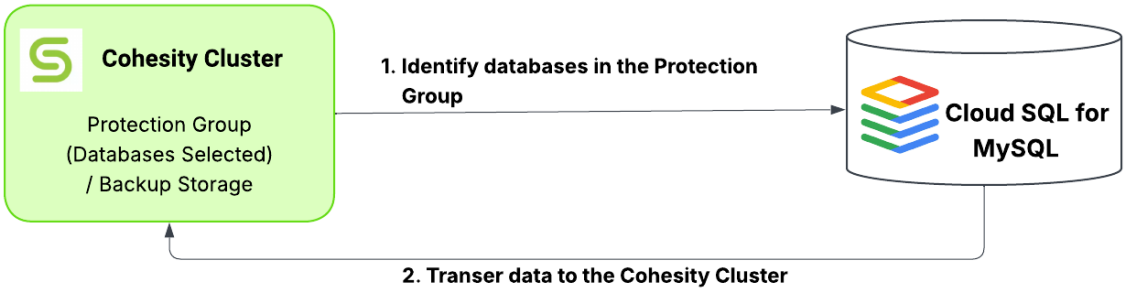

Supported Backup Types

Cohesity supports only full backups for Cloud SQL for MySQL databases. Incremental backups are not supported.

Protect Cloud SQL for MySQL

In the Cohesity cluster, Google Cloud Cloud SQL for MySQL databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites

Before you protect Cloud SQL for MySQL databases, ensure that you add the database credentials required to authenticate and connect to the MySQL server. These credentials allow the Cohesity cluster to establish a secure connection, perform the database dump, and complete backup and recovery operations.

Add Database Credentials

You must add credentials at least at the Cloud SQL for MySQL server level. You can also add credentials at any higher level in the hierarchy, such as the region. Credentials added at a higher level automatically apply to all objects beneath it.

The Cohesity cluster searches for credentials in a bottom-up manner, starting from the specific entity (for example, a server) and moving up the hierarchy until valid credentials are found. When multiple credentials exist, the credentials closest to the MySQL server take precedence.

For example:

- If **region-1** has a set of credentials and **server-1** under that region has its own, the Cohesity cluster uses the credentials on **server-1** to authenticate to **server-1**.

- Another server, **server-2**, under the same region, the Cohesity cluster uses the region-level credentials for authentication because that is the closest level where credentials are specified.

To add database credentials:

1. Navigate to **Data Protection** > **Sources**, locate the registered Google Cloud source, and click on it.
2. Click the **Cloud SQL for MySQL** tab.
3. Choose one of the following options:
    - Select the checkbox next to the MySQL server and click **Select All Child Objects**.

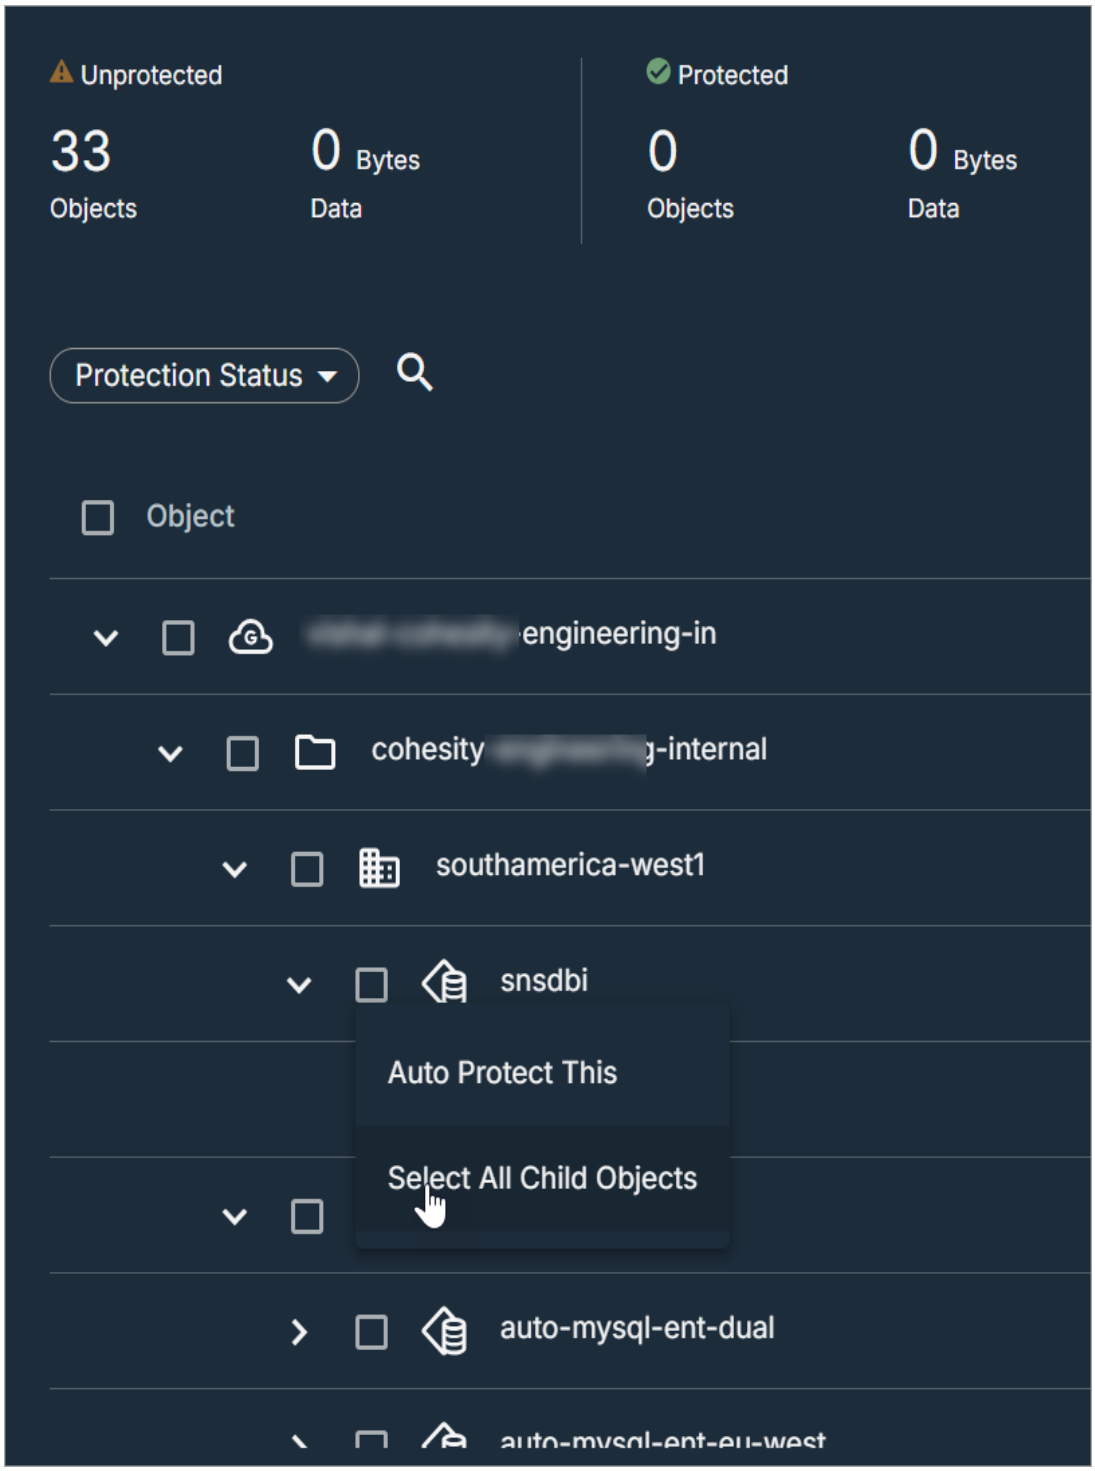

Then, click **Database Credentials** at the top left corner.

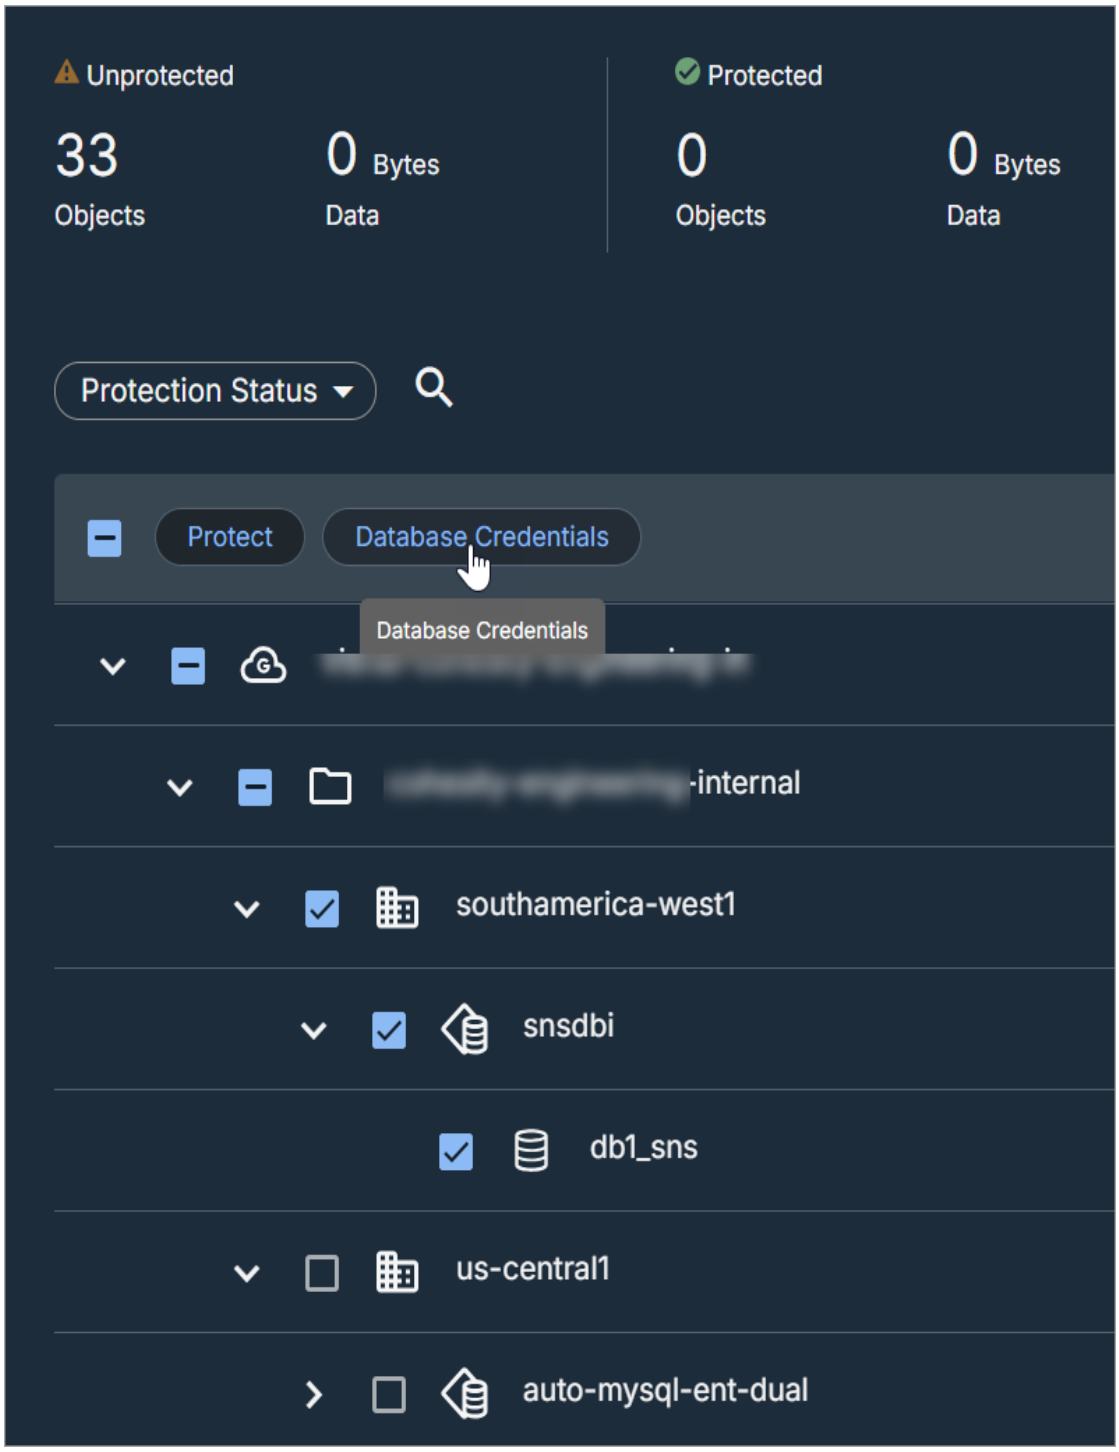

- Click the next to the required MySQL server.

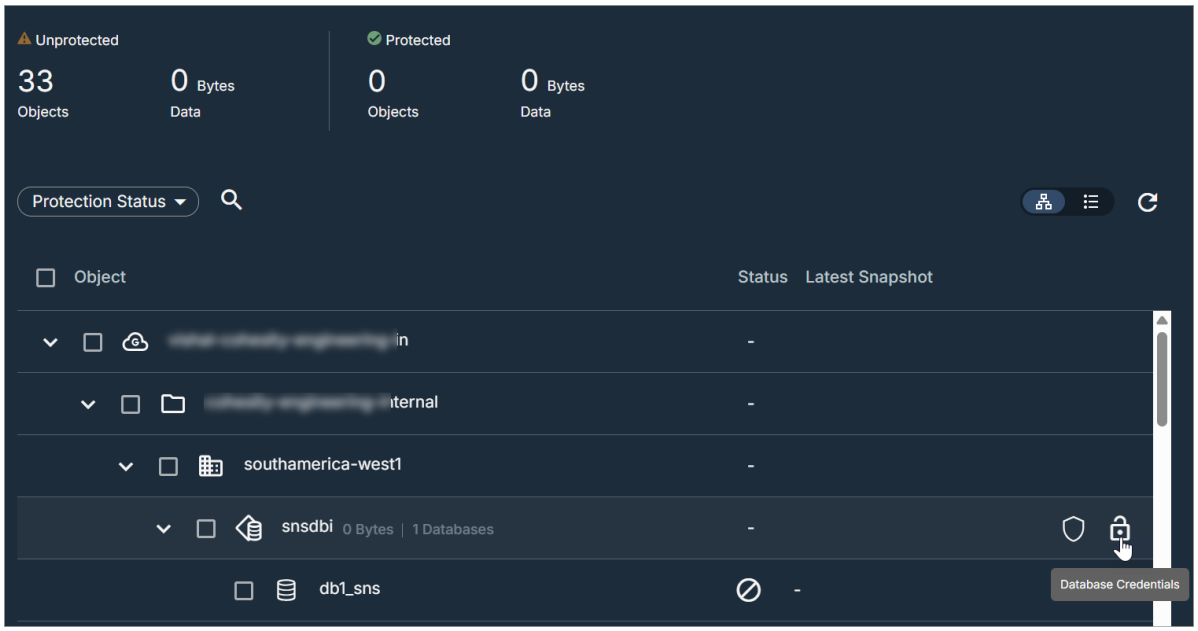

4. In the **Database Credentials** dialog box, enter the MySQL server **username** and **password**, and click **Save**.

Protect Cloud SQL for MySQL Databases

To protect Cloud SQL for MySQL databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **Cloud SQL for MySQL**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose databases using one of the following views:
   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the source, project, or region level.
   - **List View**: Select databases from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.

6. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** This option is not applicable for Cloud SQL for MySQL data protection. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

7. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

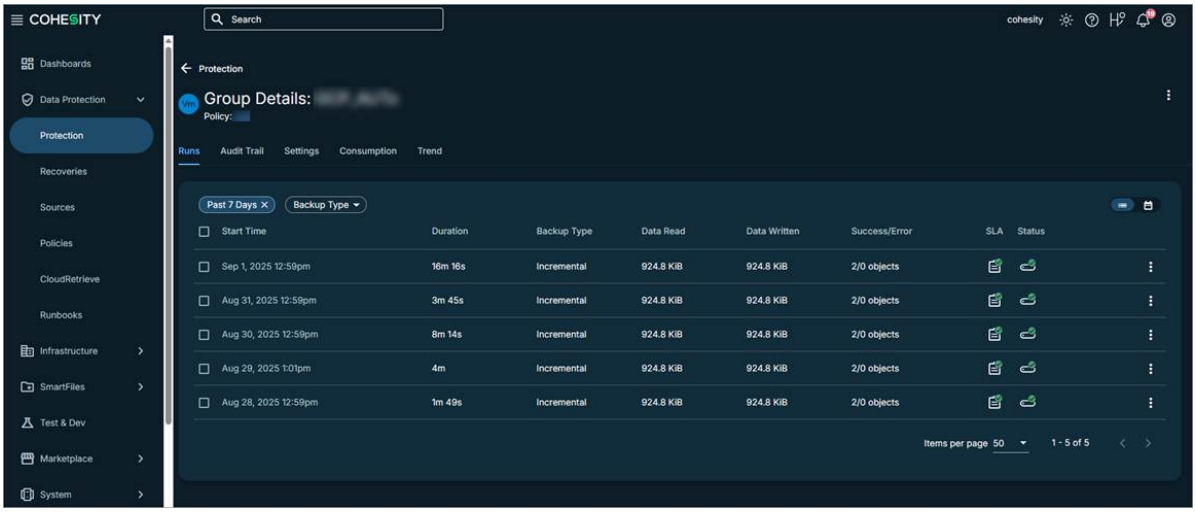

Manage Cloud SQL for MySQL Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.

5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Cloud SQL for MySQL Databases

You can restore Cloud SQL for MySQL databases either to their original GCP source or to a different GCP source and region.

Prerequisites

To recover Cloud SQL for MySQL databases, add the database credentials for the target Cloud SQL MySQL databases server if you plan to restore it to a new location.

Recover Cloud SQL for MySQL Databases

To recover Cloud SQL for MySQL databases:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **Cloud SQL for MySQL**.
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
    - By default, the latest snapshot of the database is selected for recovery.
    - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <database_name>** page appears.
    - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

      To select a snapshot from the **List** mode:

        a. On the **List** tab, snapshots available in the past month are displayed by default.

        b. To view snapshots from a different range, select a date range from the **Date** drop-down list.

        c. Select the snapshot(s) to recover.

        d. By default, the snapshot is recovered from the local Cohesity cluster.

        e. Click **Select Recovery Point** and then go to step 5.

      To select a snapshot from the **Timeline** mode:

        a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

        b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.
d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the database to the same project from which it was backed up. > **Note:** Restoring to the original location replaces any existing database with the same name. Type **YES** to confirm.
New Location	Restore the database to an alternate location in a different GCP source. Select a registered GCP source, project, and region. To overwrite an existing database with the same name, enable **Overwrite existing database(s)** and type **YES**. > **Note:** Restoring to a new location also replaces any existing database with the same name in the selected region.

7. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

8. Click **Recover**.

Cloud SQL for PostgreSQL

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including Cloud SQL for PostgreSQL.

> **Important:** The backup and recovery of Google Cloud Cloud SQL for PostgreSQL is supported only on Cohesity Cloud Edition on GCP.

This chapter includes the following topics:

Plan and Prepare for Cloud SQL for PostgreSQL Data Protection

To protect Cloud SQL for PostgreSQL databases with the Cohesity cluster, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and recovery operations. The following table lists the requirements are grouped into those needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Network Connectivity	Source registration, backup/restore	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	Source registration, backup/restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/`
Database User	Backup and restore	Cohesity requires PostgreSQL user credentials to perform backup and recovery operations on Cloud SQL for PostgreSQL databases. You can use the PostgreSQL instance admin user or create a dedicated database user with the required privileges for backup and recovery operations.

Create a Custom Database User

Cohesity requires PostgreSQL user credentials to retrieve database size information and perform backup and recovery operations. You can use the PostgreSQL instance admin credentials or create a dedicated database user with the required privileges and provide those credentials to the Cohesity cluster.

These credentials allow the Cohesity cluster to establish a secure connection to the PostgreSQL instance, perform database dumps, and complete backup and recovery operations.

To create a custom database user:

1. Log in to the PostgreSQL instance using a user with administrative privileges.
2. Create a database user by running the following command:

```
CREATE USER <username> WITH PASSWORD '<password'>;
```

3. Assign the following permissions to the database user.

```
GRANT cloudsqlsuperuser,pg_read_all_data,pg_write_all_data TO
<username>;
ALTER USER <username> WITH CREATEDB;
ALTER USER <username> WITH BYPASSRLS;
```

Create an IAM Custom Role in GCP

To enable Cloud SQL for PostgreSQL data protection, you must create an IAM custom role in Google Cloud and assign it to the service account used by the Cohesity cluster.

For detailed instructions, refer to the following topics in the Google Cloud documentation:

- Create a custom IAM role
- Assign a role to a service account

> **Note:** Before you begin creating IAM custom role, ensure that you have created a dedicated service account in Google Cloud for the Cohesity cluster.

Considerations for Cloud SQL for PostgreSQL Data Protection

Review and understand the following limitations and known issues before you protect Cloud SQL for PostgreSQL databases:

- **Security Privileges Not Restored**: Restoring database security privileges is not supported.
- **Owner and Privilege Options**: You can use the `-no-owner` and `-no-privileges` options during restore. The metadata captured during backup is displayed as the owner and ACL details in the restore activity logs.
- **Missing Roles Do Not Block Restore**: The restore operation does not fail if the owner or role from the source does not exist on the destination instance.
- **Post-Restore Role Assignment**: After restore, the database is assigned to the role configured in the Cohesity cluster for the destination instance.
- **Manual Ownership Updates**: You must manually update database ownership after the restore, if required.
- **Unsupported Characters in Database Names**: The restore workflow does not support the following characters in database names: `, @, \, [, ], !, #, %, ^, ., ,, &, *, (, ), <, >, ?, /, |, }, {, ~, :, ', ", ;, +, =, -`

- **Uppercase Usernames Not Supported**: Uppercase usernames are not supported for users added after the PostgreSQL server is created.
- **Temporary Objects May Cause Backup Failures**: Full or differential incremental backups may fail if temporary objects exist. Remove temporary objects manually and retry the backup.
- **Version Downgrade Not Supported**: Restoring a backup taken from a higher Cloud SQL for PostgreSQL database version to a lower version is not supported.
- **Objects owned by multiple users**: PostgreSQL database may contain objects owned by multiple users. Ensure that the backup user has sufficient privileges to access the tables and other objects owned by all users. Otherwise the backup operations may fail or will be incomplete.
- **Cross-region charges**: Cross-region backup and recovery operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.

Backup Cloud SQL for PostgreSQL Databases

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, it automatically discovers all projects that the service account can access, as well as all supported workloads such as Cloud SQL for PostgreSQL databases within those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab. Select the **Cloud SQL for PostgreSQL** tab to view all the discovered databases.

The **Cloud SQL for PostgreSQL** tab displays the entity hierarchy for the Cloud SQL for PostgreSQL environment, allowing you to navigate from the project level down to individual databases.

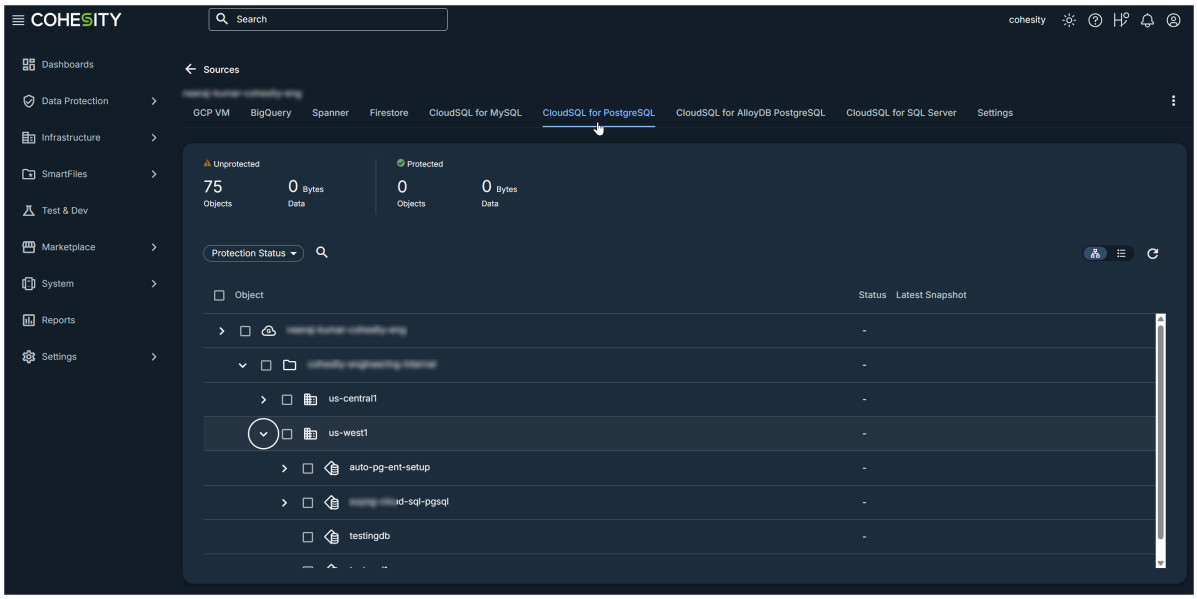

Cloud SQL for PostgreSQL Entity Hierarchy

The following hierarchy describes how Cloud SQL for PostgreSQL entities are organized and displayed in the Cohesity cluster:

Structure	Description
Service Account	The Google Cloud service account configured for the Cohesity cluster to perform Cloud SQL for PostgreSQL backup and recovery operations.
Project ID	The Google Cloud Project ID that the service account has permissions to access.
Region	The Google Cloud region configuration where the PostgreSQL instance is hosted.
Instance	The PostgreSQL instance available for backup and recovery.
Database	The individual Cloud SQL for PostgreSQL database that can be backed up or restored.

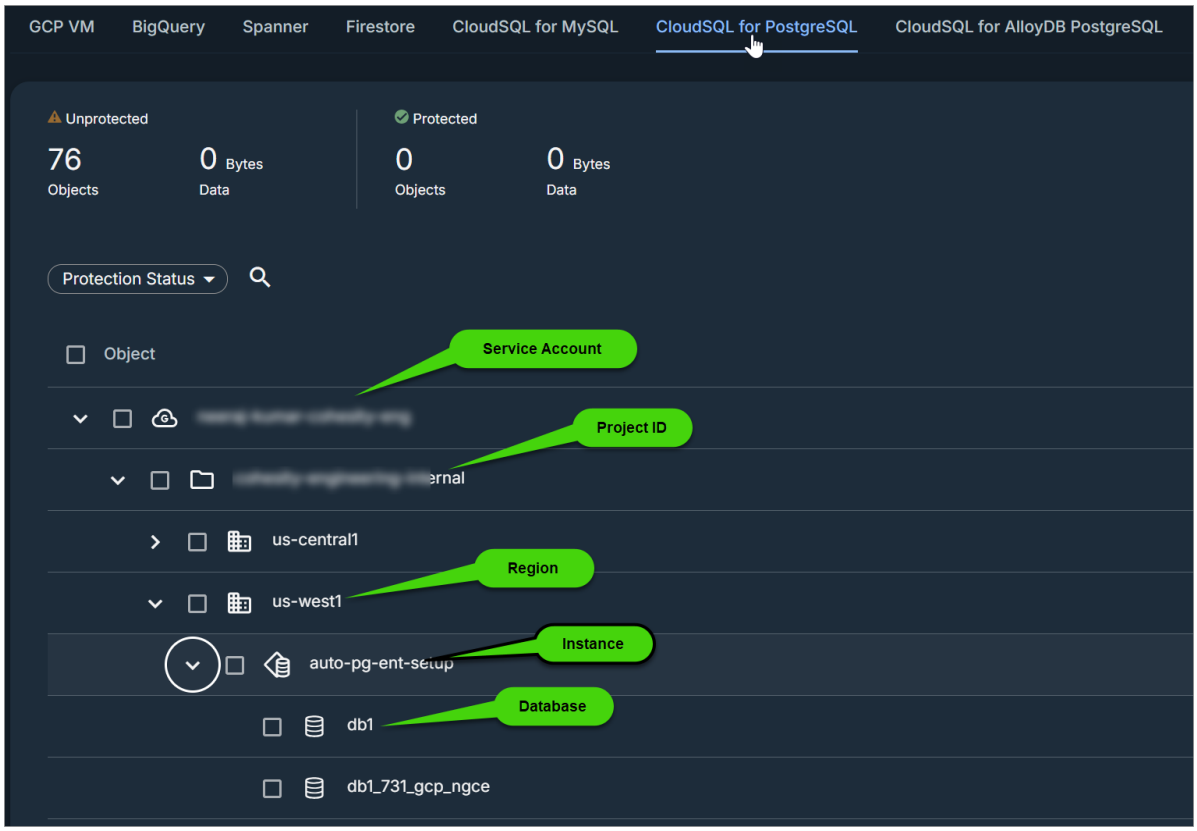

Backup Scope by Cloud SQL for PostgreSQL Hierarchy

Cohesity cluster allows you to back up Cloud SQL for PostgreSQL databases at different levels of the hierarchy. You can choose to back up:

- **Individual database**: Back up a single database.
- **Multiple databases within a region**: Back up selected databases from the same region.
- **Region level**: Back up all databases within a specific region.
- **Multiple databases across regions**: Back up selected databases from different regions.
- **Project level**: Back up all databases in all regions within the specified Project ID.

Cloud SQL for PostgreSQL Backup Process

When you run a protection job for Cloud SQL for PostgreSQL, the Cohesity cluster performs a full backup of the databases specified in the Protection Group. The Cohesity cluster connects to the selected Cloud SQL instance using the database credentials you provided during source registration and uses standard PostgreSQL backup utilities, such as `PostgreSQL` and `pgdump`, to export the complete database contents to the cluster.

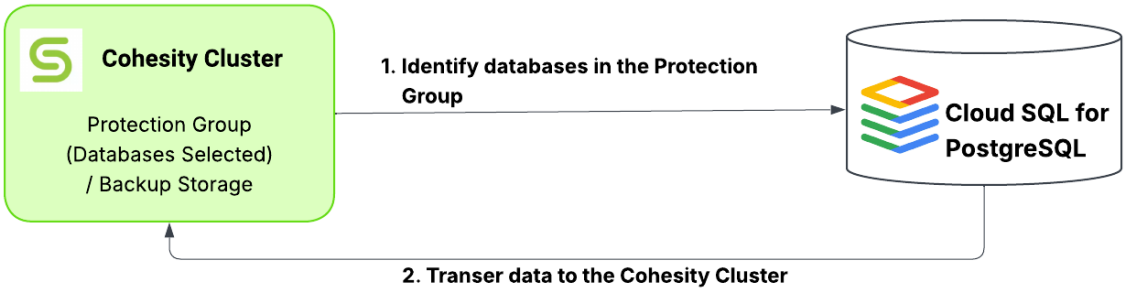

Supported Backup Types

Cohesity supports only full backups for Cloud SQL for PostgreSQL databases. Incremental backups are not supported.

Protect Cloud SQL for PostgreSQL

In the Cohesity cluster, Google Cloud Cloud SQL for PostgreSQL databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites

Before you protect Cloud SQL for PostgreSQL databases, ensure that you add the database credentials required to authenticate and connect to the PostgreSQL server. These credentials allow the Cohesity cluster to establish a secure connection, perform the database dump, and complete backup and recovery operations.

Add Database Credentials

You must add credentials at least at the Cloud SQL for PostgreSQL server level. You can also add credentials at any higher level in the hierarchy, such as the region. Credentials added at a higher level automatically apply to all objects beneath it. When multiple credentials exist in the hierarchy, the credentials closest to the PostgreSQL server take precedence. For example:

- If **region-1** has a set of credentials and **server-1** under that region has its own, the Cohesity cluster uses the credentials on **server-1** to authenticate to **server-1**.
- Another server, **server-2**, under the same region, the Cohesity cluster uses the region-level credentials for authentication because that is the closest level where credentials are defined.

To add database credentials:

1. Navigate to **Data Protection** > **Sources**, locate the registered Google Cloud source, and click on it.
2. Click the **Cloud SQL for PostgreSQL** tab.
3. Choose one of the following options:
    - Select the checkbox next to the PostgreSQL instance and click **Select All Child Objects**.

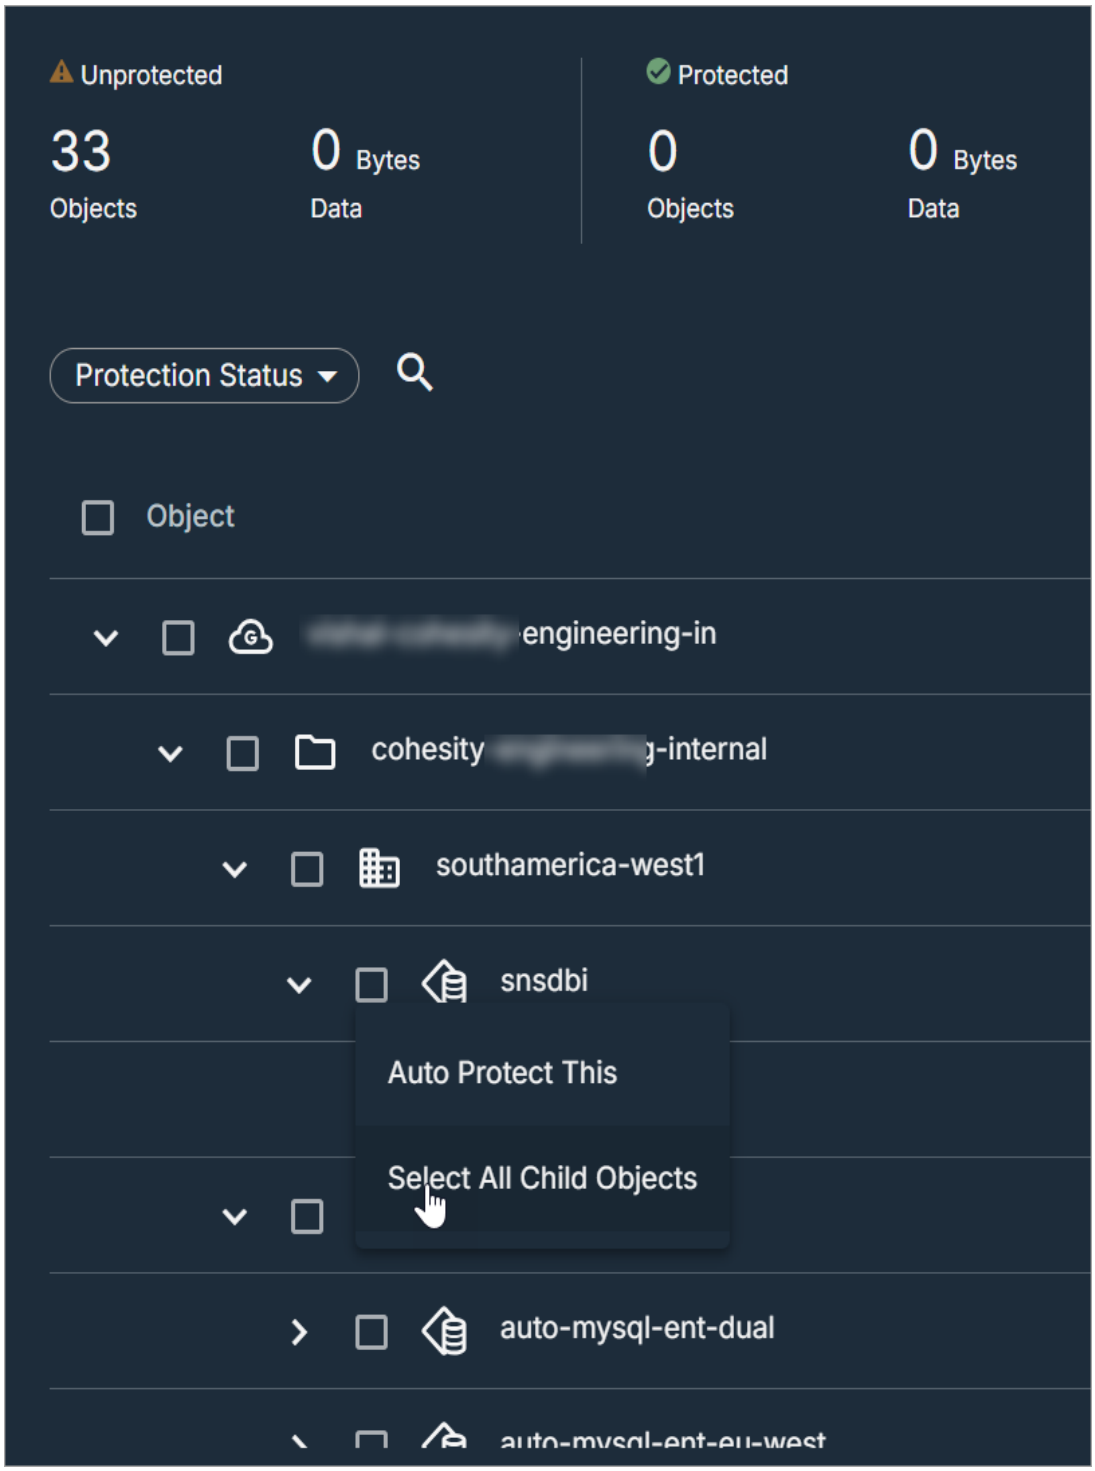

Then, click **Database Credentials** at the top left corner.

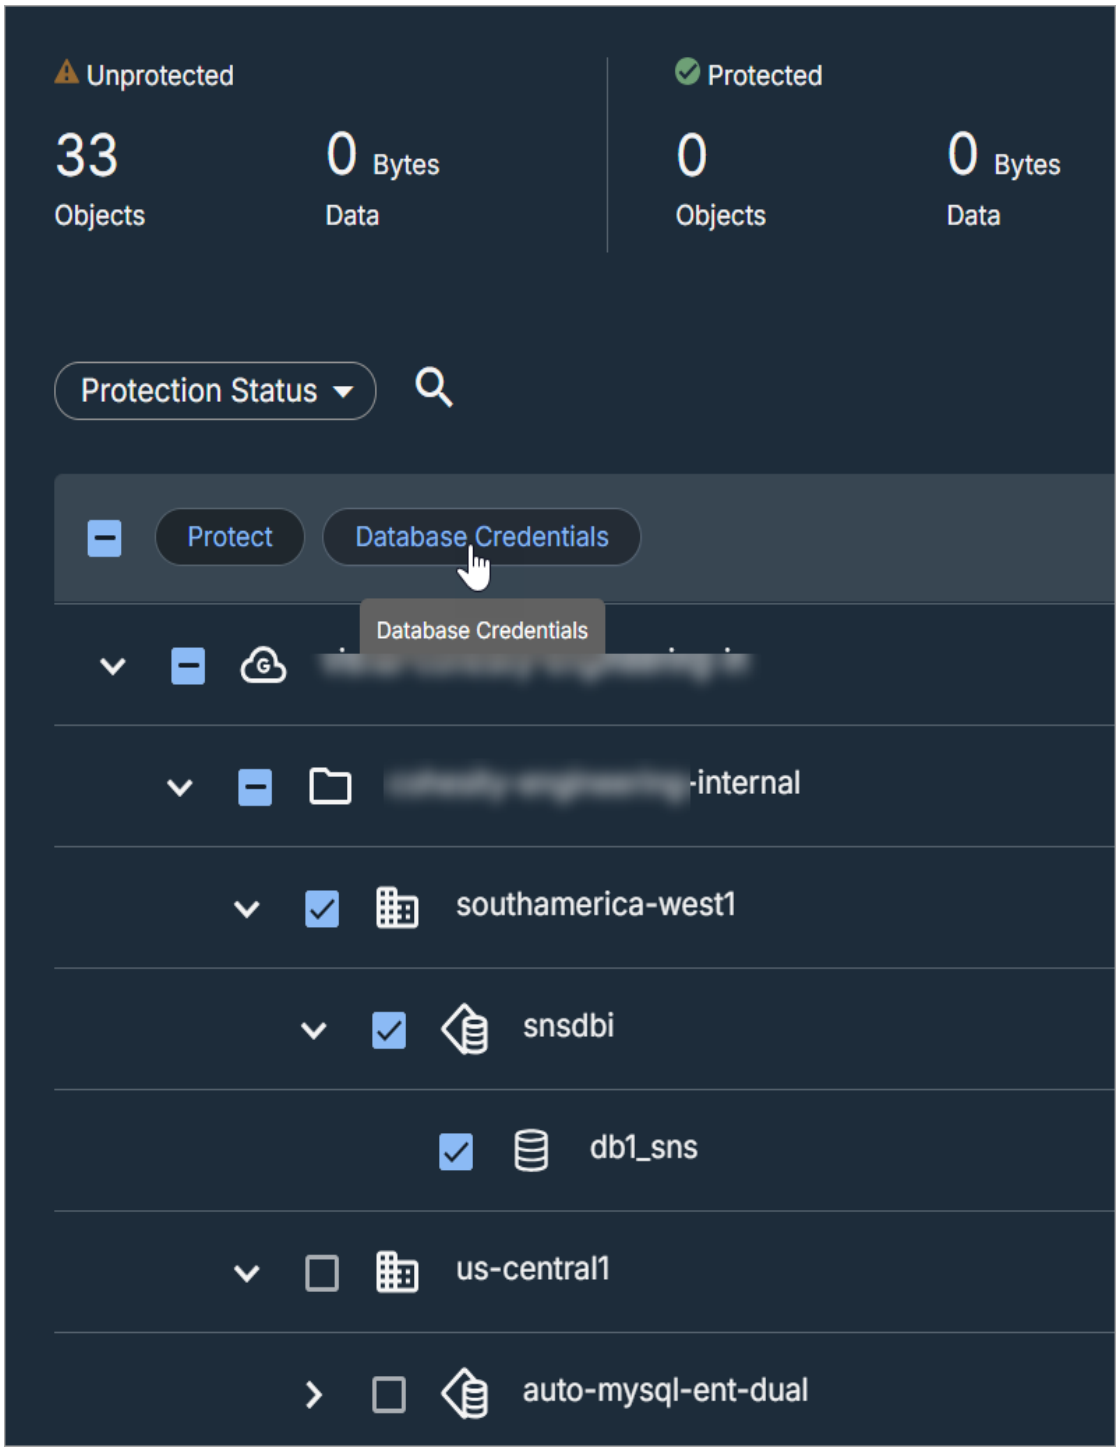

- Click the next to the required PostgreSQL instance.

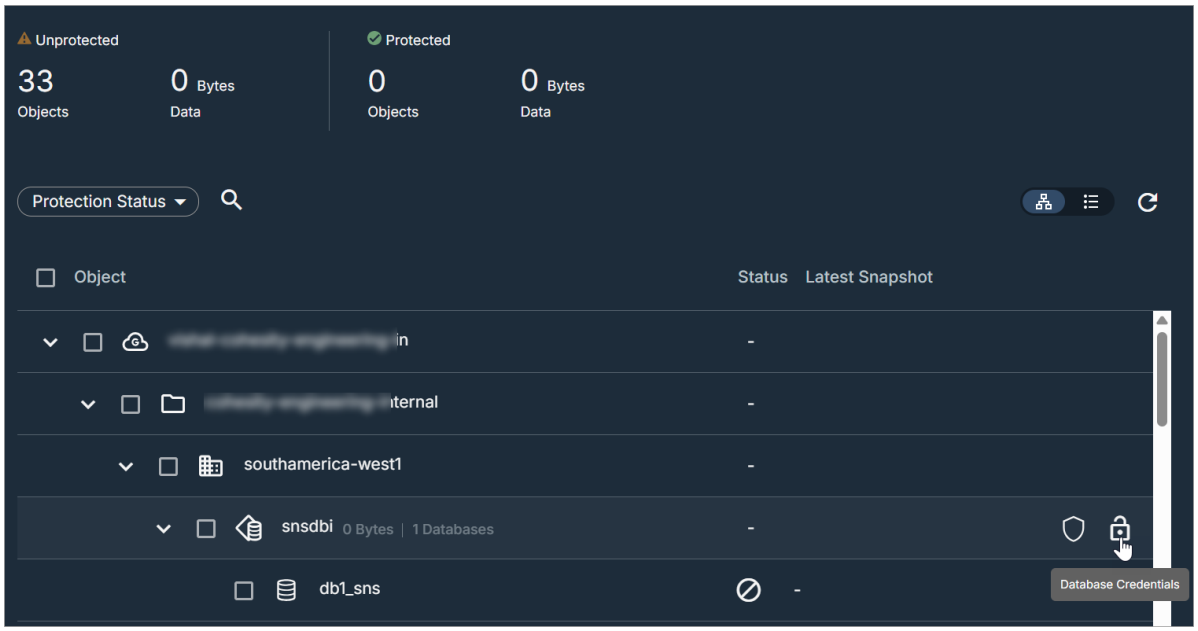

4. In the **Database Credentials** dialog box, enter the PostgreSQL instance **username** and **password**, and click **Save**.

Protect Cloud SQL for PostgreSQL Databases

To protect Cloud SQL for PostgreSQL databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **Cloud SQL for PostgreSQL**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose databases using one of the following views:
   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the source, project, or region level.
   - **List View**: Select databases from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.

6. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** This option is not applicable for Cloud SQL for PostgreSQL data protect ion. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

7. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

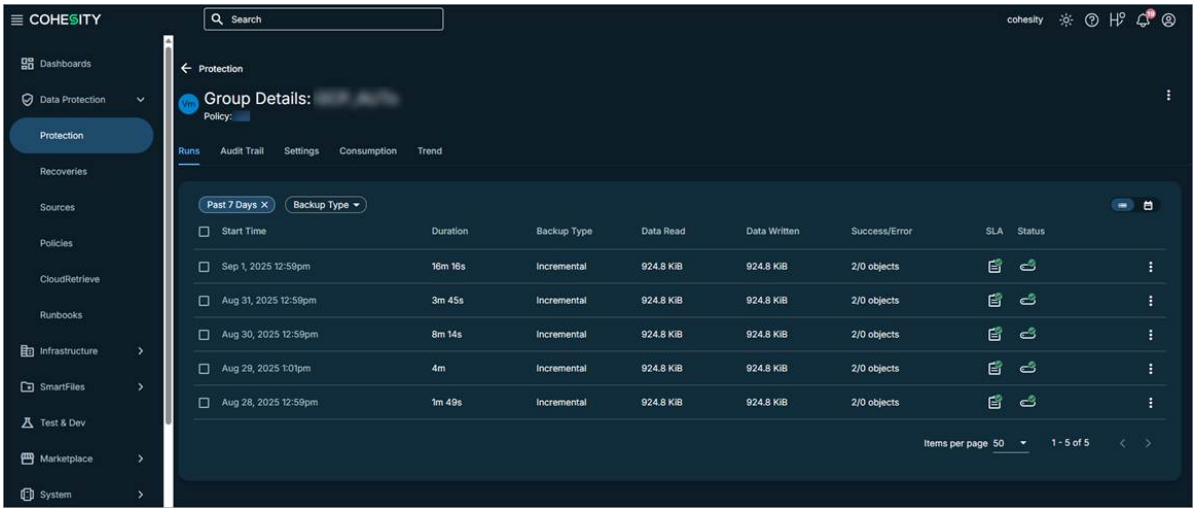

Manage Cloud SQL for PostgreSQL Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.

5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Cloud SQL for PostgreSQL Databases

You can restore Cloud SQL for PostgreSQL databases either to their original GCP source or to a different GCP source and region.

Prerequisites

To recover Cloud SQL for PostgreSQL databases, add the database credentials for the target Cloud SQL PostgreSQL instance if you plan to restore it to a new location.

Recover Cloud SQL for PostgreSQL Databases

To recover Cloud SQL for PostgreSQL databases:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **Cloud SQL for PostgreSQL**.
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
    - By default, the latest snapshot of the database is selected for recovery.
    - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <database_name>** page appears.
    - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

      To select a snapshot from the **List** mode:

        a. On the **List** tab, snapshots available in the past month are displayed by default.

        b. To view snapshots from a different range, select a date range from the **Date** drop-down list.

        c. Select the snapshot(s) to recover.

        d. By default, the snapshot is recovered from the local Cohesity cluster.

        e. Click **Select Recovery Point** and then go to step 5.

      To select a snapshot from the **Timeline** mode:

        a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

        b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the database to the same project from which it was backed up. **Note:** Restoring to the original location replaces any existing database with the same name. Type **YES** to confirm.
New Location	Restore the database to an alternate location in a different GCP source. Select a registered GCP source, project, and region. To overwrite an existing database with the same name, enable **Overwrite existing database(s)** and type **YES**. **Note:** Restoring to a new location also replaces any existing database with the same name in the selected region. If you want to restore the database to a different instance within the same GCP source, you must use this option.

7. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

8. Click **Recover**.

GCP-Postgresql-Materialized-view

Delete this text and replace it with your own content.

AlloyDB for PostgreSQL

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including AlloyDB for PostgreSQL.

> **Important:** The backup and recovery of Google Cloud AlloyDB for PostgreSQL is supported only on Cohesity Cloud Edition on GCP.

This chapter includes the following topics:

Plan and Prepare AlloyDB for PostgreSQL Data Protection

To protect AlloyDB for PostgreSQL databases with the Cohesity cluster, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and recovery operations. The following table lists the requirements are grouped into those needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Network Connectivity	Source registration, backup/restore	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	Source registration, backup/restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/`
Database User	Backup and restore	Cohesity requires PostgreSQL user credentials to perform backup and recovery operations on AlloyDB for PostgreSQL databases. You can use the AlloyDB cluster admin user or create a dedicated database user with the required privileges for backup and recovery operations.

Create a Custom Database User

Cohesity requires PostgreSQL user credentials to discover databases in an AlloyDB for PostgreSQL cluster, retrieve database size information, and perform backup and recovery operations. You can use the AlloyDB cluster admin credentials or create a dedicated database user with the required privileges and provide those credentials to the Cohesity cluster.

These credentials allow the Cohesity cluster to establish a secure connection to the AlloyDB cluster, perform database dumps, and complete backup and recovery operations.

To create a custom database user:

1. Log in to the AlloyDB for PostgreSQL cluster using a user with administrative privileges.

   ```
   Create a database user by running the following command:
   ```

2. CREATE USER <username> WITH PASSWORD '<password'>;
3. Assign the following permissions to the database user.

   ```
   GRANT alloydbsuperuser, pg_read_all_data, pg_write_all_data TO
   <username>;
   ALTER USER <username> WITH CREATEDB;
   ALTER USER <username> WITH BYPASSRLS;
   ```

Create an IAM Custom Role in GCP

To enable AlloyDB for PostgreSQL data protection, you must create an IAM custom role in Google Cloud and assign it to the service account used by the Cohesity cluster.

For detailed instructions, refer to the following topics in the Google Cloud documentation:

- Create a custom IAM role
- Assign a role to a service account

> **Note:** Before you begin creating IAM custom role, ensure that you have created a dedicated service account in Google Cloud for the Cohesity cluster.

Considerations for AlloyDB for PostgreSQL Data Protection

Review and understand the following limitations and known issues before you protect AlloyDB for PostgreSQL databases:

- **Security Privileges Not Restored**: Restoring database security privileges is not supported.
- **Owner and Privilege Options**: You can use the `-no-owner` and `-no-privileges` options during restore. The metadata captured during backup is displayed as the owner and ACL details in the restore activity logs.
- **Missing Roles Do Not Block Restore**: The restore operation does not fail if the owner or role from the source does not exist on the destination instance.
- **Post-Restore Role Assignment**: After restore, the database is assigned to the role configured in the Cohesity cluster for the destination instance.

- **Manual Ownership Updates**: You must manually update database ownership after the restore, if required.
- **Unsupported Characters in Database Names**: The restore workflow does not support the following characters in database names: `, @, \, [, ], !, #, %, ^, ., ,, &, *, (, ), <, >, ?, /, |, }, {, ~, :, ', ", ;, +, =, -`
- **Uppercase Usernames Not Supported**: Uppercase usernames are not supported for users added after the AlloyDB PostgreSQL instance is created.
- **Temporary Objects May Cause Backup Failures**: Full or differential incremental backups may fail if temporary objects exist. Remove temporary objects manually and retry the backup.
- **Version Downgrade Not Supported**: Restoring a backup taken from a higher AlloyDB for PostgreSQL database version to a lower version is not supported.
- **Cluster visibility**: Only cluster with primary instances are displayed. Clusters without any instances are not displayed, even if the cluster exists.
- **Cross-region charges**: Cross-region backup and recovery operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.

Backup AlloyDB for PostgreSQL Databases

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, it automatically discovers all projects that the service account can access, as well as all supported workloads such as AlloyDB for PostgreSQL databases within those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab. Select the **AlloyDB for PostgreSQL** tab to view all the discovered databases.

The **AlloyDB for PostgreSQL** tab displays the entity hierarchy for the AlloyDB for PostgreSQL environment, allowing you to navigate from the project level down to individual databases.

AlloyDB for PostgreSQL Entity Hierarchy

The following hierarchy describes how AlloyDB for PostgreSQL entities are organized and displayed in the Cohesity cluster:

> **Note:** Only cluster with primary instances are displayed. Clusters without any instances are not displayed, even if the cluster exists.

Structure	Description
Service Account	The Google Cloud service account configured for the Cohesity cluster to perform AlloyDB for PostgreSQL backup and recovery operations.
Project ID	The Google Cloud Project ID that the service account has permissions to access.
Region	The Google Cloud region where the CloudSQL for AlloyDB PostgreSQL cluster reside.
AlloyDB Cluster	The CloudSQL for AlloyDB PostgreSQL cluster available for backup and recovery.
Database	The databases that can be selected for backup and recovery. **Note:** Databases appear only after the required credentials are applied on the AlloyDB cluster.

Backup Scope by AlloyDB for PostgreSQL Hierarchy

Cohesity cluster allows you to back up AlloyDB for PostgreSQL databases at different levels of the hierarchy. You can choose to back up:

- **Individual database**: Back up a single database.
- **Multiple databases within a region**: Back up selected databases from the same region.
- **Region level**: Back up all databases within a specific region.
- **Multiple databases across regions**: Back up selected databases from different regions.
- **Project level**: Back up all databases in all regions within the specified Project ID.

AlloyDB for PostgreSQL Backup Process

When you run a protection job for AlloyDB for PostgreSQL, the Cohesity cluster performs a full backup of the databases specified in the Protection Group. The Cohesity cluster connects to the selected Cloud SQL instance using the database credentials you provided during

source registration and uses standard PostgreSQL backup utilities, such as `pgdump`, to export the complete database contents to the cluster.

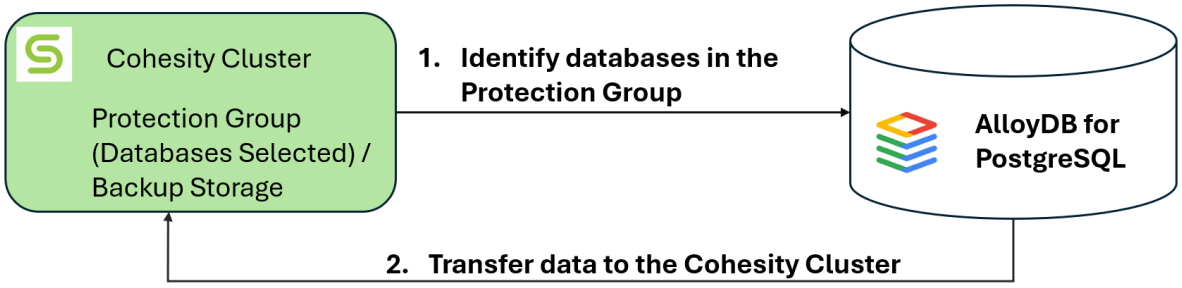

Supported Backup Types

Cohesity supports only full backups for AlloyDB for PostgreSQL databases. Incremental backups are not supported.

Protect AlloyDB for PostgreSQL

In the Cohesity cluster, Google Cloud AlloyDB for PostgreSQL databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites

Before you protect AlloyDB for PostgreSQL databases, ensure that you add the database credentials required to authenticate and connect to the AlloyDB cluster. These credentials allow the Cohesity cluster to establish a secure connection, perform the database dump, and complete backup and recovery operations.

Add Database Credentials

You must add credentials at least at the AlloyDB cluster level. You can also add credentials at any higher level in the hierarchy, such as the region. Credentials added at a higher level automatically apply to all objects beneath it. When multiple credentials exist in the hierarchy, the credentials closest to the AlloyDB cluster take precedence. For example:

- If **region-1** has a set of credentials and **cluster-1** under that region has its own, the Cohesity cluster uses the credentials on **cluster-1** to authenticate to **cluster-1**.
- Another cluster, **cluster-2**, under the same region, the Cohesity cluster uses the region-level credentials for authentication because that is the closest level where credentials are defined.

To add database credentials:

1. Navigate to **Data Protection** > **Sources**, locate the registered Google Cloud source, and click on it.
2. Click the **AlloyDB - PostgreSQL** tab.
3. Choose one of the following options:
    - Select the checkbox next to the AlloyDB cluster and click **Select All Child Objects**.

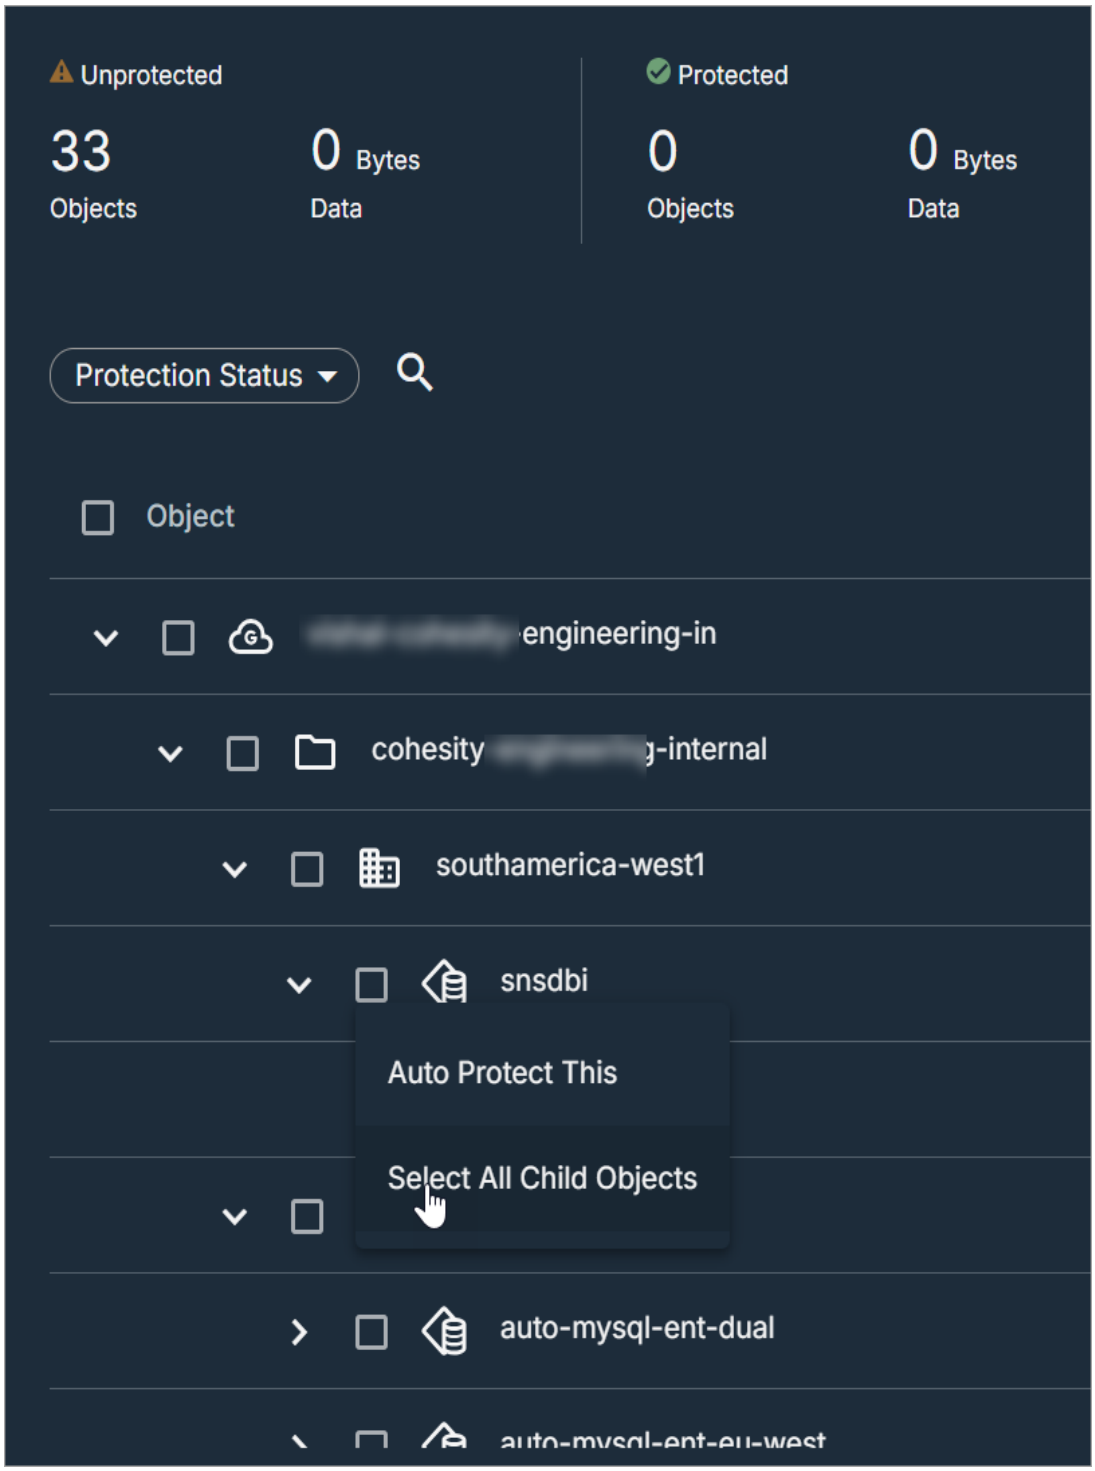

Then, click **Database Credentials** at the top left corner.

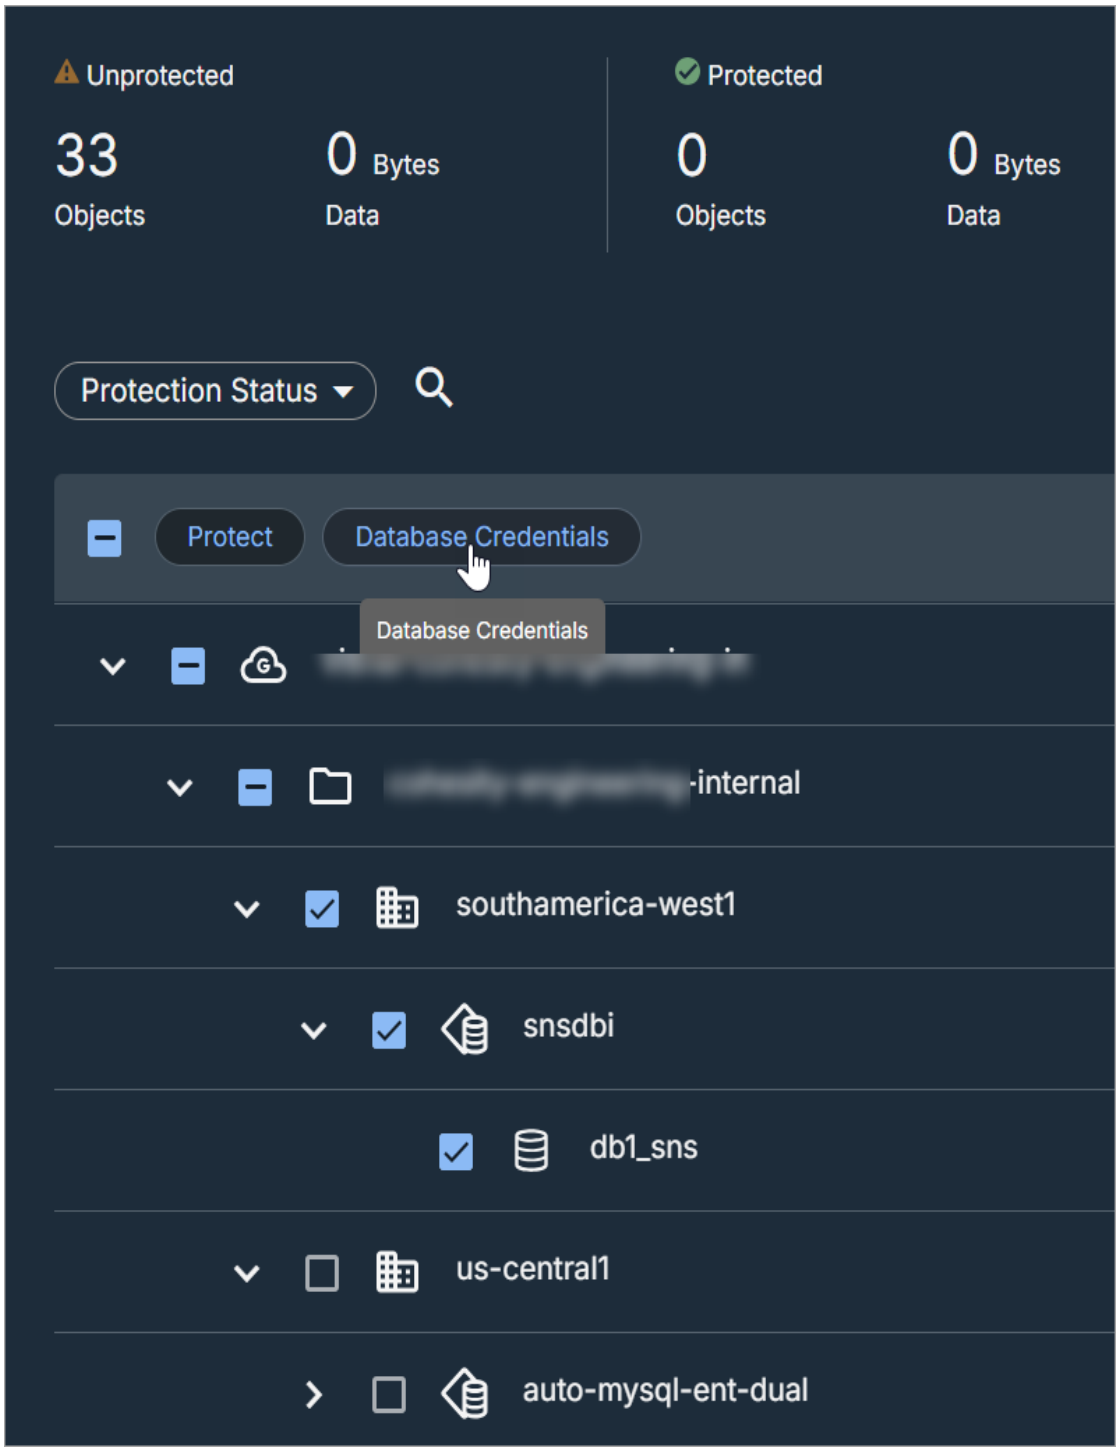

- Click the next to the required AlloyDB cluster.

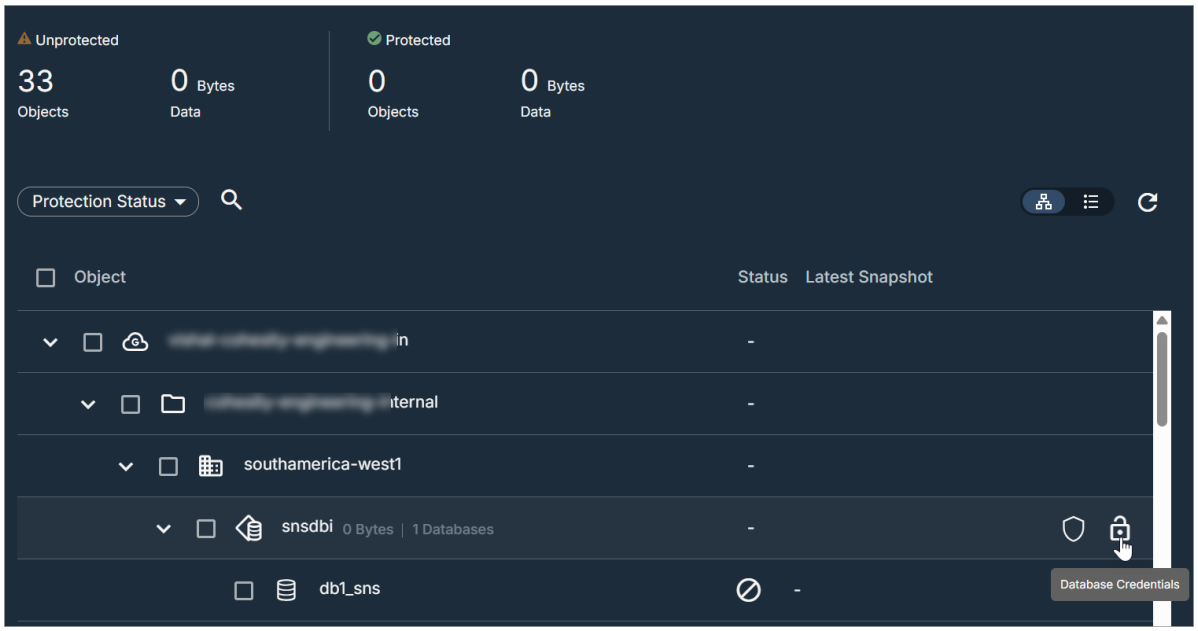

4. In the **Database Credentials** dialog box, enter the AlloyDB cluster **username** and **password**, and click **Save**.

Protect AlloyDB for PostgreSQL Databases

To protect AlloyDB for PostgreSQL databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **AlloyDB - PostgreSQL**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose databases using one of the following views:
   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the source, project, or region level.
   - **List View**: Select databases from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.

6. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** This option is not applicable for AlloyDB for PostgreSQL data protection . • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

7. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

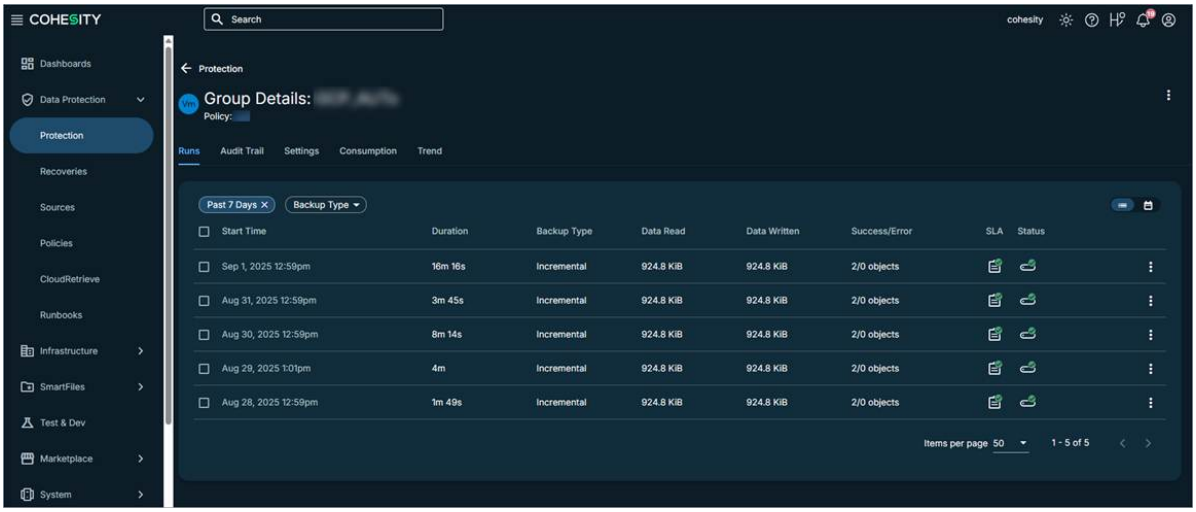

Manage AlloyDB for PostgreSQL Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for AlloyDB for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.

5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for AlloyDB for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for AlloyDB for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for AlloyDB for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
   - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover AlloyDB for PostgreSQL Databases

You can restore AlloyDB for PostgreSQL databases either to their original GCP source or to a different GCP source and region.

CloudSQL for AlloyDB PostgreSQL Recovery Process

When you run a recovery job for AlloyDB for PostgreSQL, Cohesity restores the database(s) specified in the recovery job. The Cohesity cluster connects to the selected Cloud SQL instance using the database credentials you provide after source registration and imports the backup data into the target instance.

Cohesity retrieves the backup from Cohesity-managed storage, transfers the backup files to the recovery target, and uses standard PostgreSQL restore utilities, such as `pg_restore`, to import the database contents into the target instance.

Prerequisites

To recover AlloyDB for PostgreSQL databases, add the database credentials for the target Cloud SQL AlloyDB PostgreSQL instance if you plan to restore it to a new location.

Recover AlloyDB for PostgreSQL Databases

To recover AlloyDB for PostgreSQL databases:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **AlloyDB - PostgreSQL**.
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
    - By default, the latest snapshot of the database is selected for recovery.
    - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <database_name>** page appears.
    - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

      To select a snapshot from the **List** mode:

        a. On the **List** tab, snapshots available in the past month are displayed by default.

        b. To view snapshots from a different range, select a date range from the **Date** drop-down list.

        c. Select the snapshot(s) to recover.

d. By default, the snapshot is recovered from the local Cohesity cluster.

e. Click **Select Recovery Point** and then go to step 5.

To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster.

d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the database to the same project from which it was backed up. **Note:** Restoring to the original location replaces any existing database with the same name. Type **YES** to confirm.
New Location	Restore the database to an alternate location in a different GCP source. Select a registered GCP source, project, and region. To overwrite an existing database with the same name, enable **Overwrite existing database(s)** and type **YES**. **Note:** Restoring to a new location also replaces any existing database with the same name in the selected region.

7. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

8. Click **Recover**.

Cloud SQL for SQL Server

Cohesity natively integrates with Google Cloud to provide a simple, fast, and cost-effective solution for backup, recovery, and data management of Google Cloud databases, including Cloud SQL for SQL Server.

> **Important:** The backup and recovery of Google Cloud SQL for SQL Server is supported only on Cohesity Cloud Edition on GCP.

This chapter includes the following topics:

Plan and Prepare for Cloud SQL for SQL Server Data Protection

To protect Cloud SQL for SQL Server databases with the Cohesity cluster, you must register your Google Cloud environment as a source in the Cohesity cluster and perform backup and recovery operations. The following table lists the requirements are grouped into those needed only for source registration and those needed for both backup and restore.

Requirement	When Needed	Details
Network Connectivity	Source registration, backup/restore	Ensure that you open the specified ports in the firewall to allow the Cohesity cluster to transmit and receive data.
Allowlisted Domains	Source registration, backup/restore	Allowlist the following FQDNs for source registration and data protection operations: • `https://accounts.google.com/` • `https://www.googleapis.com/` • `https://storage.googleapis.com` • `https://vpcaccess.googleapis.com/` • `https://cloudfunctions.googleapis.com/` • `https://cloudresourcemanager.googleapis.com/`

Requirement	When Needed	Details
Staging Bucket for Backup and Restore	Backup/Restore	You must have a staging bucket in Google Cloud Storage to enable backup and restore operations for Cloud SQL for SQL Server databases. • **Auto Create**: Cohesity can automatically create a staging bucket when you configure backup and recovery for Cloud SQL for SQL Server databases. • **Manual Create**: You can manually create one or more staging buckets in Google Cloud Storage. Use different staging buckets for protecting different databases or for recovery operations. Each bucket can be in the same or a different region as the database. **Note:** Using a bucket in a different region than the database may incur additional Google Cloud Storage costs.

Create an IAM Custom Role in GCP

To enable Cloud SQL for SQL Server data protection, you must create an IAM custom role in Google Cloud and assign it to the service account used by the Cohesity cluster.

For detailed instructions, refer to the following topics in the Google Cloud documentation:

- Create a custom IAM role
- Assign a role to a service account

> **Note:** Before you begin creating IAM custom role, ensure that you have created a dedicated service account in Google Cloud for the Cohesity cluster.

Create Staging Bucket on Google Cloud Storage

To enable Cloud SQL for SQL Server data protection with the Cohesity cluster, you need a staging bucket in Google Cloud Storage. The staging bucket is used to temporarily store exported Cloud SQL for SQL Server databases during backup and recovery operations.

You can choose one of the following options when setting up the staging bucket:

- **Auto Create**: Cohesity can automatically create a staging bucket for you as part of the backup configuration workflow in the Cohesity cluster.
- **Create Manually**: Manually create and specify an existing Google Cloud Storage bucket to be used for staging.

> **Note:** Create one or more staging buckets in Google Cloud Storage. You can use different staging buckets for protecting different databases or for recovery operations. Each bucket can be in the same or a different region as the database. Note that using a different region may incur additional Google Cloud Storage costs.

Manually Create a Staging Bucket

To manually create a staging bucket:

1. In the **Google Cloud** console, go to the Cloud Storage **Buckets** page.
2. Click **Create**.
3. On the **Create a bucket** page, provide the following information. After completing each section, click **Continue** to proceed:
    1. **Get started**: Enter a unique name for the staging bucket. Follow naming guidelines to ensure compatibility.
    2. **Choose where to store your data**: Select either **Region** or **Multi-region** as the location type.
    3. **Choose how to store your data**: Select **Set a default class**, then choose **Standard** as the default storage class.
    4. **Choose how to control access to objects**:
        - Enable **Enforce public access prevention**.
        - Select **Uniform as the access control model**.
    5. **Choose how to protect object data**:
        - Under **Data protection**, enable **Soft delete policy (For data recovery)** and select **Use default retention duration**.
        - Under **Data encryption**, select **Google-managed encryption key** as the encryption method.
    6. Click **Create**.

Considerations for Cloud SQL for SQL Server Data Protection

Review and understand the following limitations and known issues before you protect Cloud SQL for SQL Server databases:

- **Logical Backups Only**: Snapshot-based backups are not supported.
- **Backup Coverage**: Only user databases are supported. Backups of read-only databases or databases running in single-user mode are not supported.

- **Recovery Behavior**: Databases can be recovered only as new databases. Cross-region and cross-account recoveries are not supported.
- **Version Downgrade Not Supported**: Restoring a backup taken from a higher Cloud SQL for SQL Server database version to a lower version is not supported.
- **Cross-region charges**: Cross-region backup and recovery operations will incur additional charges as per Google Cloud Platform (GCP) pricing policies and may result in higher overall costs due to data transfer and storage in multiple regions.

Backup Cloud SQL for SQL Server Databases

After registering your Google Cloud source in the Cohesity cluster, it uses the provided service account credentials (JSON key or email and private key) to authenticate with Google Cloud.

Once authenticated, it automatically discovers all projects that the service account can access, as well as all supported workloads such as Cloud SQL for SQL Server databases within those projects. The discovered workloads appear under **Data Protection** > **Sources**. When you open the registered Google Cloud source, each workload type is displayed on its own tab. Select the **Cloud SQL for SQL Server** tab to view all the discovered databases.

The **Cloud SQL for SQL Server** tab displays the entity hierarchy for the Cloud SQL for SQL Server environment, allowing you to navigate from the project level down to individual databases.

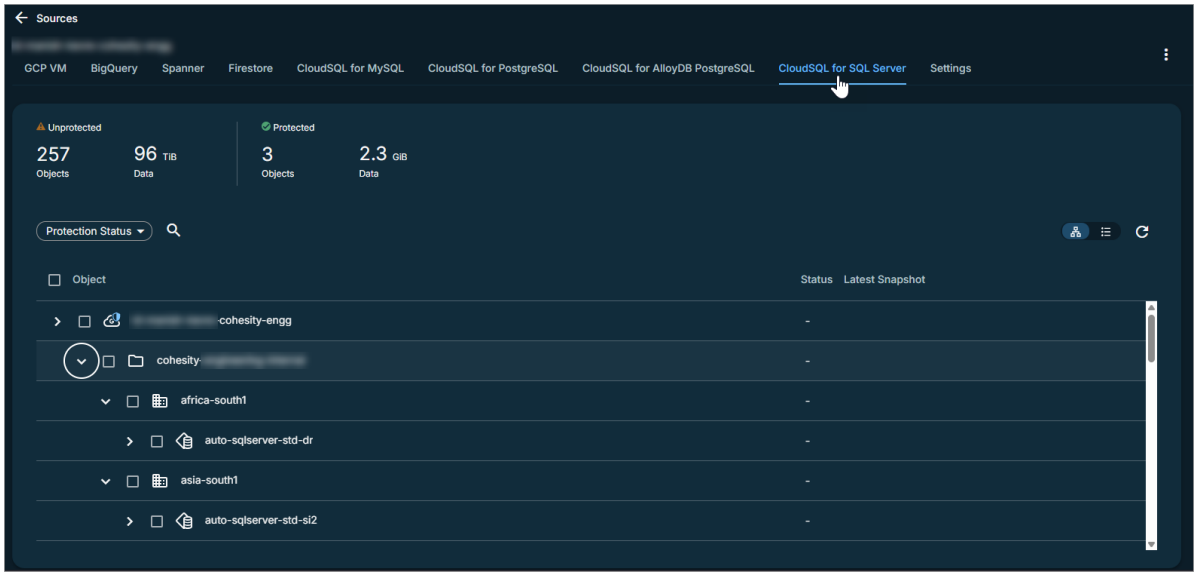

Cloud SQL for SQL Server Entity Hierarchy

The following hierarchy describes how Cloud SQL for SQL Server entities are organized and displayed in the Cohesity cluster:

Structure	Description
Service Account	The Google Cloud service account configured for the Cohesity cluster to perform Cloud SQL for SQL Server backup and recovery operations.
Project ID	The Google Cloud Project ID that the service account has permissions to access.
Region	The Google Cloud region where the Cloud SQL for SQL Server databases are hosted.
Instance	The Cloud SQL for SQL Server instance available for backup and recovery.
Database	The database that can be selected for backup and recovery.

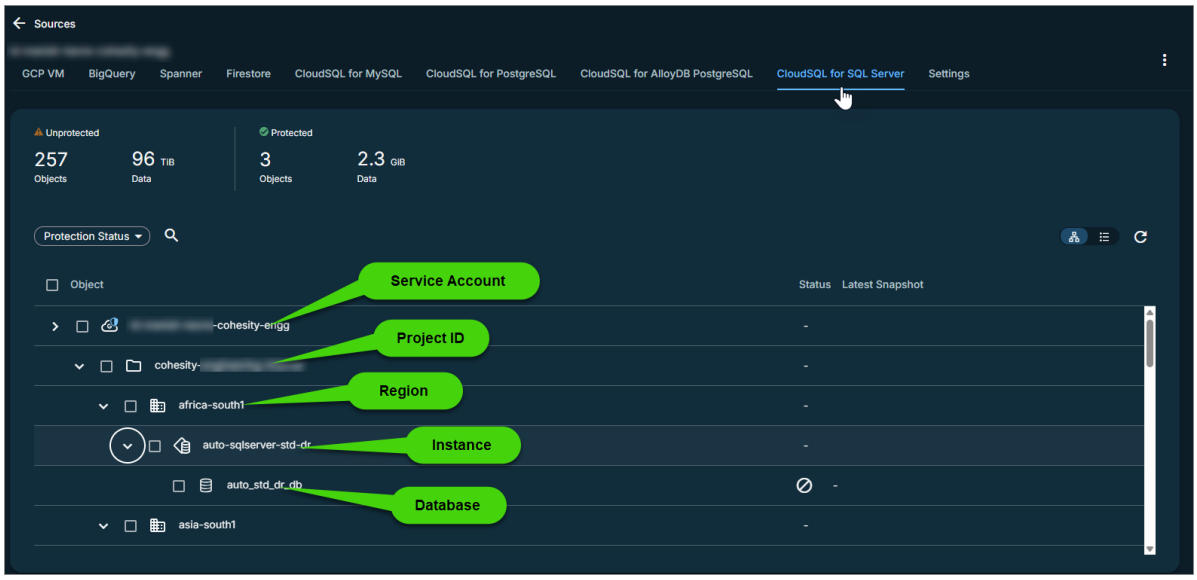

Backup Scope by Cloud SQL for SQL Server Hierarchy

Cohesity cluster allows you to back up Cloud SQL for SQL Server databases at different levels of the hierarchy. You can choose to back up:

- **Individual database**: Back up a single database.
- **Multiple databases within a region**: Back up selected databases from the same region.

- **Region level**: Back up all databases within a specific region.
- **Multiple databases across regions**: Back up selected databases from different regions.
- **Project level**: Back up all databases in all regions within the specified Project ID.

Cloud SQL for SQL Server Backup Process

When you run a protection job in the Cohesity cluster to back up Cloud SQL for SQL Server databases, the databases specified in the Protection Group are first exported from Cloud SQL for SQL Server to a staging bucket in Google Cloud Storage. The Cohesity cluster then copies the data from the staging bucket to the cluster. After the backup completes successfully, the contents of the staging bucket are deleted to free up storage.

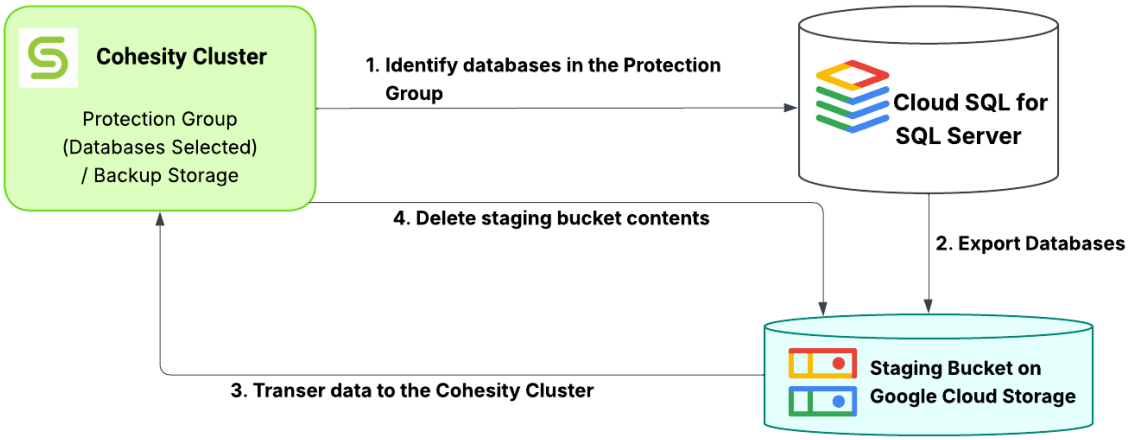

Supported Backup Types

Cohesity supports both full backup and incremental backups for Cloud SQL for SQL Server databases.

Protect Cloud SQL for SQL Server

In the Cohesity cluster, Google Cloud Cloud SQL for SQL Server databases are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Prerequisites

Before you protect Cloud SQL for SQL Server databases, ensure that:

- A staging bucket exists in Google Cloud Storage (GCS) for storing exported data during backups. You can:
  - If you choose to **Auto Create** the staging bucket:

    Cohesity automatically creates a staging bucket in the same region as your database(s) when you start a backup operation.
  - If you choose to **specify an existing bucket**:

    You can specify an existing bucket located in the same region as the databases being protected. Using a bucket in a different region is supported but may incur additional Google Cloud Storage charges.

  > **Note:** After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

- Add the database credentials required for the Cohesity cluster to authenticate and connect to the SQL server instance. These credentials allow the Cohesity cluster to establish a secure connection, perform the database dump, and complete backup and recovery operations.

Add Database Credentials

You must add the credentials at the Cloud SQL for SQL Server instance level. The same credentials are used whenever the Cohesity cluster connects to the database for backup or restore.

To add the database credentials:

1. Navigate to **Data Protection** > **Sources**, locate the registered Google Cloud source, and click on it.
2. Click the **Cloud SQL for SQL Server** tab.
3. Choose one of the following options:
   - Select the checkbox next to the Cloud SQL for SQL Server instance and click **Select All Child Objects**.

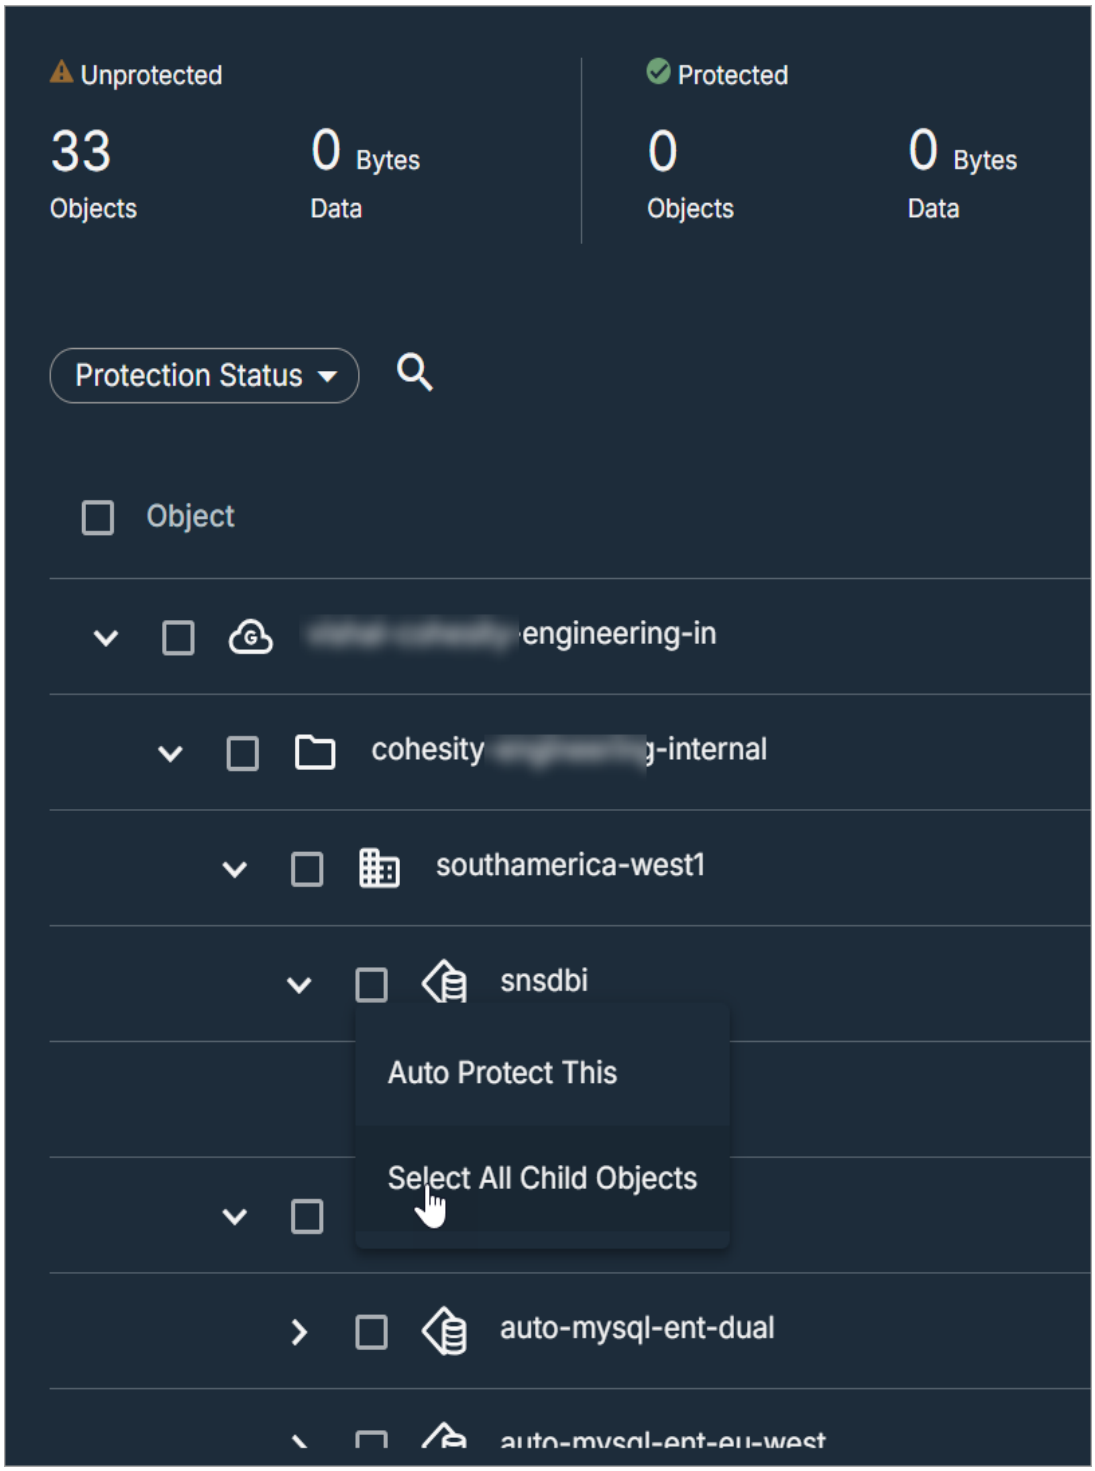

Then, click **Database Credentials** at the top left corner.

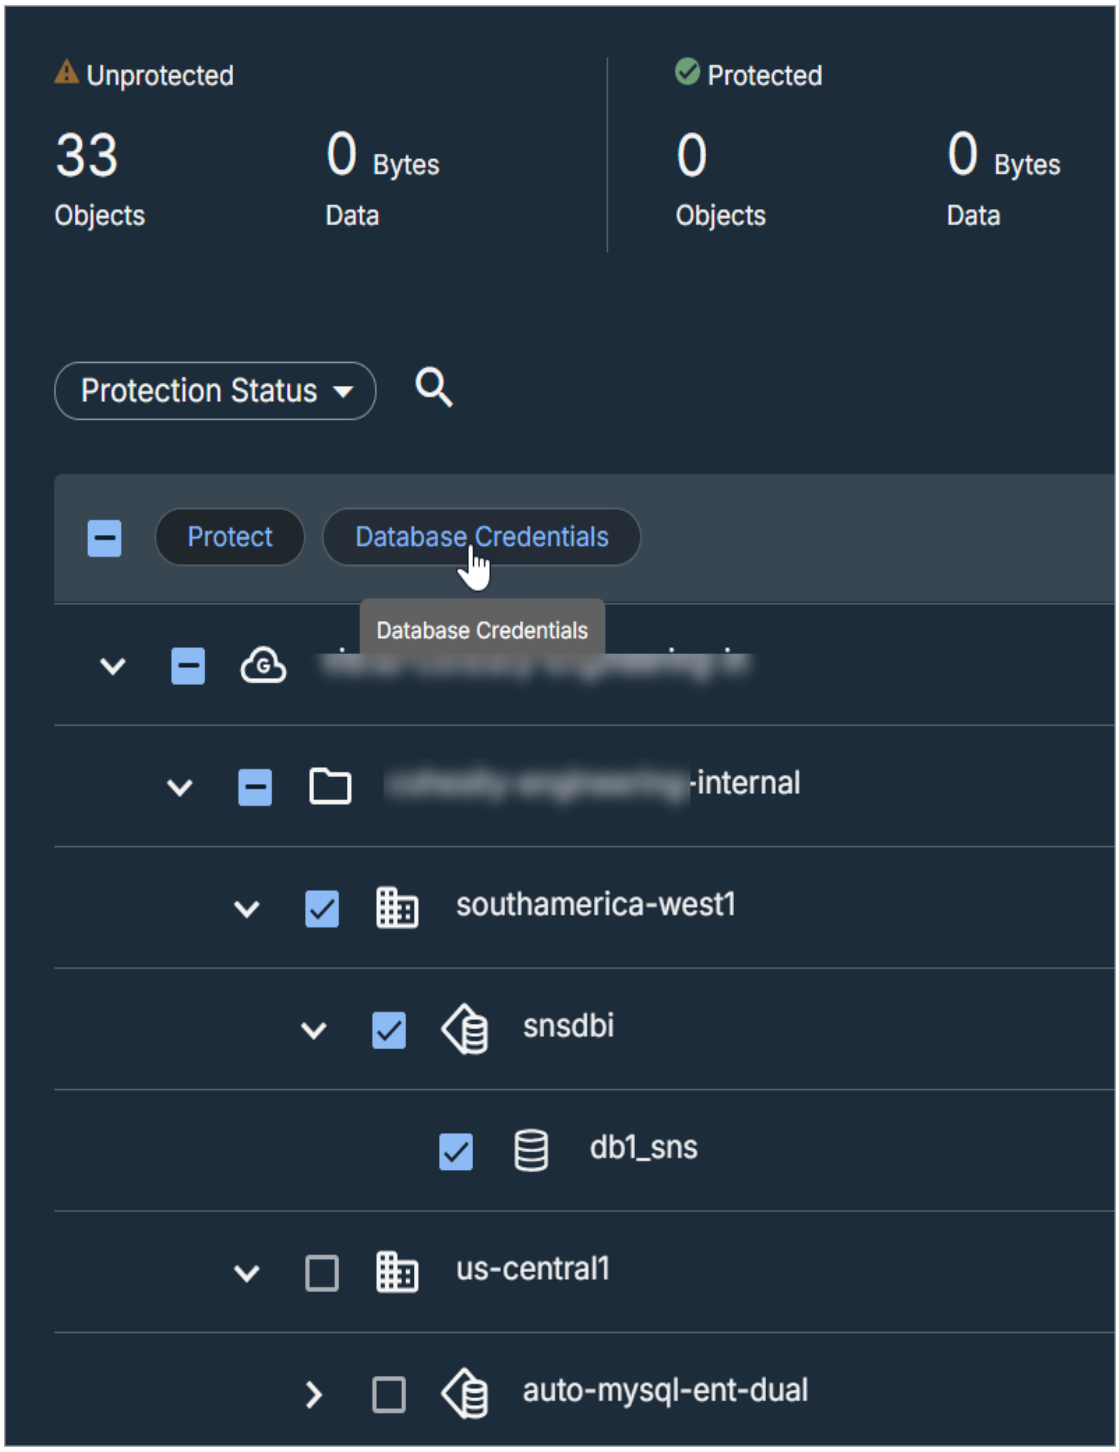

- Click the next to the required Cloud SQL for SQL Server instance.

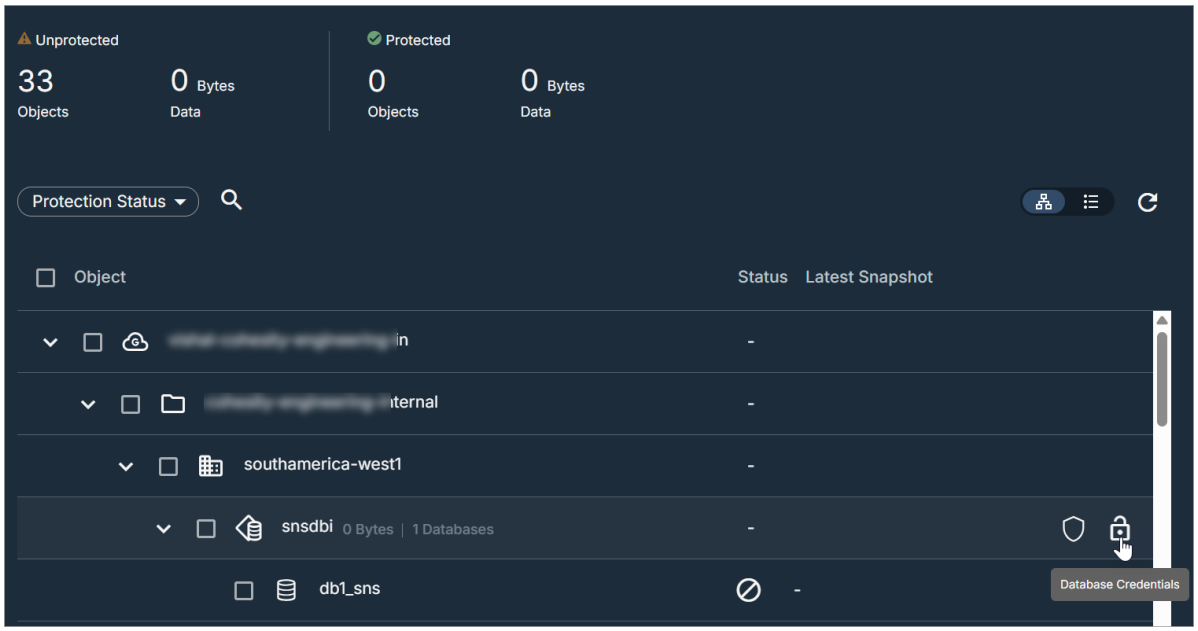

4. In the **Database Credentials** dialog box, enter the Cloud SQL for SQL Server instance **username** and **password**, and click **Save**.

Protect Cloud SQL for SQL Server Databases

To protect Cloud SQL for SQL Server databases:

1. Navigate to **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Google Cloud** > **Cloud SQL for SQL Server**.
2. From the **Source** drop-down list, select a registered Google Cloud source, or click **Register Source** to add a new one.
3. In the **Objects** section, choose databases using one of the following views:
   - **Hierarchy View (default)**: Browse objects by backup scope. You can enable **Auto Protect** at the source, project, or region level.
   - **List View**: Select databases from a flat, searchable list.
4. In the **Protection Name** field, enter a name from the Protection Group.
5. In the **Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your database(s). The bucket follows the naming format `cohe-<project-id >-<region-name>`.

     Where:
     - `project-id` is the Google Cloud project ID.
     - `region-name` is the Google Cloud region where the databases reside.

- **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for backups.

After each protection job run, Cohesity deletes only the temporary data stored in the staging bucket during backup operations. The staging bucket itself is retained and reused for future backup operations to minimize storage costs and avoid cross-region charges.

6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.
7. In the **Settings** section, configure the following:

Settings	Description
Storage Domain	Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
Start Time	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).

8. In the **Additional Settings** section, configure the following:

Additional Settings	Description
Pause Future Runs	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.

Additional Settings	Description
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. The Cohesity cluster supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	Optional. Enter a brief description about the Protection Group.

View Protection Group Details

After creating a Protection Group, you can view key details including:

- Status of protection runs
- Audit trails
- Storage consumption
- Configuration settings
- Protection run trends

To access these details, navigate to **Protection** and click the Protection Group you created.

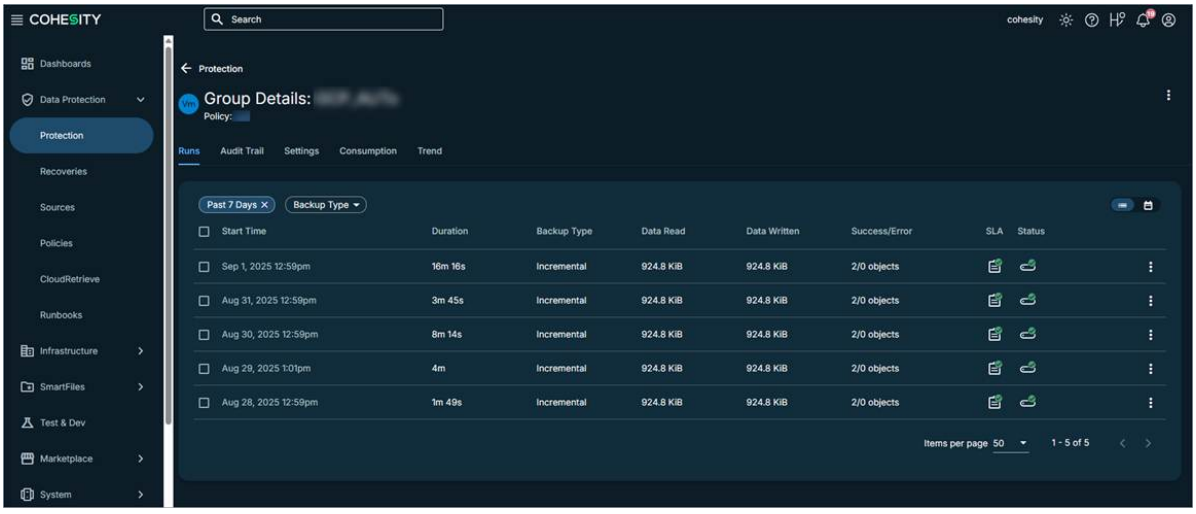

Manage Cloud SQL for SQL Server Backup Configuration

In the Cohesity cluster, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new object(s) to the existing Protection Group for backup. The newly added object is backed up from the next protection run.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for SQL Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for SQL Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for SQL Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all associated local and archived snapshots. Replicated snapshots must be deleted from the remote cluster. Snapshots under DataLock cannot be deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cloud SQL for SQL Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Keep the **Delete Snapshots** toggle disabled to delete the Protection Group only.
    - Enable the **Delete Snapshots** toggle to delete the Protection Group and all associated snapshots.
5. Type **YES** to confirm that you want to delete the Protection Group.
6. Click **Delete Group**.

Recover Cloud SQL for SQL Server Databases

You can restore Cloud SQL for SQL Server databases either to their original GCP source or to a different GCP source and region.

Cloud SQL for SQL Server Recovery Process

When you run a recovery job in the Cohesity cluster, one or more selected Cloud SQL for SQL databases backups are exported to a Google Cloud Storage staging bucket. The staging bucket temporarily stores the exported data before it is ingested into Cloud SQL for SQL deployment.

After the export completes, Google Cloud automatically copies the data from the staging bucket to the Cloud SQL for SQL deployment. Once ingestion is successful, the contents of the staging bucket are then automatically deleted to free up storage and reduce costs.

Prerequisites

Before you restore Cloud SQL for SQL Server databases, ensure the following:

- A staging bucket is available in Google Cloud Storage. You can either use an existing bucket or allow Cohesity to automatically create one during the restore process.
    - If you choose to **Auto Create** the staging bucket:

      Cohesity automatically creates a staging bucket in the same region as your database(s) when you start the recovery.
    - If you choose to **specify an existing bucket**:

You can specify an existing bucket located in the same region as the recovery target. Using a bucket in a different region is supported but may incur additional Google Cloud Storage charges.

> **Note:** After each recovery job run, Cohesity deletes only the temporary data stored in the staging bucket. The staging bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.

- **If recovering to a new location**: Ensure that you have added the database credentials for the Cloud SQL for SQL Server instance to which the data will be restored.

Recover Cloud SQL for SQL Server Databases

To recover Cloud SQL for SQL Server databases:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Google Cloud** > **Cloud SQL for SQL Server**.
3. Search for the database by entering characters of the database name, or by specifying the wildcard character *.
4. From the list of search results, select the database to recover.
   - By default, the latest snapshot of the database is selected for recovery.
   - To recover from a different snapshot, hover over the snapshot and click the Edit icon. The **Edit Recovery Point for <database_name>** page appears.
   - By default, snapshots are displayed in **List** mode. You can switch to **Timeline** mode by clicking the **Timeline** tab in the top-right corner of the page.

     To select a snapshot from the **List** mode:

     a. On the **List** tab, snapshots available in the past month are displayed by default.

     b. To view snapshots from a different range, select a date range from the **Date** drop-down list.

     c. Select the snapshot(s) to recover.

     d. By default, the snapshot is recovered from the local Cohesity cluster.

     e. Click **Select Recovery Point** and then go to step 5.

     To select a snapshot from the **Timeline** mode:

a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
b. Select the snapshot to be recovered by clicking the corresponding blue dot.
c. By default, the snapshot is recovered from the local Cohesity cluster.
d. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. In the **Recover To** section, select any one of the following options based on your requirements:

Option	Description
Original Location	Restore the database to the same project from which it was backed up. **Note:** Restoring to the original location replaces any existing database with the same name. Type **YES** to confirm.
New Location	Restore the database to an alternate location in a different GCP source. Select a registered GCP source, project, and region. To overwrite an existing database with the same name, enable **Overwrite existing database(s)** and type **Yes**. **Note:** Restoring to a new location also replaces any existing database with the same name in the selected region.

7. In **Restore Staging Bucket** section, choose one of the following options:
   - **Auto Create**: Cohesity automatically creates a temporary staging bucket in the same region as your database(s) for the recovery operation. The bucket follows the naming format: `cohe-<project-id>-<region-name>`.

     Where:
     - `project-id` is the Google Cloud project ID.
     - `region-name` is the Google Cloud region where the databases reside.
   - **Specify Bucket to Use**: Enter the name of an existing staging bucket in Google Cloud Storage (GCS) that you created for recovery.

   Cohesity deletes only the temporary data stored in the staging bucket during each recovery operation. The bucket itself is retained and reused for future recovery operations to minimize storage costs and avoid cross-region charges.

8. In the **Recovery Options** section, configure the following options:

Options	Description
Rename	Enter a new name for the database if you want to rename it after recovery.
Task Name	Update the default name of the recovery task.

9. Click **Recover**.

Chapter 11: S3-Compatible Storage

S3-compatible storage is a cloud storage system or service that supports the Amazon S3 (Simple Storage Service) API. Cohesity provides enterprise-grade backup and recovery of S3-compatible storage.

Plan and Prepare for S3-Compatible Storage Protection

Before you protect your S3-compatible storage, ensure you have met the prerequisites and reviewed the considerations.

Supported S3-Compatible Storage

For the list of supported S3-compatible Storage, see Supported Software.

Ports Used for Communication

Ensure the port used by the S3-compatible is open for communication with the Cohesity cluster.

S3-Compatible Storage Source	Port
Dell EMC ECS	9021
Pure Storage Flashblade	443
NetApp StorageGrid	443
IBM Cloud Object Storage	443
MinIO	9000

You need to mention these ports when registering the S3-compatible with the Cohesity cluster.

Required IAM Permissions

The user account used to register an S3-compatible storage with the Cohesity cluster must have an IAM policy attached with the following permissions.

Allow Permissions

These permissions are required for Cohesity to read metadata, list objects, and write backup data to the target storage.

Resource Type	Permission	Note
S3	s3: Get* s3: List* s3: AbortMultipartUpload s3: PutObject*	
	s3:PutOverwriteObject	Required only when registering NetApp StorageGRID as the S3-compatible storage.

Deny Permissions

To prevent accidental deletion or modification of the storage infrastructure, the following permissions must be explicitly denied:

Resource Type	Permission
S3	s3:DeleteObject* s3:DeleteBucket* s3:CreateBucket s3:PutBucket* s3:PutBucket*

Example

```
{
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:Get*",
```

```
                "s3:List*",
                "s3:AbortMultipartUpload",
                "s3:PutObject*"
            ],
            "Resource": "arn:aws:s3:::*"
        },
        {
            "Effect": "Deny",
            "Action": [
                "s3:DeleteObject*",
                "s3:DeleteBucket*",
                "s3:CreateBucket",
                "s3:PutBucket*"
            ],
            "Resource": "arn:aws:s3:::*"
        }
    ]
```

Considerations

- Cohesity does not support:
  - Auto-protection of S3-compatible bucket.
  - Backup and recovery of metadata and tags for S3-compatible buckets.
  - Recovery of S3-compatible bucket to AWS S3 bucket.
  - Browse and recover an object in an S3-compatible bucket. However, you can recover multiple objects by specifying the object prefix in the recovery task under the S3 Prefixes to Recover option.
  - The backup of older versions of the S3-compatible versioned bucket. Only the latest version of the versioned S3-compatible bucket is backed up.
- While creating the Protection Group for Dell EMC ECS, the UI displays only up to 1,000 buckets from the S3-compatible storage for selection.
- Cohesity supports CloudArchive Direct for S3-compatible storage configured with non-archive storage tiers.
- Cohesity supports all object stores that are compatible with the Amazon S3 API version dated 2006-03-01.

Register and Manage S3-Compatible Storage

To protect your S3-compatible storage bucket, you need to register your S3-compatible storage as a source on the Cohesity cluster. After registration, you can manage the source on the Cohesity cluster in the following ways:

- Update the S3-compatible storage source configuration.
- Refresh the source details.
- Unregister the S3-compatible storage source from the Cohesity cluster.

Before you Begin

Before you register S3-compatible storage as a source on the Cohesity cluster, ensure that you:

1. Review the requirements, considerations, and best practices.
2. Perform the prerequisite tasks.

For more information, see Plan and Prepare for S3-Compatible Storage Protection.

Register S3-compatible storage as a Source

Perform the following steps to register S3-compatible storage as a source:

1. In the product UI, select **Data Protection** > **Sources**.
2. Click Register on the top-right of the page and then select **S3 Compatible**.

   The Register S3-compatible Source page appears.
3. Provide the following information, and then click **Register**:
   - **Endpoint**: The IP address or Fully Qualified Domain Name (FQDN) of S3-compatible storage.

     > **Note:** Do not include 'http://' or 'https://'.
   - **Port**: The S3-compatible port that Cohesity communicates with.
   - **Access Key ID**: Access key ID of the user used to access the storage.
   - **Secret Access Key**: Secret access key address of the user to access the storage.

When registration finishes, a new entry appears on the Sources page. You can now create a Protection Group to back up S3-compatible buckets. For more information, see Protect S3-Compatible Storage.

Update the S3-Compatible Source Configuration

You can update the S3-compatible storage source details that you had provided during the registration process with the latest S3-compatible storage configuration.

To edit the S3-compatible storage source configuration:

1. Select **Data Protection** > **Sources**.
2. Click the actions menu next to the AWS source and select **Edit**.
3. In the **Edit S3 Compatible Storage Source** page, update the respective

configurations.

4. Click **Update**.

Refresh the S3-Compatible Storage Source

You can refresh the S3-compatible storage source configuration and fetch the latest changes on the S3-compatible storage source.

To refresh the AWS source configuration:

1. Select **Data Protection** > **Sources**.
2. Click the actions menu next to the S3-compatible storage source and select **Refresh**.

Unregister S3-Compatible Storage Source

If you plan to stop taking backups of the S3-compatible storage, you can unregister the S3-compatible storage source from the Cohesity cluster. Before you unregister, you must delete all the Protection Groups associated with the S3-compatible storage source. For more information, see Delete the Protection Group.

To unregister the S3-compatible storage source on the Cohesity cluster:

a. Select **Data Protection** > **Sources**.
b. Click the actions menu next to the S3-compatible storage source and select **Unregister**.
c. Click **Unregister**.

Protect S3-Compatible Storage

Once you register the S3-compatible storage source with the Cohesity cluster, you can create a Protection Group to back up the S3 buckets of the S3-compatible storage. A Protection Group uses the schedules and settings defined in the protection policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Policies.

Before you Begin

Before you protect an S3-compatible storage as a source on the Cohesity cluster, ensure that you:

1. Review the requirements, considerations, and best practices.
2. Perform the prerequisite tasks.
3. Register the S3-compatible storage as source with the Cohesity cluster.

Protect S3-Compatible Storage Buckets

Perform the following steps to protect the S3-compatible storage buckets:

1. Select **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then click **Protect** > **S3 Compatible**.

   The New Protection page appears.
3. Click **Add Objects**, and from the **Registered Source** drop-down list, select the S3-compatible source that contains the buckets you want to back up. You can also register a S3-compatible source with the Cohesity cluster so you can protect it.

   To register a new S3-compatible source, click Register Source from the drop-down. For more information, see Register and Manage S3-Compatible Storage.
4. From the buckets, select the buckets you want to back up.
5. Click **Continue**.
6. Provide a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. From the Policy drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
8. Select **More Options** to configure additional settings and continue with the next steps. Else, click **Protect**.
9. Optionally, from the **Protection Group** field, you can select an existing Protection Group. By selecting an existing Protection Group, the additional settings configured in the existing Protection Group will be applied to this Protection Group. To select an existing Protection Group, perform the following steps or else, go to step 10:
   a. From the **Protection Grou**p field, enable **Existing Group**.
   b. Select an existing Protection Group from the **Group** drop-down list.
   c. Click **Protect**.
10. Select the appropriate Storage Domain. For more information, see Create or Edit Storage Domains.
11. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
12. In the **SLA** field, select the expected time to complete either a full backup or incremental backup.

- **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
- **Full**: Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).

13. If you need to change any of the additional settings on the Protection Group page, click the down arrow icon next to the **Advance Settings** and click **Edit** of the respective field. For more information, see Additional Settings.
14. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: • High • Medium • Low

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Cancel Runs at Quiet Time Start	(Available only if the selected policy has at least one Quiet Time) When enabled, all the protection runs that are currently executing will cancel when the Quiet Time period starts. By default, this setting is disabled, meaning that after a protection run starts, it continues to execute even when a Quiet Time period starts. However, new protection runs will not start during a Quiet Time.
Alerts	Select one or more of the following Alerts On settings if you want alerts to be created for the following triggers: • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, seeAlerts. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **Success**: Create an informational alert when the Protection Group completes successfully. Emails are not sent when informational alerts are created. To send emails when alerts are triggered for the Protection Group, in the **Email Recipients** field, add the email addresses of the recipients.
Description	Enter a brief description about the Protection Group.

Policy Settings for S3-Compatible Storage

The following table lists the applicable policy settings for S3-compatible storage:

Field	Description
Primary Copy	Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster. Select any external target for performing CloudArchive Direct wherein the primary copy of the data is stored on the external targets. For more information, see CloudArchive Direct.
DataLock	If applicable, add a DataLock for compliance and regulatory requirements, to ensure that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users, including those who have the Data Security role in Cohesity, from modifying or deleting any Snapshots that were generated by the Protection Groups that use this Policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the Protection Group fails. The default time between retries is five minutes. You can customize the number of retries and how long to wait between each attempt.
Periodic full backup	Use this option to choose when to run a periodic full backup.
Quiet Times	Quiet Times defines time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturday or Sunday. Quiet Times only prevents new protection runs from starting during the specified time period. The protection runs that started before the quiet time period will be dependent on the quiet time settings configured in the additional settings of the protection Group. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	To replicate snapshots to another cluster, click the **Replication** icon from the floating menu. The copied snapshots are stored in the storage domain of the remote cluster. .
Archival	To add an archive schedule, click **archival** icon from the floating menu.

Use an Existing Protection Policy

You can use an existing Protection Policy protecting S3-compatible storage If you are selecting an existing protection policy, then the following options are ignored if they are already configured in the protection policy:

- Log Backup (Databases)
- Add BMR Backup (Physical Server)
- CloudSpin

Manage S3-Compatible Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

You cannot update the registered S3-compatible storage and the assigned Storage Domain in the Protection Group settings.

Add S3-Compatible Storage Buckets to the Protection Group

You can add new S3-compatible storage buckets to an existing Protection Group for backup. The newly added buckets are backed up from the next protection run, and it will be a full backup.

To add S3-compatible storage buckets to the Protection Group:

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you created for protecting the S3-compatible storage.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Select the S3-compatible storage buckets you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added S3-compatible storage buckets will be backed in the next scheduled protection run.

Remove S3-Compatible Storage Buckets from the Protection Group

From the Edit Protection page, you can remove S3-compatible storage buckets from the Protection Group that you do not want to back up in the next protection run.

To remove the S3-compatible storage buckets from the Protection Group:

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you created for protecting the S3-compatible storage buckets.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Unselect the S3-compatible storage buckets that you want to remove from the Protection Group.
6. Click **Save**.

The removed S3-compatible storage buckets will not be backed in the next scheduled protection run.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirement. In the Protection Policy, you can edit the settings listed in the Create or Use an Existing Protection Policy.

To the policy settings:

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for S3-compatible storage buckets.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the Edit Protection page, in the **Policy** field, click **Edit**.

   The Edit Policy page appears.
5. Edit the settings listed in the Create or Use an Existing Protection Policy and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you created for protecting the S3-compatible storage buckets.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the Edit Protection page, from the drop-down list in the policy field, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-Demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to backup S3-compatible storage buckets. It provides you options to back up all the objects or only selected objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- If configured with replication targets, you can select those targets.
- If configured with archival targets, you can select those targets.

To start a protection run:

1. Select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for protecting the S3-compatible storage buckets.
3. Click the action icon next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    a. In the Backup drop-down list, select any one of the following:
        - Backup all Objects in the Protection Group
        - Backup only selected Objects in the Protection Group
        - All objects that failed last run
    b. In the **Backup Type** drop-down list, select the backup type.
    c. In the **Replicate** section, select the replication target, the retention duration, and enable the **Replicate only fully successful runs** checkbox if you want to replicate the data only after a successful protection run.
    d. In the Archive section, select the archive target, the retention duration, and enable the **Archive only fully successful runs** checkbox if you want to archive the data only after a successful protection run.
5. Click **Run**.

Pause Future Protection Runs

You can stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you created for protecting the S3-compatible storage buckets.
3. Click the action icon next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run.**Pause Future Runs**.

To resume a paused protection run:

1. Select **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had paused.
3. Click the action icon next to the Protection Group and select **Resume**.

Delete the Protection Group

Cohesity provides the following ways to delete the Protection Group associated with the AS3-compatible storage buckets:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Protection Group are not deleted
- **Delete Protection Group and Snapshots**: Delete the Protection Group and all Snapshots, any local Snapshots captured by this Protection Group and stored on this cluster are deleted. In addition, any archived Snapshots stored on external targets that were captured by this Protection Group are also deleted.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for protecting the S3-compatible storage buckets.
3. Click the action icon next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group**? dialog box, select any one of the following options:
    - **Delete Protection Group Only**
    - **Delete Protection Group and Snapshots**

If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Group Run.

If the Protection Group is replicating to a remote cluster, the inactive job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the inactive job.

Recover S3-Compatible Storage Buckets

After you protect your S3-compatible storage, you can recover them to their original location or a new location using Cohesity.

Recover S3-Compatible Storage Buckets to the Original Location

To recover your protected S3-compatible storage buckets to their original location:

1. Navigate to **Data Protection** > **Recoveries**.

   Click **Recover** on the top-right of the page and then select **S3 Compatible** as the recovery source.

2. Search for a S3-compatible storage bucket or Protection Group that contains the snapshots to recover from by entering characters of the S3 bucket name, Protection Group name, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, or date.

3. From the list of search results, select an S3-compatible storage bucket you want to recover.

   By default, the latest snapshot is selected for recovery.

   To recover from a different snapshot, perform the following steps:

   a. From the selected lists, click the edit icon.

      The Edit recovery point for <S3bucket_name> page appears. By default, the page displays the latest snapshot available.

   b. Optionally, you can also select a snapshot from a different date by performing one of the following:

      - Selecting the date from the **Choose a date** drop-down list. Once you select the date, the Edit recovery point page displays the available snapshots on the specified date.
      - Clicking the **List** tab. The List tab displays all the available snapshots of the selected S3-compatible storage bucket.

   c. Select the snapshots to be recovered.

      By default, the snapshot is recovered from the local cluster.

   d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud icon if available in the **Location** option.

   e. Click **Select Recovery Point** and then go to step 5.

4. Click **Next: Recover Options**.

5. Under **Recover To**, select **Original Location**.

6. Enable **Overwrite Existing Bucket** if you want to recover the S3-compatible storage bucket by overwriting the original S3-compatible storage bucket. By default, this option is disabled.

   If the data in the source S3 object and target S3 are the same, then the target S3 will not be updated even if the Overwrite Existing Bucket option is enabled. If the target

bucket is versioned and the S3 object name of the target and the source match, then a new version of the S3 object will be created.

7. In the **Target Prefix** field, add prefix value to the names of the S3 objects that are recovered by this task.
8. Select your **Recovery Options**:
   - **S3 Prefixes to Recover**: To perform granular recovery of the S3-compatible bucket, enable the Turn ON to recover specific prefixes option and provide the prefix of the objects to be recovered.

     For example, `Finance/statement/document`

     You can click the add icon to provide multiple prefixes.

     The maximum number of prefixes that you can provide per recovery task is 100.
   - **Task** : Change the default name of the recovery task.
9. Click **Recover**.

Cohesity begins to restore the selected S3-compatible bucket.

Recover S3-Compatible Storage Buckets to New Location

To recover your protected S3-compatible storage buckets to a new location:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click Recover on the top-right of the page and then select S3-compatible as the recovery source.
3. Search for a S3 bucket or Protection Group that contains the snapshots to recover from by entering characters of the S3 bucket name, Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select an S3-compatible storage bucket you want to recover.

   By default, the latest snapshot is selected for recovery.

   To recover from a different snapshot, perform the following steps:

   a. From the selected lists, click the **edit** icon.

      The Edit recovery point for <S3bucket_name> page appears. By default, the page displays the latest snapshot available.

   b. Optionally, you can also select a snapshot from a different date by performing one of the following:

- Selecting the date from the **Choose a date** drop-down list. Once you select the date, the Edit recovery point page displays the available snapshots on the specified date.
- Clicking the **List** tab. The List tab displays all the available snapshots of the selected S3-compatible storage bucket.

c. Select the snapshots to be recovered.

By default, the snapshot is recovered from the local cluster.

Optionally, you can also select snapshots archived to the cloud by clicking the cloud icon if available in the **Location** option.

d. Click Select **Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. Under **Recover To**, select **New Location** and provide the following information:
   - **Register Source**: Select a registered S3-compatible source as the new recovery destination.
   - **Region**: Select a destination region.
   - **Bucket**: Select the S3-compatible bucket to which you want to recover.
7. Enable **Overwrite Existing Bucket** if you want to recover the S3-compatible bucket by overwriting the original S3-compatible bucket. By default, this option is disabled.

   If the data in the source S3 object and target S3 are the same, then the target S3 will not be updated even if the **Overwrite Existing Bucket** option is enabled. If the target bucket is versioned and the S3 object name of the target and the source match, then a new version of the S3 object will be created.
8. In the **Target Prefix fi**eld, add prefix value to the names of the S3 objects that are recovered by this task.
9. Select your **Recovery Options**:
   - **S3 Prefixes to Recover**: To perform granular recovery of the S3-compatible bucket, enable the Turn ON to recover specific prefixes option and provide the prefix of the objects to be recovered.

     For example, `Finance/statement/document`

     You can click the add icon to provide multiple prefixes.

     The maximum number of prefixes that you can provide per recovery task is 100.
   - **Task Name**: Change the default name of the recovery task.
10. Click **Recover**.

Cohesity begins to restore the selected S3-compatible bucket.

Chapter 12: Virtual Machines

Cohesity unifies fragmented data protection solutions for virtualized environments. With Cohesity, organizations no longer need to deal with complex and expensive protection solutions that result in multiple infrastructure silos and copies of data.

Cohesity supports backing up Virtual Machines (VMs) from different hosts in a single Protection Group. Protection Groups can run repetitively on a schedule.

To learn more about the relationship between Protection Groups and Policies, see Manage Protection Policy. For instructions on how to use the Cohesity Dashboard to back up VMs, see Manage Protection Groups.

For example, you could create a Protection Group with a protection schedule that backs up all critical VMware VMs every hour as represented by the *Critical* Protection Group in the following figure and a second Protection Group that backs up important systems every night as represented by the *Important* Protection Group in the following figure. Note that the Critical Protection Group is a mix of VMs from the Engineering ESXi Host and the IT ESXi Host.

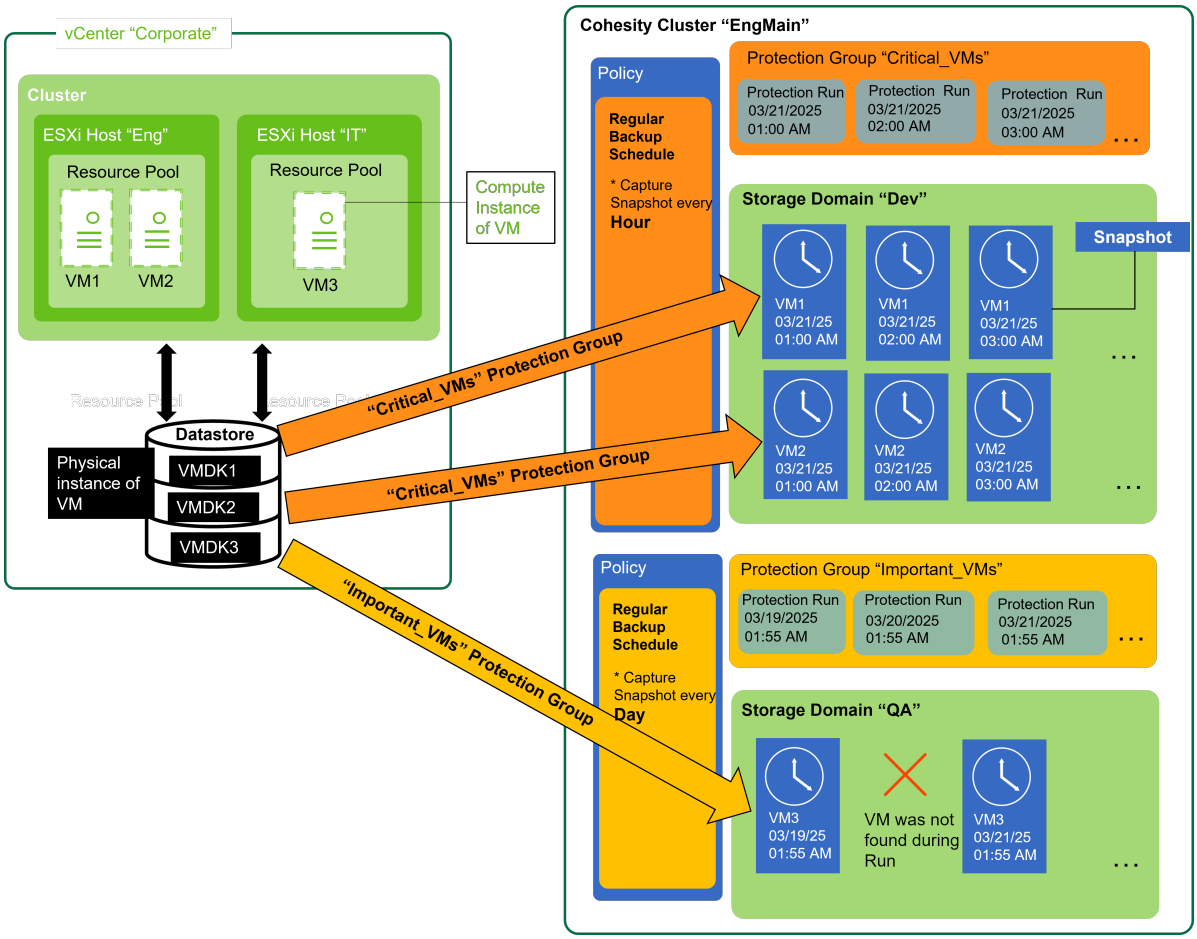

During a Protection Group Run, Snapshots of the VMs are saved on the Cohesity cluster with a date and time stamp. If a VM is unavailable when the Protection Group runs, a Snapshot is not created for that VM in the Protection Group as shown by the 03/20/2015 01:55 AM protection run in the Important Protection Group. In this example, a Snapshot of VMDK3 is not created by the 03/20/2015 01:55 AM protection run because VM3 was unavailable when the Protection Group ran. If the VM becomes available at a later date, the next time the Protection Group runs a Snapshot of the VM is included in the Protection Group as shown by the 03/21/2015 01:55 AM Protection Group Run for the Important Protection Group.

For instructions on how to use the Cohesity Dashboard to back up VMs, see Add or Edit a Protection Group for Virtual Servers.

About Encrypted VM Backup

Cohesity supports backing up VMware VMs encrypted by a VMware storage policy. A single Protection Group can back up both encrypted and unencrypted VMs.

The VMware user account designated for data protection requires additional privileges to back up encrypted VMs. For more information, see Role Privileges for vCenter Server.

A Protection Group performs the following actions:

1. Connects to the vCenter using an account with the proper privileges.
2. Takes a snapshot of encrypted VMs.
3. Decrypts and backs up data to the cluster.
4. Optionally, the backed-up data is encrypted.

The recovered VM from the backup will be decrypted. You must re-enable encryption after recovery by applying an appropriate storage policy.

Plan and Prepare for VM Protection

This topic provides information on minimum permissions, prerequisites and considerations to protect VMware VMs and Hypervisors with the Cohesity cluster.

Minimum Permissions

Ensure you have the required minimum permissions to register VMware VMs and Hypervisors as a source with the Cohesity, and to perform backup and recovery. For more information, see Ensure Adequate Privileges for Cohesity on the Source.

Support Matrix

Before you register your VMware VMs and Hypervisors with the Cohesity cluster, ensure that Cohesity supports the VMware VMs and Hypervisor versions. For more information, see Supported Software.

Ports Used for Communication

Ensure all necessary ports are open in the firewall between the Cohesity cluster and VMware VMs and Hypervisors. For more information, see Manage Firewall Ports.

Prerequisites and Considerations

Standalone ESXi Hosts

Review the following considerations:

- Throttling
- SQL backup, clone or restore
- File recovery to source
- Instant volume mount
- You must manually configure Storage vMotion.

VMware Snapshots

Prerequisite

Cohesity and VMware recommend limiting the number of delta files in a snapshot chain. For more information, see Best practices for virtual machine snapshots in the VMware environment (1025279). You may want to consider disk consolidation.

Considerations

The Cohesity cluster creates snapshots of VMware VMs. The limitations that apply to VMware snapshots also apply to Cohesity cluster snapshots. These limitations are described in VMware Snapshot limitations.

Review the following snapshot considerations:

- During Protection Group runs, the Cohesity cluster can report `The virtual machine requires manual consolidation` alerts when attempts to consolidate data in VMs fails due to VMware issues such as locking issues and timeouts. During the next run, the Cohesity cluster will attempt to consolidate the data in each VM but if the cluster is unable to consolidate the data in a VM, the cluster will continue to report `The virtual machine requires manual consolidation` alerts and skip backing up those VMs with this issue. However, the rest of the VMs in the Protection Group will be backed up. For the Cohesity cluster to start backing up the VMs with this issue, you must perform a manual workaround as described in the Unable to delete the virtual machine snapshot due to locked files (2017072) Knowledge Base issue. Cohesity recommends attempting consolidation by right-clicking the VM and selecting **Snapshot >Consolidate** as explained in the note of step#2. If that does not fix the problem, try restarting the management agents on the ESXi hosts where the VMs are running as described in step#5 and right-click the VM and select **Snapshot >Consolidate** as explained in the note of step#2.
- Registering a VMware source or protecting VMware VMs on Wide Area Networks (WAN) is not supported. Cohesity recommends using a Cohesity cluster (Virtual or Cloud Edition) with LAN connectivity to the vCenter and ESXi Host to register and protect the VMs.
- Registering a vCenter source and using it for Data Protection on multiple Cohesity clusters is not recommended. Cohesity clusters running in parallel may deplete vCenter's resources (such as API sessions, API calls, network bandwidth, computing power, memory, IO resources, number of tasks, and much more) as resource usage is not coordinated across Cohesity clusters. If you want to protect portions of your vCenter, Cohesity suggests registering the corresponding Standalone ESXi host(s) to the Cohesity clusters rather than registering the entire vCenter.
- Registering VMware vCenter or ESXi hosts through intermediaries is not supported.
- Creating snapshots of fault-tolerant VMs is not supported by VMware vSphere or the Cohesity cluster. Cohesity clusters will not display fault-tolerant VMs among the list of protectable VMs discovered from the vCenter.

- Creating snapshots of VMs with PCI vSphere DirectPath I/O devices is not supported by VMware vSphere or the Cohesity cluster.
- Creating snapshots of VMs with SR-IOV (Single Root I/O Virtualization) is not supported by the Cohesity cluster.
- When a VM with independent disks is protected, the independent disks are automatically skipped and the rest of the disks are backed up. After recovering or cloning these VMs from a Cohesity backup, you must manually attach the independent disks.
- Recovery of a VMware VM to a different vCenter does not attach tags of the original VM to the restored VM.
- If a Protection Group fails with the following error: "Failed with error: kVSphereError, detail: [Cannot complete the operation because the file or folder already exists]", see Unable to create snapshot: Operation failed because file already exists or Cannot complete the operation because the file or folder [DatastoreName] VMname/VMname.vmx already exists (308942) for information about a possible resolution.
- If the backup of a standalone ESXi host fails with one of the following errors:
  - "The operation is not allowed in the current state."

    This may be caused if ESXi is in maintenance mode.
  - "Backup task failed with error: [kInvalidLogin]: SOAP 1.1 fault: "":ServerFaultCode[no subcode] "Permission to perform this operation was denied." Detail:"

    This means the same ESXi host was registered on vCenter in "lockdown" mode.
- VMware folder-based protection is currently qualified with vCenter 6.7, 6.5, 6.0 and 5.5 only.
- If a Protection Group backs up multiple vApps, each vApp must be recovered individually.
- VMware does not support taking the snapshots or backups of the VMs having RDM (physical and virtual) on Shared Bus.
- Cohesity recommends increasing the default HTTPS connection limit after upgrading to vSphere 7.0 and later. For details, see the KB article.
- VMware Site Recovery Manager (SRM) creates placeholder VMs that are used for disaster recovery of the vCenter. Although Cohesity clusters display these placeholder VMs, the option to select or add them to a Protection Group is disabled.
- Review the following considerations when leveraging Pure Storage snapshots:
  - This feature is supported for vCenter only; it is not supported for standalone ESXi sources.

    - It is not supported for VMware VVols.
- Known VMware limitations when SAN Transport mode is enabled:
    - Protection of VMs on VVol datastores is unsupported.
    - VMware vSAN, a network-based storage solution with direct attached disks, does not support SAN transport.
    - Differential restore of VMware VMs is not supported.
    - Backing up VMware VMs using SAN transport mode cannot be leveraged if VMs are hosted on Pure ActiveCluster.
    - Backing up VMware VMs using SAN transport mode cannot be leveraged if VMs are hosted on IBM FlashSystem MetroMirror, GlobalMirror, HyperSwap, Policy-based Replication, or Policy-based HA systems.
    - VMware SAN transport mode is not supported with Cohesity virtual nodes. To use SAN transport mode for backing up VMware VMs, zoning and granting datastores access to the Cohesity nodes has to be done manually.

  For more information, see the Virtual Disk Development Kit Programming Guide.
- VMware - Storage Array Snapshot
    - VMware VM backup using Leverage Storage Snapshots is not supported for NetApp volumes enabled with Load-Sharing Mirror (LSM).
    - N_Port ID Virtualization (NPIV) Fibre Channel (FC) standard and Virtual World Wide Name (WWN) identifier are not supported when backing up VMware VMs using Leverage Storage Snapshots for Data Protection.

      > **Note:** This consideration is not applicable for IBM FlashSystem storage systems.

    - In Cohesity release 7.2 and later, Cohesity uses SAN connectivity (iSCSI or FC) with Storage Arrays (Pure, HPE Alletra, Nimble, and IBM) to transfer VMware VM backup data when the Leveraging Storage Snapshots for Data Protection option is enabled.
    - In Cohesity releases 7.2 and later, SAN connectivity between Cohesity nodes and the Storage Array (Pure, HPE, and IBM) is a prerequisite for VM backups. VM backups will fail without SAN connectivity.
    - If you are leveraging storage snapshots for VMware VM backups in releases prior to 7.2 and plan to upgrade the Cohesity cluster to release 7.2 or later, ensure that SAN connectivity between Cohesity nodes and the Storage Array (Pure, HPE Alletra, Nimble, and IBM) is in place before the upgrade to prevent VM backup failures.

- VMs provisioned on NetApp NFS datastores are backed up over NBDSSL through the temporary VM (tmp VM).
- Backing up VMware VMs using Storage Snapshot cannot be leveraged if VMs are hosted on Pure ActiveCluster.
- Backing up VMware VMs using Storage Snapshot cannot be leveraged if VMs are hosted on IBM FlashSystem MetroMirror, GlobalMirror, HyperSwap, Policy-based Replication, or Policy-based HA systems.
- Backing up VMware VMs using Storage Snapshot cannot be leveraged when VMs have disks in different storage arrays.

- Review the following considerations when leveraging the Nutanix AHV storage snapshots:
  - This feature is supported for vCenter only. It is not supported for standalone ESXi sources.
  - Cohesity does not support application-consistent backups of VMware VMs that leverage Nutanix as the storage snapshot provider.
  - Cohesity does not support differential restore of VMware VMs backed by leveraging the storage snapshots.
  - If the protection run of VMware VMs is switched from VADP-based backup to storage snapshot-based backup, then the Cohesity cluster will initiate a full backup.
  - While performing storage snapshot-based backup, if a VMware VM snapshot is present on the ESXi host, then Cohesity will perform VADP backup.
  - Cohesity does not support leveraging storage snapshots for backing up VM templates.

Change Block Tracking

Prerequisite

You must specify a VMware vCenter user that has adequate privileges on the VMware vCenter server. The role assigned to the specified VMware vCenter user should include privileges to turn on Change Block Tracking (CBT) and other privileges required by the Cohesity cluster.

> **Note:** When creating snapshots, the Cohesity cluster attempts to turn on the Change Block Tracking (CBT) to improve performance and reduce disk space usage.

Consideration

For older versions of ESXi 5.5, if Change Block Tracking (CBT) is enabled for a VM and the VM is moved (or migrated) using VMware Storage vMotion, the Change Block Tracking gets reset so when the Cohesity cluster creates the next backup of the VM, a full backup is

triggered instead of an incremental backup. For more information see Resetting Changed Block Tracking for VMware vSphere virtual machines (2139574).

Microsoft Hyper-V

Prerequisites

- Any one of the following:
  - Hyper-V version 8.0 or 9.0 VMs running on Windows Server 2019 managed by System Center Virtual Machine Manager (SCVMM) 2016 or 2019.
  - Hyper-V version 8.0 VMs running on Windows Server 2016 managed by SCVMM 2012 R2 or 2016.
  - Hyper-V version 5.0 VMs running on Windows Server 2012 R2 managed by SCVMM 2012 R2 or 2016.
  - Standalone Hyper-V 2019 or 2016 or 2012 R2 host or cluster (cluster can have mixed versions).
- If antivirus software is installed and running on a Hyper-V host, then you must configure several exclusions and options of the antivirus software for optimal operation of Hyper-V and the running virtual machines. For more information, see the Microsoft documentation.

Considerations

> **Note:** Hyper-V backups that use VSS snapshots may encounter issues if VSS Legacy Tracing is enabled on the applicable Hyper-V hosts. Disable VSS Legacy Tracing unless it is explicitly required. For more information, refer to the Microsoft documentation.

- Cohesity does not support the backup of Hyper-V VMs with shared disks.
- Deleting the most recent Hyper-V snapshot from the Cohesity cluster causes the object's next Protection Group run to be a full backup instead of incremental.
- Hyper-V 2012 R2 VMs can be protected by only one Cohesity cluster at a time. Backup by multiple clusters of a Hyper-V 2012 R2 VM results in a failure with the following error: [kAgentErrorNotRetryable]: [kInvalid]: FetchVolumeSnapshotInfo called for an incorrect snapshot set id.
- Performing VM instant recovery requires the Cohesity cluster to be joined to the same domain as the Hyper-V cluster.
- If an instant recovery with storage migration is in progress, you cannot restart the VM. The VM cannot be powered on until the storage migration completes.

- When canceling an in-progress instant VM recover task, the created Cohesity SMB internal view must be torn down after VM storage migration is complete.
- A Hyper-V VM that is recovered using instant recovery shows an "Unsupported VM Configuration" status in SCVMM during migration and after migration finishes. After VM migration finishes, you must manually refresh the recovered VM in SCVMM to get it out of this state and make it usable from SCVMM. A VM in the "Unsupported VM Configuration" state, however, can still be accessed directly from the Hyper-V host.
- CBT information on a Hyper-V 2012 R2 host does not persist after a reboot. The next backup after a reboot is full instead of incremental.
- Hyper-V 2012 R2 supports SMB that was already preconfigured to back up Hyper-V VMs using agentless VSS; array-based SMB is not supported. Only VMs residing on Windows Native SMB servers are supported.
- File-level recovery to a Hyper-V VM requires the following:
    - For Windows target VM:
        - Target VM belongs to Active Directory.
        - The Windows Remote Management (WinRM) service is enabled on the target VM: winrm quickconfig.
        - Remote powershell must be enabled from Hyper-V host to VM.
        - The Hyper-V Guest Service Interface is enabled on the target VM: Enable-VMIntegrationService -Name 'Guest Service Interface.
    - For Linux target VM:
        - Cohesity Linux agent must be installed on the target VM.
- Review the following considerations when performing instant volume mount to Hyper-V VMs:
    - Only Windows VMs are supported.
    - Dynamic Disks (LDM and LVM) are not supported.
    - The Bring Disks Online option requires the following:
        - VM must be part of Active Directory, the VM and the Hyper-V host must be in the same AD.
        - Users must execute "winrm quickconfig" to enable winrm on the target VM and remote powershell must be enabled from Hyper-V host to VM.
    - Instant volume mount and file level recovery from Gen 1 to Gen 2 type VMs is not supported.
    - If SCVMM is unregistered from the Cohesity cluster, ensure you tear down all instant volume mounts. Not tearing them down can prevent the VM from being backed up when the source is registered with a different Cohesity cluster.

    - Instant volume mounting Hyper-V 2012 R2 VMs without a SCSI controller is not supported. This is because Hyper-V disallows dynamically adding a SCSI controller, which is required to add the virtual disks.
    - On 2012 R2 VMs, if an instant volume mount disk is attached to a VM during a Protection Group run, that snapshot cannot be application-consistent. If this occurs, the event viewer may contain a VSS-catastrophic error or similar message.
- Standalone Hyper-V hosts and failover clusters support backing up VMs on SMB, but recovering to SMB volumes is not supported.
- If Hyper-V 2012 R2 VSS group snapshots are enabled, the following limitations apply:
    - All Hyper-V VMs on the same Hyper-V host must belong in the same Protection Group.
    - Each Protection Group can only include Hyper-V VMs belonging to the same Hyper-V host.
- Cloned VMs in a saved state prevent VM tear down.
- The Protection Group for Hyper-V VMs will fail if you select to auto-protect the VMs using more than one tag.
- When recovering Hyper-V VMs with Tags (including custom properties) to an alternate location, the restoration of Tags will be skipped. Tags are only restored when recovered to the original location.
- Consider the following when recovering a VM with a thin-provisioned disk otherwise, the restored VM will have full allocated disk size:
    - The backed-up VM has thin-provisioned disks.
    - The Recovery Type is Instant Recovery.
    - The power state of the VM is On.

Nutanix Acropolis

Prerequisites

- The data services IP address must be set in the AHV cluster.
- A Linux or a Windows agent is required for file or folder recovery.

Considerations

- AFS NFS shares are not supported.
- The instant recovery of Nutanix Cloud Clusters (NC2) AHV on AWS and Azure is not supported.

> **Note:** Volume Group backups are crash-consistent (Nutanix AHV does not support application-consistent VG snapshots).

Nutanix AHV Prism Central

Considerations

- If any of the Prism Elements have incorrect credentials, a warning will be logged for those credentials, but the other Prism Element sources will still be registered.
- You can recover VMs from one PC to another, as well as recover VMs within the same Prism Central instance, but on a different Prism Element (cluster).
- You cannot perform mass instant recovery across different clusters. All VMs must be recovered to the same cluster for mass instant recovery or done individually.

RHV/OLVM

Considerations

- Cohesity supports RHV only until version 6.8.1. While you can still protect your RHV sources, the Cohesity Support team will not provide assistance.
- VM tags are not supported.
- File recovery to a VM is not supported.
- Non-data storage domains display as recovery destinations when recovering to an alternate location.
- RHV/OLVM VM snapshots are not application consistent.

Troubleshooting

- Canceling a recover task may not clean up the VMs that were being recovered.

  Workaround - Delete the VMs manually.
- Original VM console settings are not recovered.

  Workaround - Change them manually after recovery. Some settings cannot be changed with the VM running and may require you to power off the VM or recover with a power off to change before booting.
- The VM object hierarchy may show non-VMs during job creation. These can be ignored.
- When searching for files/folders to recover, the UI shows their Protection Group as "Deleted Job." This can be ignored; the jobs have not been deleted.
- Cohesity does not support the backup of RHV/OLVM VMs that are using compressed `qcow2` disk images.

Workaround - Decompress the `qcow2` images using the `qemu-img convert` command.

Register or Edit a Hypervisor Source

After a source is added (registered), its objects are eligible to be protected or backed up on the Cohesity cluster. For example, if a VMware vCenter Server is registered as a source, its VMs are eligible to be backed up.

To register a source to back up:

1. Navigate to **Data Protection > Sources**.
2. Select **Register** > **Virtual Machines**.
3. (*Applies only in multi-tenant environment with Data Source Connection enabled*.) From the **Data Source Connection** drop-down, select a Data Source Connection, and then click **Continue**.
4. Select a **Hypervisor Source Type** from the drop-down.
5. Enter the information for your source type.

   VMware: vCenter or Standalone ESXi Host

   > **Note:** VMware supports asynchronous source registration, allowing you to register a VMware environment with minimal information. This method is beneficial for large environments because it speeds up the registration process and performs object discovery in the background.

   > **Tip:** You can also watch the Registering a VMware Source video to learn how to register a VMware vCenter Server or a Standalone ESXi host as a source.

Field Name	Description
Hostname or IP Address	Specify the hostname or IP address (IPv4 or IPv6) of the source that contains the VMs to protect.
Username	You must specify a VMware vCenter user that has adequate privileges on the VMware vCenter server. The role assigned to the specified VMware vCenter user should include privileges to turn on Change Block Tracking (CBT) and other privileges required by the Cohesity cluster. When specifying a Windows Active Directory username, use the following format: <**domain**>\<**username**> or <**username**>@domain. For all other cases, just enter: <**username**>.

Field Name	Description
Password	Enter the password for the specified user.
CA Certificate	For VMware vCenter only. Toggle this setting on if you want Cohesity to validate the connection to vCenter for operations. Copy and paste certificate content in the field.

Field Name	Description
Latency Based Adaptive Throttling	For VMware vCenter only. Toggle this setting on if you want the Cohesity cluster to adaptively throttle the running of tasks created by the Protection Groups protecting this source. This throttling helps to ensure that the datastores on the source (vCenter Server) are not adversely affected by the Protection Groups. You specify latency thresholds and if the observed latency on the datastore is greater than the specified threshold, the Cohesity cluster throttles the processing of the backup using that datastore as explained below. **Note:** • Enabling throttling can increase how much time a Protection Group run takes to complete, resulting in SLA violations. You may want to increase the SLA for the Protection Group. For more information, see Manage Protection Groups. • Latency-based adaptive throttling is ineffective while leveraging storage array snapshots for VMware backups. Specify the latency thresholds in milliseconds. If you enable throttling, the following two fields appear: • **Latency Threshold for throttling the new tasks of Protection Group Runs**–If the observed latency of a datastore is greater than the specified latency threshold, the Cohesity cluster throttles the processing of new backup tasks using that datastore. Throttling can occur in the middle of a protection run. For example, tasks creating the first snapshots of the protection run could complete without throttling but after the threshold is reached, the tasks creating the remaining Snapshots in the protection run are throttled. • **Latency Threshold for throttling the currently running tasks of Protection Group Runs**–If the observed latency of a datastore is greater than the specified latency threshold, the Cohesity cluster throttles any currently running backup tasks using that datastore. If throttling is enabled, the vCenter user specified for the source must have administrative datastore privileges. When throttling is enabled and the vCenter user has the necessary administrative datastore privileges, the Cohesity cluster enables the collection of Storage IO Control (SIOC) statistics on the datastores if not already enabled. The Cohesity cluster uses the statistics from SIOC to calculate the observed latency of datastores. If throttling and the collection of SIOC statistics were initially enabled and then throttling is disabled, the Cohesity cluster does not change the SIOC collection statistic settings on the datastores.

Field Name	Description
Cap concurrent streams per Datastore	Toggle this setting on if you want to limit the number of streams concurrently backed up per datastore. This helps to reduce any performance impact on the datastore. When registering a vCenter, you can configure the **Number of Streams** for all datastores in the vCenter only. After a vCenter is registered and the datastores in the vCenter are discovered, you can add overrides for individual datastores by clicking **Add Datastore Override**. Select a datastore, toggle on **Cap** and enter the number of **Streams**. If you toggle off **Cap** for a datastore, it is not subject to the vCenter's global concurrent streams cap.
Cap concurrent VM Backups	Toggle this setting on if you want to limit the number of VMs of a vCenter to be backed up concurrently. In the **Number of VM Backups** field, you can set the number of VMs that can be backed up concurrently.
Auto Cancel Backups if Datastore is running low on space	Turn this toggle ON if you want to set the minimum free space that should always be available in the datastore. Once set, Cohesity automatically prevents backups from starting if the datastore does not meet this requirement. Enter the minimum free space needed in the **GiB** or **%** fields. For example, if you enter 10 in the **%** field, 10% of the datastore space will always be available. The percentage option is more flexible because it auto-adjusts when the datastore shrinks or expands, even if a VM has multiple datastores. Which means, if two virtual disks of a VM reside in datastores of size 500 GB and 1 TB each, 10% of free space will apply to both the datastores.

Field Name	Description
Set Last Backup detail	For VMware vCenter only: Toggle this setting ON to set the last backup time and backup status for a VM in vCenter. These are published as custom attributes. By default, the option is turned OFF. When a backup is performed on a vCenter with the feature enabled, the custom attributes "Last Cohesity Backup Attempt Time" and "Last Cohesity Backup Status" will be set or updated to reflect the backup time and status. • **Last Cohesity Backup Attempt** - The exact timestamp when the last backup attempt succeeded. • **Last Cohesity Backup Status** - Status of the last backup for each VM. **Note:** When performing VM-based backups, if the backup is successful but has warnings caused by implicit fallbacks such as failure to utilize storage snapshots or issues with deleting the VM snapshot (which is a time-consuming operation), we will set the last backup status as 'Success' since the VM was successfully backed up. However, if there are errors or warnings due to explicit fallbacks (for example, transitioning from SAN to NBDSSL, or from app consistent to crash consistent), we will update the backup status to accurately reflect the error or warning.
Use VM BIOS UUID to track virtual machines	Toggle this setting on if you want the Cohesity cluster to track VMs using a unique BIOS UUID. You must enable the BIOS UUIDs toggle to support integration with Zerto–a continuous data protection (CDP) product. In addition, you must configure Zerto to preserve the BIOS UUIDs upon recovery. If both products are using BIOS UUIDS to uniquely identify VMs, VMs can be restored by either of the products and the link back to the original VM is not lost. To protect a VM in the vCenter with this toggle on, all VMs must have a unique BIOS UUID. If the Cohesity cluster finds more than one VM with the same BIOS UUIDs, these VMs are not listed when selecting VMs to protect. **Important:** This option must be set when you initially register a vCenter. After the vCenter has been registered, you cannot edit the hypervisor source and change this setting. To configure Zerto to also use BIOS UUIDs, select **Site Settings**. Under **Policies > Replication** select **For incoming replication, copy the BIOS UUID of the protected VM to the recovered VM**.

Field Name	Description
Network for Data Transfer	(Applies only to VMware vCenter) Toggle this option on if you want to configure a secondary network for the backup and copy recovery of VMs on the ESXi host of the VMware vCenter, while also using NBDSSL as the transport mode. By default, this option is turned off. If this option is turned off, the default network will be used for backup and recovery. **Configure a Secondary Network** To configure the secondary network, specify the subnet IP in the **Subnet IP** field and **CIDR prefix** in the CIDR prefix field. When the ESXi host has multiple VMkernel adapters, then Cohesity cluster will use the IP address that belongs to the specified subnet for the backup and recovery of VMs on the ESXi host. To view the updated ESXi hostname to IP address mapping, from the Cohesity Dashboard, select **Settings** > **Cluster** > **Networking** and click the **Host Mapping** tab. > **Note:** When ESXi hosts are registered to VMware vCenter using IP addresses instead of hostnames, Host Mapping does not work. **Add Multiple Subnets** To add multiple subnets, click the **Add** icon, and then specify the subnet IP and CIDR prefix. If you add more than one subnet, then the ESXi IP address selected for host mapping will be the one which matches the first subnet in the order they are provided. Example 1: Consider a VMware vCenter having two ESXi hosts and each ESXi has two VM kernel adapters. ESXi-1: • VM kernel adapter 1: 111.255.190.57 • VM kernel adapter 2: 10.10.33.130 ESXi-2: • VM kernel adapter 1: 111.255.190.51 • VM kernel adapter 2: 10.10.33.120 Consider you have provided the subnet IP as 10.10.33.0 and CIDR as 24. During the backup and restore of VMs under both the ESXi hosts, Cohesity cluster will resolve the hostname of ESXi-1 and ESXi-2 to IP address 10.10.33.130 and 10.10.33.120 respectively since these IP addresses belong to the same subnet you have provided.

Field Name	Description
	Example 2: Consider a VMware vCenter having two ESXi hosts and each ESXi has two VM kernel adapters. ESXi-1: • VM kernel adapter 1: 111.255.190.57 • VM kernel adapter 2: 10.10.33.130 ESXi-2: • VM kernel adapter 1: 111.255.190.51 • VM kernel adapter 2: 10.10.148.40 Consider you have provided the subnet information in the order listed below: 1. Subnet IP: 10.10.33.0 CIDR: 24 2. Subnet IP: 10.10.148.0 CIDR: 24 During the backup and restore of VMs under ESXi-1, the Cohesity cluster will resolve the hostname of ESXi-1 to IP address - 10.10.33.130 since the IP address belongs to 10.10.33.0/24 subnet. Similarly, during the backup and restore of VMs under ESXi-2, the Cohesity cluster will resolve the hostname of ESXi-2 to IP address - 10.10.148.40 since the IP address belongs to 10.10.148.0/24 subnet. **Delete an Added Subnet** To delete an added subnet, click the **cross** icon next to the subnet. When an ESXi host is added or deleted from a VMware vCenter, the corresponding host mapping will be updated during the next source refresh.

Field Name	Description
Detect VM migration across vCenter to Preserve backup chain	Toggle this option on if you want your Cohesity cluster to detect the VMs that are migrated from the current VMware vCenter to another VMware vCenter. With this option enabled: • The VMware vCenter (the VMware vCenter to which the VMs are migrated) hierarchy view (**Data Protection** > **Sources** > <VMware source> **All Objects**) in the Cohesity cluster will display a tag with the label, vCenter Migrated next to the migrated VMs. • When a Protection Group is run on the VMware vCenter to which the VMs are migrated, the Cohesity cluster eliminates the need for a full backup of the migrated VMs as the VMs are protected before the migration. Only an incremental backup is needed. > **Note:** If you are enabling this option to an existing VMware vCenter source, then you must refresh (**Data Protection** > **Sources** > <VMware source> **Action menu**>**Refresh**) the source for this feature to function. For more information, see Detect VM migration across vCenter to Preserve backup chain.

VMware: vCloud Director

Field Name	Description
Hostname or IP Address	Specify the hostname or IP address of the vCloud Director that contains the vCenters to protect.
Username	Specify the local administrator user. The minimum privileges needed to register vCD as a source are: • vCD user • vCenter user • User role that is used within the vCD Extension to start the mapping between the vCD organization and the Cohesity organization • User role that is used within the vCD Extension to complete the mapping between the vCD organization and the Cohesity organization
Password	Enter the password for the specified user.
vCenter Details:	

Field Name	Description
Username	You must specify a VMware vCenter user that has adequate privileges on the VMware vCenter server. The role assigned to the specified VMware vCenter user should include privileges to turn on Change Block Tracking (CBT) and other privileges required by the Cohesity cluster. When specifying a Windows Active Directory username, use the following format: <**domain**>\<**username**> or <**username**>@domain. For all other cases, just enter: <**username**>.
Password	Enter the password for the specified user.
Detect VM migration across vCenter to Preserve backup chain	Toggle this option on if you want your Cohesity cluster to detect the VMs that are migrated from the current VMware vCenter to vCD. With this option enabled: • The vCD (the vCD to which the VMs are migrated) hierarchy view (**Data Protection** > **Sources** > <VMware source> **All Objects**) in the Cohesity cluster will display a tag with the label, vCenter Migrated next to the migrated VMs. • When a Protection Group is run on the vCD to which the VMs are migrated, the Cohesity cluster eliminates the need for a full backup of the migrated VMs as the VMs are protected before the migration. Only an incremental backup is needed. **Note:** If you are enabling this option to an existing vCD source, then you must refresh (**Data Protection** > **Sources** > <VMware source> **Action menu**>**Refresh**) the source for this feature to function. For more information, see Detect VM migration across vCenter to Preserve backup chain.

Hyper-V: SCVMM Server, Standalone Host or Failover Cluster

> **Important:** Before registering SCVMM, you must install the Cohesity agent on the SCVMM server or on an existing proxy endpoint that is connected to an SCVMM server. Registration will fail without the agent installed. Ensure you install the Cohesity agent on any Hyper-V hosts that you want Cohesity to protect. The Cohesity agent version cannot be higher than the Cohesity cluster version. For more information, see Agent Compatibility Within a Cluster. For agent installation steps, see Download and Install the Agent on the SCVMM Server and Download and Install the Agent on Hyper-V Hosts. If a Hyper-V host is managed by SCVMM, do not register that host as a standalone host with the Cohesity cluster. Failover clusters require additional privileges. For more information, see Microsoft Hyper-V.

Field Name	Description
Hostname or IP Address	Specify the FQDN of the SCVMM cluster. Otherwise, you can specify the hostname or IP address of the source that contains the VMs to protect.
Username	Specify a domain user's account that has been granted administrative access to the SCVMM application.
Password	Enter the password for the specified user.
The Cohesity Agent has a different hostname or IP address than this SCVMM	If toggled off, the Cohesity agent is installed on the SCVMM server. If toggled on, the Cohesity agent is installed on an existing proxy endpoint that is connected to an SCVMM server. Specify the **Agent Hostname or IP Address**. The proxy endpoint must already have the VMM Administrator Console installed.
Cap concurrent streams per Datastore	Toggle this setting on if you want to limit the number of streams concurrently backed up per datastore. This helps to reduce any performance impact on the datastore. When registering a Hyper-V SCVMM Server, you can configure the **Number of Streams** for all datastores in the Hyper-V. After a Hyper-V is registered and the datastores in the Hyper-V are discovered, you can add overrides for individual datastores by clicking **Add Datastore Override**. Select a datastore, toggle on **Cap** and enter the number of Streams. If you toggle off **Cap** for a datastore, it is not subject to the Hyper-V's global concurrent streams cap.

Ensure you open necessary firewall ports as described in Manage Firewall Ports.

When editing registration (such as moving the SCVMM proxy agent endpoint to a different proxy or to the SCVMM cluster; or moving from the SCVMM cluster to a proxy), perform the following steps:

1. Copy the old agent registry values from the Cohesity folder and subfolders.
2. When moving to a proxy, paste the registry values to the new endpoint.

   When moving to SCVMM, paste the registry values to the active SCVMM node (RDPing into the SCVMM cluster redirects to the active master node).

Nutanix AHV Prism Central

Field Name	Description
Hostname or IP Address	Specify the hostname or IP address of Nutanix Prism Central.
Prism Central Username	Specify a Prism username that has adequate privileges to perform actions on the source. For more information, see Nutanix Acropolis.
Prism Central Password	Enter the password for the specified user.
Cap concurrent streams per Datastore	Toggle this setting on if you want to limit the number of streams concurrently backed up per datastore. This helps to reduce any performance impact on the datastore. When registering a Prism Central, you can configure the **Number of Streams** for all datastores in the vCenter only. After a Prism Central is registered and the datastores in the Prism Central are discovered, you can add overrides for individual datastores by clicking **Add Datastore Override**. Select a datastore, toggle on **Cap** and enter the number of **Streams**. If you toggle off **Cap** for a datastore, it is not subject to the Prism Central's global concurrent streams cap.
Prism Element Username	(*Optional*) **Note:** If the same user account is used for both Prism Central and Prism Element credentials, you do not need to manually re-enter the Prism Element credentials. Specify a Prism element username that has adequate privileges to perform actions on the source. It is used to access the Nutanix Prism element.

Field Name	Description
Prism Element Password	(*Optional*) Enter the password for the specified user.

Configure Data Services IP for iSCSI Transport

You can configure a segmented Data Services IP address for iSCSI transport at both the Prism Element and Prism Central levels to control how backup and recovery traffic flows between the Cohesity cluster and AHV sources. This enhancement provides greater flexibility in separating management and data transport networks for backup and recovery operations.

This option is available after the AHV source is registered and can be configured from the **Sources** > **Edit** page.

To Configure Data Services IP for iSCSI transport:

Prerequisite

Ensure that the Nutanix AHV sources (Prism Central and/or Prism Element) are already registered on the Cohesity cluster.

1. Navigate to **Data Protection** > **Sources**.
2. Locate the registered AHV source (Prism Central or Prism Element).
3. Click the Actions (⋮) menu for the source and select **Edit**.
4. On the Edit Source page, select either **Use iSCSI Data Services IP** (default) or **Use segmented iSCSI Data Services IP**.
5. For a Prism Central source, select the Prism Elements that should use the segmented Data Services IP, and leave others unselected to continue using the default IP.
6. For a Prism Element source, select the required option, **Use iSCSI Data Services IP** (default) or **Use segmented iSCSI Data Services IP**.
7. Enter the required credentials if prompted.
8. Click **Save**.

   During backup, recovery, and copy operations, the selected Data Services IP is used for iSCSI transport. For Prism Central, the transport selection applies at the Prism Element level. Backup and recovery operations record the transport IP used in the audit logs.

Acropolis: Acropolis Standalone Cluster and Nutanix Cloud Clusters AHV on AWS and Azure

To register, follow the steps mentioned above for both Acropolis Standalone Cluster and Nutanix Cloud Clusters (NC2) AHV on AWS.

After you complete the iSCSI configuration and have an iSCSI target set/enabled on the source, verify the iSCSI connectivity between the Cohesity cluster and the source. You can manually log in to one of the configured iSCSI targets from one of the Cohesity cluster nodes.

Field Name	Description
Hostname or IP Address	Specify the hostname or IP address of Nutanix Prism.
Username	Specify a Prism username that has adequate privileges to perform actions on the source. For more information, see Nutanix Acropolis.
Password	Enter the password for the specified user.

GCP

For information on registering GCP, see Register or Edit a GCP Source.

AWS

For information on registering AWS, see Register AWS Source.

Azure

For information about required permissions, see Azure Permissions.

Field Name	Description
Category	Select one of the following options: • **Standard**–Standard Azure. • **Gov**–Azure Government. • **Stack**–Azure Stack Hub. You must specify the authentication type, region, and the domain name. For more information, see Azure Stack Hub.
Subscription ID	This is the subscription ID from which you want to back up data. Log in to the Azure portal. In the left panel, click **Subscriptions**. Copy the SUBSCRIPTION ID from the table.

Field Name	Description
Application ID	The Application ID assigned by Azure during the service principal creation process. For instructions see Get application ID and authentication key in the Microsoft Azure documentation.
Application Key	The Azure Application key generated using Legacy App Registration during the service principal creation process that is used for authentication. For instructions see Get application ID and authentication key in the Microsoft Azure documentation.
Tenant ID	The unique tenant ID assigned by Azure. For instructions see Get tenant ID in the Microsoft Azure documentation. **Note:** The **Tenant ID** field is not displayed for Azure Stack Hub Active Directory Federation Services (AD FS) authentication type.
Stack Type	If you select Stack, select the authentication type: • Azure Active Directory • Active Directory Federation Services Azure Stack Hub uses either Azure Active Directory or AD FS for authentication. AD FS is mostly used for disconnected deployments of Azure Stack Hub. For more information, see Microsoft Azure Stack Hub documentation.
Region	If you select **Stack**, specify the Azure Stack Hub region. For instructions, see Region management in Azure Stack Hub in the Microsoft Azure documentation.
Domain Name	If you select **Stack**, specify the Azure Stack Hub domain name.

RHV: oVirt Manager

> **Note:** OLVM VMs can be protected by configuring the oVirt Manager as an RHV: oVirt Manager host.

> **Note:** When you register the oVirt-Manager with a Cohesity cluster, ensure you install the Cohesity agent on the RHV/OLVM hosts that are joined to the oVirt Manager. The Cohesity agent version cannot be higher than the Cohesity cluster version. For more information, see Agent Compatibility Within a Cluster. For agent installation steps, see Download and Install the Cohesity Agent on RHV/OLVM Hosts.

To register the oVirt-Manager with a different Cohesity cluster:

1. Unregister the current Cohesity cluster from the oVirt-Manager.

   > **Note:** If you skip this step, protection runs fail in the current Cohesity cluster.

2. Uninstall and reinstall the Cohesity agent on each RHV/OLVM host.
3. Register the oVirt-Manager with the new Cohesity cluster.

Field Name	Description
Hostname or IP Address	Specify the RHV/OLVM oVirt Manager hostname or IP address.
Username	Specify a username that has adequate privileges (Cohesity recommends a user with the "admin" role) to perform actions on the source. Use the format: username@profile, for example, admin@internal.
Password	Enter the password for the specified user.

6. Click **Register**.

   When source registration finishes, a new entry appears in the sources table.

Manage Application Registration on a VM

After you register applications with the Cohesity cluster, they can be protected. Currently, MS SQL Server and MS Exchange are supported.

For information about MS SQL, see Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs.

For information about MS Exchange, see Register Exchange on a VM.

Edit Application Registration

1. Navigate to **Data Protection > Sources**.
2. Navigate the source hierarchy to the server, mouse over its actions menu and click **Edit Application Registrations**.
3. Edit the desired information.
4. Click **Save**.

Unregister Applications

You can unregister MS SQL Server only if it is not currently protected by an MS SQL Server Protection Group.

1. Navigate to **Data Protection > Sources**.
2. Navigate the source hierarchy to the server, mouse over its actions menu and click **Unregister All Applications**.
3. Click **Unregister**.

Register a Hyper-V Source

To register and back up a Microsoft Hyper-V source, follow the steps described below.

> **Note:**
> - Cohesity does not support the backup of Hyper-V VMs with shared disks.
> - If antivirus software is installed and running on a Hyper-V host, you must configure several exclusions and options for optimal agent operation. For more information, see Antivirus Software Exclusions.

Download and Install the Cohesity Agent on the SCVMM Server

Before registering SCVMM, you must install the Cohesity agent on the SCVMM server or on an existing proxy endpoint that is connected to an SCVMM server.

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges, run the executable and complete the installation wizard.

   Install the agent without additional components.
4. The agent starts automatically. Now you can register the Hyper-V SCVMM server as a source.

Download and Install the Cohesity Agent on Hyper-V Hosts

Ensure you install the Cohesity agent on any Hyper-V hosts that you want Cohesity to protect.

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate hosts.
3. As an administrator with local system privileges, run the executable and complete the installation wizard.

   For Hyper-V 2016 servers, install the agent without additional components.

   For incremental backups of individual VMs hosted by Hyper-V on Windows Server 2012 R2, install the **File System CBT (Changed Block Tracker)** component.
4. The agent starts automatically.

Guest Windows VM minimum recommended specifications: 2 GB RAM and the equivalent of a 1 GHz processor.

Best practice: For Hyper-V 2016, configure all VMs' Automatic Stop Action to shut down or turn off, instead of save. This results in all powered-on VMs having minimal size .vmrs files. Because Cohesity does not back up .vmrs files that are greater than 10 MB (.vmrs file size, if set to save, scales to the amount of memory assigned to the VM). VMs in the saved state, by definition, generally have .vmrs files greater than 10 MB. This Cohesity Windows Agent registry value entry determines the maximum size of the .vmrs file supported per backup: agent_hyperv_max_vm_config_file_size_mb. Set it to zero to back up .vmrs files of all sizes.

Register a Hyper-V Source with the Cohesity cluster

1. Navigate to **Data Protection > Sources**.
2. Select **Register** > **Virtual Machines**.
3. From the **Select Hypervisor Source Type** drop-down, select the hypervisor source type and complete the fields. The available options are:
   - HyperV: SCVMM Server
   - HyperV: Standalone Host
   - HyperV: Failover Cluster

   For detailed registration steps, see Register or Edit a Hypervisor Source.
1. Click **Register**.

   After registration is completed, ensure the Cohesity cluster is joined to an Active Directory domain as described in Join Active Directory, which is required for Hyper-V VM backup and recovery.

   Now Cohesity can back up any Hyper-V hosts where the Cohesity agent is installed.

User Account Requirements for VM Protection

Specific user account configuration is required for the following operations:

- The user account that is used for auto-deploying agents in a Windows VMware environment must have administrator privileges on the VM where the agent is deployed. The account must have the "sysadmin" Server Role.

  If the Windows Firewall is used, add an inbound rule to accept the Cohesity Agent and TCP connections on local port 50051.
  Port 50051 may have limited requirements. It is needed for application integration, but not required for File-Level Recovery or Instant Volume Mount operations.

If the operation fails with an error indicating insufficient privileges, then check if the user account used for this operation allows remote command execution without a UAC prompt. Otherwise, use a persistent agent or other suitable recovery methods, if applicable.

Register a Storage Snapshot Provider

Cohesity can leverage the snapshots from the following storage snapshot providers:

- Cisco HyperFlex
- Nimble Storage, HPE Alletra 5000, and/or HPE Alletra 6000
- NetApp ONTAP
- Nutanix
- Pure Storage Array
- IBM FlashSystem

You can leverage snapshots only for VMware backups from the preceding storage snapshot providers. To leverage the storage snapshots, you must:

1. Register the snapshot providers with the Cohesity cluster. See the following sections on how to register the above snapshot providers.
2. Enable the **Leverage Storage Snapshots for Data Protection** field and select the storage snapshot provider from the drop-down when creating the Protection Group. For more information, see Add or Edit a Protection Group for Virtual Servers.

Considerations

- You can leverage only one storage snapshot provider for a Protection Group.
- For IBM FlashSystem, Nimble Storage, HPE Alletra 5000, HPE Alletra 6000 and Pure Storage, VMFS datastores are supported.
- For NetApp, NFS 3 and NFS 4.1 datastores are supported.
- Cohesity cannot leverage storage snapshots when VM has disks on volumes of different storage systems.

- The Cohesity cluster periodically queries the registered NetApp source(s) to obtain volume information. The VM protection job does not query NetApp for volume information at the time it is run. If new NetApp volumes were presented to the vCenter as NFS datastores and VMs from such datastores are part of a protection job, you must manually refresh the NetApp source on the Cohesity cluster or wait for the periodic refresh to ensure volume information is updated for leveraging this integration.
- The Cohesity cluster periodically queries the inventory of the vCenter and the Nutanix cluster. The information from these two sources may be out of sync. Therefore, if new VMs have been added or VMs have been moved, you may need to manually refresh the vCenter and Nutanix source on the Cohesity cluster to leverage Nutanix snapshots for VMware backup.
- The leverage storage snapshot feature doesn't work when multitenancy is enabled.
- To complete the storage snapshot, you must register the source with the same IP address or FQDN that is specified in vCenter for the datastore.

Register Cisco HyperFlex

To register Cisco HyperFlex as a storage snapshot provider for VMware, follow the steps described below.

1. Navigate to **Data Protection > Sources**.
2. Click **Register > Storage Snapshot Provider**.
3. From the **Select Snapshot Storage Provider Type** drop-down, select **Storage Snapshot Provider: Hyperflex**.
4. Enter the following information:

Field	Description
Hostname or IP Address	Specify the hostname or IP address of the Cisco HyperFlex system to protect. Cohesity recommends using the HyperFlex Management IP address (SCVM floating IP address). **Note:** You must also register the corresponding VMware vCenter with the Cohesity cluster.
Username	Specify a Cisco HyperFlex admin user to ensure adequate privileges.
Password	Enter the password for the specified user.

5. Click **Register**.

   When registration finishes, a new entry appears in the Sources table.

After registering the snapshot provider, to leverage the VMware snapshots from Cisco HyperFlex, you must enable the **Leverage Storage Snapshots for Data Protection** field and select **HyperFlex** when creating the virtual Server Protection Group. For more information, see create a Protection Group for VMware.

Register Nimble Storage Array, HPE Alletra 5000, and/or HPE Alletra 6000

> **Note:**
> - Registering HPE Alletra 5000 and/or HPE Alletra 6000 as a source is an Early Access feature.
> - You can register Nimble Storage, HPE Alletra 5000, and/or HPE Alletra 6000 for leveraging storage snapshots for VMDKs. Cohesity does not support the backup of Nimble Storage, HPE Alletra 5000, and/or HPE Alletra 6000 storage volumes.

Before registering the Nimble Storage Array as a storage provider, ensure it meets firewall port requirements.

Registering a Nimble Storage as a Storage array makes its LUN (both iSCSI and FC) content eligible to be protected on the Cohesity cluster.

After you complete the iSCSI configuration and have an iSCSI discovery set/enabled on the Nimble array, verify the iSCSI connectivity between the Cohesity cluster and the Nimble array. To verify the connectivity, you can manually log in to one of the configured iSCSI targets on the Nimble array from any one of the Cohesity cluster nodes.

Run the following command to check the iSCSI connectivity:

```
iscsiadm -m discovery -t sendtargets -p <ipaddress>
```

For additional iSCSI commands, see Additional iSCSI commands.

Perform the following steps to register Nimble Storage, HPE Alletra 5000, and/or HPE Alletra 6000 Storage Array:

1. Navigate to **Data Protection > Sources**.
2. Click **Register > SAN**.
3. Select the source type as **Nimble**.

> **Note:** You can register HPE Alletra 5000 and HPE Alletra 6000 by selecting the source type as Nimble.

4. Enter the following information:

> **Note:** You must also register the corresponding VMware vCenter where Storage Array Datastore is mounted with the Cohesity cluster.

Field	Description
Management Host or IP Address	Specify the Management Host or IP Address of the Nimble or HPE Alletra 5000 and HPE Alletra 6000 Storage to register.
Username	Specify the user name used for accessing the Nimble or HPE Alletra 5000 and HPE Alletra 6000 Storage.
Password	Enter the password for the specified user.

5. Click **Register**.

   When registration finishes, a new entry appears in the Sources table.

Additional iSCSI Commands

The table below lists the additional iSCSI commands:

Command	Description
iscsiadm -m node -T <targetname> -p <ipaddress> -l	Use this command to log in to a specific target.
iscsiadm -m node -T <targetname> -p <ipaddress> -u	Use this command to log out from a specific target.
iscsiadm -m node -T< targetname> -p <ipaddress>	Use this command to display the target information.
iscsiadm -m node -s -T <targetname> -p <ipaddress>	Use this command to display the statistics of a target.
iscsiadm -m session	Use this command to display all the current logged-in sessions.

Command	Description
iscsiadm -m discovery -o show	Use this command to display the iSCSI database regarding discovery.
iscsiadm -m node -o show	Use this command to display the targets to log into.
iscsiadm -m session -o show	Use this command to display the sessions logged into.

After registering the snapshot provider, to leverage the VMware snapshots from Nimble Storage, HPE Alletra 5000, and/or HPE Alletra 6000, you must enable the **Leverage Storage Snapshots for Data Protection** field and select **Storage Array/NAS** when creating the virtual Server Protection Group. For more information, see create a Protection Group for VMware.

Using the **Leverage Storage Snapshots for Data Protection** option with the **Leverage SAN Transport for Data Protection** option is not recommended.

Register NetApp ONTAP

To register NetApp ONTAP as a source as a storage snapshot provider for VMware, follow the steps described below.

Minimum Permissions

Before you register NetApp ONTAP cluster or SVM with Cohesity cluster as a storage snapshot provider for VMware, you must have the required permissions:

Permission for registering NetApp ONTAP cluster as a storage snapshot provider for VMware:

Access-level	Command	Description
All	vserver export-policy	Adds the Cohesity cluster subnet to the export-policy so that Cohesity can mount volumes.
	volume snapshot	Allows fetching, creating, and deleting snapshots for volumes.
	volume	Fetches information about volumes and deletes flex clone.
ReadOnly	cluster identity	Fetches information about the cluster.
	network interface	Fetches information about network interfaces to which Cohesity connects for mounting volumes.
	vserver	Fetches information about Vservers.

Permission for registering NetApp ONTAP SVM as a storage snapshot provider for VMware:

Access-level	Command	Description
All	vserver export-policy	Adds the Cohesity subnet to the export policy so that the Cohesity cluster can mount volumes.
	volume snapshot	Allows fetching, creating, and deleting snapshots for volumes.
	volume	Fetches information about volumes and deletes flex clone.
ReadOnly	network interface	Fetches information about network interfaces to which the Cohesity cluster connects for mounting volumes.
	vserver	Fetches information about Vservers.

To register NetApp ONTAP:

1. Select **Data Protection > Sources**.
2. Click **Register > NAS**.

   The Register NAS page appears.
3. From the **NAS Source** drop-down list, select **NetApp**.
4. Enter the following information.

Field Name	Description
NetApp Source Type	The NetApp ONTAP cluster or Vserver to register.
Host or IP Address	The Host or IP address of the NetApp ONTAP cluster or Vserver to register.
Username	Username used to access the NetApp ONTAP cluster or Vserver. **Note:** The username used to register the NetApp ONTAP cluster and Vserver is case sensitive.
Password	Enter the password for the specified user.

Field Name	Description
Back Up SMB Volumes	Not applicable. This field is not applicable for registering NetApp ONTAP as a storage snapshot provider
Filter IPs	Not applicable. This field is not applicable for registering NetApp ONTAP as a storage snapshot provider.

5. Click **Register**.

   When registration finishes, a new entry appears in the Sources table.

After registering the snapshot provider, to leverage the VMware snapshots from NetApp ONTAP, you must enable the **Leverage Storage Snapshots for Data Protection** field and select **Storage Array/NAS** when creating the virtual Server Protection Group. For more information, see create a Protection Group for VMware.

Register Nutanix

To leverage the Nutanix Prism element as the storage snapshot of VMware VMs, you must first register the Nutanix Prism element as a source with the Cohesity cluster. To register Nutanix as a storage snapshot provider for VMware, follow the steps described below:

1. Navigate to **Data Protection > Sources**.
2. Click **Register > Virtual Machines**.

   The Register Virtual Machines page appears.
3. From the **Source Type** drop-down list, select **Acropolis: Other Hypervisors**.
4. Enter the following information:

Field Name	Description
Hostname or IP Address	Specify the hostname or IP address of Nutanix Prism.
Username	Specify a Prism username that has adequate privileges to perform actions on the source. For more information, see Nutanix Acropolis.
Password	Enter the password for the specified user.

5. Click **Register**.

   When registration finishes, a new entry appears on the Sources page.

After registering the snapshot provider, to leverage the VMware snapshots from Nutanix, you must enable the **Leverage Storage Snapshots for Data Protection** field and select **Nutanix** when creating the virtual Server Protection Group. For more information, see create a Protection Group for VMware.

Considerations

Review the following considerations of Nutanix Storage Snapshot:

- Cohesity does not support application-consistent backups of VMware VMs that leverage Nutanix as the storage snapshot provider.
- Cohesity does not support differential restore of VMware VMs backed by leveraging the storage snapshots.
- If the protection run of VMware VMs is switched from VADP-based backup to storage snapshot-based backup, then the Cohesity cluster will initiate a full backup.
- While performing storage snapshot-based backup, if a VMware VM snapshot is present on the ESXi host, then Cohesity will perform VADP backup.
- Cohesity does not support leveraging storage snapshots for backing up VM templates.

Register Pure FlashArray

Before registering or editing a Pure FlashArray, ensure it meets prerequisites, minimum permissions, and firewall port requirements in the Plan and Prepare for Pure FlashArray.

To register Pure FlashArray as a storage snapshot provider for VMware, see Register or Edit a Pure FlashArray.

To leverage the VMware snapshots from Pure FlashArray, you must enable the **Leverage Storage Snapshots for Data Protection** field and select **Pure** when creating the virtual Server Protection Group. For more information, see Add or Edit a Protection Group for Virtual Servers.

Register IBM FlashSystem

To register IBM Storage FlashSystem as a storage snapshot provider for VMware, see Register or Edit an IBM Storage FlashSystem.

To leverage the IBM FlashSystem Storage Snapshots for VMware backups, after registering the IBM FlashSystem, you must also enable the **Leverage Storage Snapshots for Data Protection** field and select **Storage Array/NAS** when creating the virtual Server Protection Group. For more information, see Add or Edit a Protection Group for Virtual Servers.

Register HPE Alletra Storage MP

Before registering the HPE Alletra Storage MP as a storage provider, ensure it meets firewall port requirements.

To register HPE Alletra Storage MP as a storage snapshot provider for VMware, see Register or Edit HPE Alletra Storage MP.

To leverage the HPE Alletra Storage MP for VMware backups, after registering the HPE Alletra, you must also enable the Leverage Storage Snapshots for Data Protection field and select Storage Array/NAS when creating the virtual Server Protection Group. For more information, see Add or Edit a Protection Group for Virtual Servers.

Related Topic

Ensure Adequate Privileges for Cohesity on the Source

Add or Edit a Protection Group for Virtual Servers

Register or Edit a Pure FlashArray.

Add or Edit a Protection Group for Virtual Servers

A Protection Group that protects VMs captures snapshots, which are required for the following tasks:

- clone VMs
- recover VMs
- recover files
- instant volume mount
- replicate snapshots
- archive snapshots

> **Tip:** You can also watch the VMware Protection Groups video to learn how to create a Protection Group.

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit a Hypervisor Source
- Create or Edit a Standard Policy
- Create or Edit Storage Domains

Important:

- The options and features available in the following procedure are dependent on the selected source and the Protection Group type. For example if a **Cloud** > **AWS: IAM User** source is selected, the **Cohesity Protection Service** and **Native Snapshot** Protection Group types appear.
- (Applies only to Hyper-V) To prevent the AVHDX snapshot files from being deleted, merged, or cleaned up after successful backup, Microsoft recommends the following workaround:
  - Change the Cohesity Windows agent service user account to LOCAL SYSTEM.
  - Clear all manually set accounts under Log On As A Service/Deny Log On As A Service from both local and group policies.

To add or edit a Protection Group for virtual servers:

1. Navigate to **Data Protection > Protection** and do one of the following:
   - To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
   - To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. If you are creating a Protection Group, click **Protect > Virtual Machines**.

   The New Protection page appears.
3. Click **Add Objects**, and from the **Registered Source** drop-down list, select the Virtual server whose VMs you want to protect. You can also Register a new source with the Cohesity cluster by clicking the **Register Source** option.

   Best practice: For environments with Hyper-V 2012 R2 (v5.0) VMs running on 2016 hosts, create separate Protection Groups for Hyper-V 2012 R2 VMs and Hyper-V 2016 (v8.0) VMs. This separation enables the Protection Group s protecting only Hyper-V 2016 VMs to use resilient change tracking (RCT) backup.

   Note: Cohesity can exclude the backup of placeholder VMs in a VMware VMs Protection Group. This is an Early Access feature and you must contact Cohesity Support to enable this feature.
4. Select the source objects to back up or exclude.

For VMware vCenter, you can also select VM templates as the objects to backup. For VM templates only full backup is performed.

By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.

- Indicates an auto protected object.
- Excluded objects are not protected. The icon indicates that an ancestor of the object is auto protected but this object is explicitly excluded.
- Excluded objects are not protected. The icon indicates that an ancestor of the object is auto protected but this object is excluded through inheritance (an ancestor is explicitly excluded).
- Click this icon to view the general information—type, server size, and version of the selected object.
- vCenter Migrated indicates the VMware VM is migrated from another VMware vCenter.

When protecting vApp and vApp template, you must enable auto-protect. If you are not enabling auto-protect, then Cohesity cluster will only back up the underlying objects of the vApp and vApp template. When protecting vApp VMs, Cohesity recommends using one protection group exclusively for all VMs within a vApp and not mixing vApp VMs with other VMs.

Display modes can show VMs in a hierarchy or a flat list of all VMs.

For VMware vSphere 6.0 onwards, AWS, Azure, and Microsoft Hyper-V, you can also display only VMs with specific tags. Click the **Tags** button, select a tag from the drop-down and click **Add**. You can choose to **Auto Protect** or **Exclude** the VMs that have this tag. After clicking **Auto Protect** or **Exclude**, the source tree is updated with your selection. If a VM has multiple tags with **Auto Protect** and **Exclude** applied, exclusion takes precedence. These excluded VMs are also excluded even if a parent object is auto protected.

The **Tags** button will not be displayed if the Cohesity cluster detects that no VMs have tags assigned to them. For the **Tags** button to appear, log in to the hypervisor and assign a tag to at least one VM. Once the tag is assigned to the VMs, you must manually refresh the registered hypervisor in the Cohesity cluster. To refresh the hypervisor, select **Data Protection** > **Sources**, navigate the source hierarchy to the hypervisor server, click the actions menu, and then click **Refresh**.

AHV Prism Central supports tag-based functionality, where the Tags view displays only VMs associated with tag:

1. Navigate to **Data Protection** > **Protection**.
2. From the **Protect** drop-down, select **Virtual Machines**.

   The **New Protection** page appears.
3. Click **Add Objects**, and from the **Registered Source** drop-down list, select a Prism Central VM.
4. Click the **Tags** button. Selecting a value will display all VMs tagged with that particular value.

You can choose to **Auto Protect** or **Exclude** the VMs that have this tag. After clicking **Auto Protect** or **Exclude**, the source tree is updated with your selection. If a VM has multiple tags with **Auto Protect** and **Exclude** applied, exclusion takes precedence. These excluded VMs are also excluded even if a parent object is auto-protected.

Protection Groups of Hyper-V VMs allow auto-protection using multiple selections in the include and exclude lists and multiple tags in each selection.

You can choose which disk to back up or exclude in a Hyper-V VM or Protection Group, VMware VM or Protection Group, or an AHV VM or Protection Group.

To include or exclude disks for a Hyper-V, AHV VM, or VMware when creating a Protection Group:

a. Navigate to **Data Protection** > **Protection**.
b. From the **Protect** drop-down, select **Virtual Machines**.

   The **New Protection** page appears.
c. Click **Add Objects**, and from the **Registered Source** drop-down list, select a VM.
d. In the **Add Objects** dialog, select the checkbox next to that VM and click the Edit ( ) icon.

   > **Note:** To enable the **Disk Inclusion/Exclusion** option to an existing Protection Group, click the Edit icon under Objects in the Protection Group settings.

e. In the **Options** dialog, enable **Enable Disk Inclusion/Exclusion**.
f. You can choose to include or exclude the disks. After selecting **Include** or **Exclude**, select the disks from the drop-down.

> **Note:** When using VMware Continuous Data Protection (CDP), disk exclusion settings do not prevent data capture on excluded disks. CDP operates at the I/O level and captures changes across all disks in the VM. CDP backups may fail when disk exclusion is used.

g. Click **Save**.

Once saved, the **Customized** link appears next to the Edit icon and lets you modify the disk inclusion/exclusion settings.

To include or exclude disks for a Hyper-V, AHV, or VMware Protection Group as a whole:

a. Navigate to **Data Protection** > **Protection**.

b. From the Actions menu (⋮) of the Protection Group, select **Edit**.

c. Under **Additional Settings** of the Protection Group, select **Include** or **Exclude Disks** and enable **Enable Disk Inclusion/Exclusion**.

d. Select **Include** or **Exclude**.

e. Click **+ Add**.

f. Enter the details of the disk:

AHV: Select the Controller Type and enter the Unit Number of the disk. The Unit Number (or Index) is the unique identifier assigned to that disk when it's added to the VM in the Nutanix Prism Web Console.

Hyper-V: Select the Controller Type and enter the Controller Number and Unit Number of the disk. The Controller Number and Unit Number (or Location) are the unique identifiers assigned to the Controller and Disk by the Hyper-V Manager.

VMware VM: Select the Controller Type and enter the Bus Number and Unit Number for each disk. The Bus Number and Unit Number (or Location) are the unique identifiers assigned to the Controller and Disk by VMware.

g. You can add more disks until you have all the disks you want included or excluded for that Protection Group.

h. Click **Protect.**

> **Note: Disk Inclusion** or **Exclusion** settings of a VM take precedence over the settings of its Protection Group.

The VMware user account designated for data protection requires the following additional privileges to back up encrypted VMs:

- Cryptographic Operations > Add Disk
- Cryptographic Operations > Direct Access
- Cryptographic Operations > Encrypt
- Cryptographic Operations > Encrypt new

When configuring an AWS Protection Group, you can search VMs with specific tags, if the VM Protection Group type is **Native Snapshot**.

You can optionally filter the objects displayed.

In the search field, enter part of the object name.

Use the filtering options:

- In the **Filter** drop-down, select one of the following options:
  - **Unprotected**—Displays the servers that are not protected by any Protection Groups in the Cohesity cluster—not protected by Cohesity.
  - **Protected by Cohesity**—Displays the servers that are protected by all the Protection Groups in the Cohesity cluster.
  - **MS SQL Servers**—Displays only object hierarchy objects that contain servers registered as MS SQL Server.
  - **Protected by this Job**—Displays the servers that are currently selected to be backed up by this Protection Group. (This option is only available when editing an existing job.)

  To reset the filter and show all objects, click **Show All**.
- For virtual servers only, the **Expand to** drop-down allows you to expand the object hierarchy to a specific level.

  Selecting the check box adjacent to **Objects** selects all filtered objects.
- For physical servers only, you can choose to show just Windows or Linux servers. All servers are shown by default.

  Selecting the check box adjacent to **Servers** selects all filtered servers.
- Click the arrow adjacent to object (**Expand All**) to expand the entire object hierarchy. Click the arrow again (**Collapse All**) to collapse the entire object hierarchy.

The graphical representations in the table are described below.

- If you are creating a new Protection Group and the shield appears next to the object, this object is already protected by another Protection Group.

- If you are editing an existing Protection Group and the shield appears next to the object and the object is selected ( ), this object will be protected by the current Protection Group. (It may also be protected by another Protection Group.)

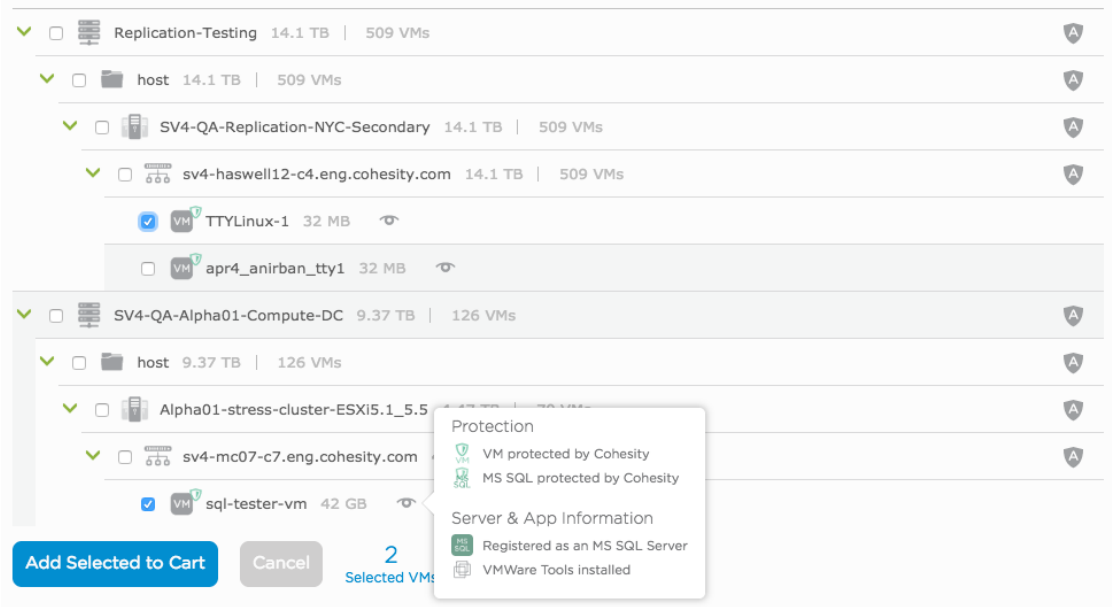

- Click the icon to see details about the object.

  The MS SQL Server application is registered.

  VMware Tools is installed on the VM. If VMware Tools are installed, this VM can be backed up by an application-consistent Protection Group. For more information, see Creating Application-Consistent Snapshots.

- If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.

If appropriate for your organization, enable auto protect.

Click to enable auto protect. Any new VMs added to the parent object in the future will automatically be protected by this Protection Group.

Click to disable auto protect for an object. The icon turns red.

> **Note:** A Protection Group that auto-protected VMware VMs will run successfully without any critical failure alerts, even if it has no objects, with the following message in the Protection Run's Details dialog: "Protection Group has no eligible source to back up.".

The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click .

> **Warning:** You may unintentionally back up cloned VMs with this option. For example, if you select this option for a Protection Group that backs up an ESXi host and that ESXi host is where cloned VMs are created, the new cloned VMs are also backed up the next time the Protection Group runs. You may want to edit the original Protection Group and explicitly exclude the new cloned VMs.

> **Note:** In the Cohesity Dashboard, users with a restricted role (such as operator) cannot access the runs of a Protection Group for leaf objects (such as VMs) that were backed up because a higher level object (such as a datacenter) was auto-protected, even if that user has permissions to access those leaf objects. Cohesity supports assigning permissions to leaf level object such as VMs. Cohesity currently does not support assigning permissions to higher level objects such as a datacenter and since auto-protection is applied higher level objects, access is denied.

5. Click **Continue** to go back to the New Protection page.
6. Select **Protection Group** and in the **Name** field, enter a name from the Protection Group.
7. **Protection Group Type:** This option is available only if the selected source is a cloud source that supports backing up with both agents or native snapshotting.

   Select a Protection Group type:

- **Cohesity Protection Service**— A Cohesity cluster supports backing up VMs in a cloud service (such as Azure or AWS) using an agent. You must install an agent on the VMs with this option.

  For more information, see and .
- **Native Snapshot**—A Cohesity cluster supports backing up VMs in a cloud service (such as AWS) using the native snapshotting of the cloud service. An agent does not need to be installed for this option.

  For more information, see AWS Backup Use Cases.
- **Snapshot Manager**—A Cohesity cluster supports the backup of VMs in a cloud service (Azure and AWS) using the Cloud Snapshot Manager APIs of the cloud service. For more information, see AWS Backup Use Cases.

8. Select **Policy** and from the **Policy** drop-down list, select an existing protection policy or create a new one by selecting **New Policy**. For more information, see Create or Edit a Standard Policy.
9. Select **Storage Domain** and select an existing storage domain or create a new one by selecting Add Storage Domain.
10. Selection one of the following options:
    - **Protect**: Select this option if you want to complete the Protection Group creation without changing or configuring any of the additional settings.
    - **More Options**: Select this option and go to step 12 if you want to change or configure any of the additional settings.
11. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
12. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
13. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.

- **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
- **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.

14. **Defer Incomplete Objects in Concurrent Runs**: (Applies only to VMware VMs) A Protection Group can protect both small (for example, 1 GB VM) and large VMs (for example, 1 TB VM). While a short VM usually completes its protection runs within a few minutes, a large VM might take longer and still be running when it's time for the next run. By default, when it's time to start the next run, the Cohesity cluster checks for any in-progress protection run in that Protection Group. If there are in-progress protection runs (typically a large VM), the Cohesity cluster cancels the next run for that Protection Group, causing the small VMs to also miss that run and the SLA. When the **Defer Incomplete Objects in Concurrent Runs** option is enabled, if there is an in-progress protection run, the Cohesity cluster postpones the subsequent run only for that VM. The postponed run resumes as soon as the current run is completed. The protection runs of the remaining VMs (typically small VMs) in that Protection Group are not affected and run on schedule.

    For example, Protection Group PG1 runs every 8 hours to protect VMs A, B, and C. When the first protection run starts, A and B complete running in a few minutes. VM C, a large VM, needs 12 hours to complete running. If **Defer Incomplete Objects in Concurrent Runs** is disabled for PG1, when it's time for the second run, the Cohesity cluster cancels the second run in all three VMs because C is still in progress. However, the third run completes in all three VMs. Because C is a large VM, A and B miss every alternate run.

If **Defer Incomplete Objects in Concurrent Runs** is enabled for PG1, when it's time for the second run, the Cohesity cluster starts the second run in A and B but postpones the run in C. C's second run resumes when its first run completes. As a result, A, B, and C do not miss any scheduled runs.

If PG1 also protects VM D, another large VM, which needs 20 hours to complete, when it's time for the second run, the Cohesity cluster postpones D's second run. When it's time for the third run, because D's first run is still in progress, the Cohesity cluster cancels its second run but postpones its third run. D's next run will resume automatically when its ongoing run completes.

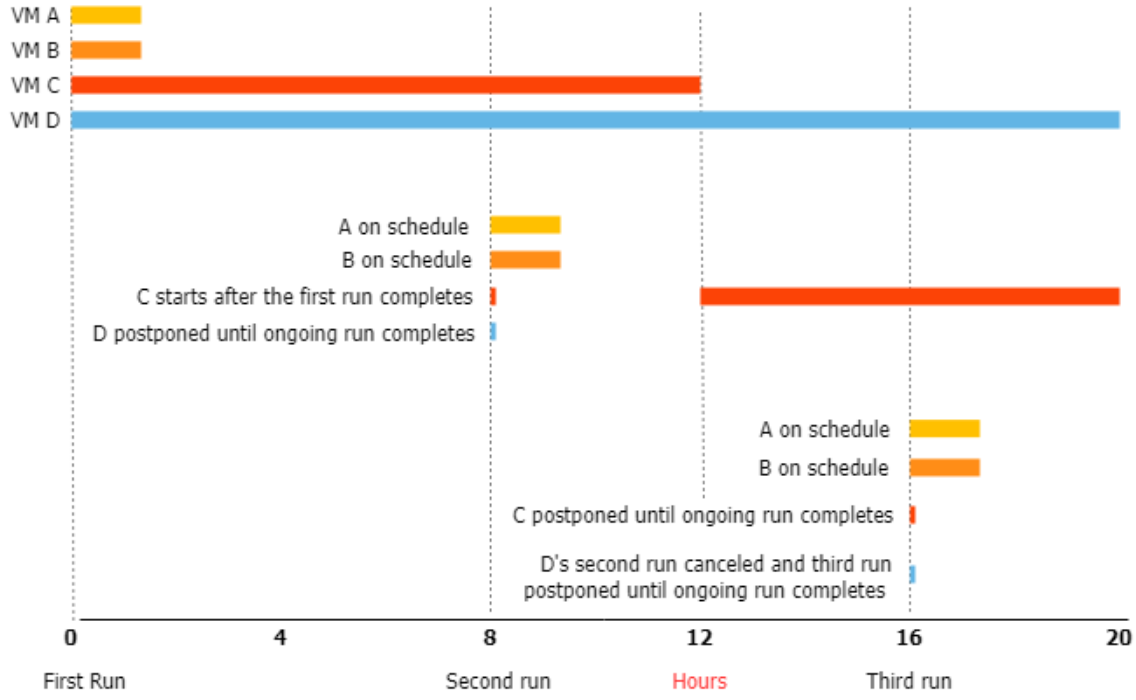

When the Cohesity cluster cancels the backup task of an object (such as D) in the Protection Group (such as PG1), the **Protection Group Details** page displays the following message at the corresponding protection run:

```
"Failed to back up some of the objects in the protection run as they are already backed in another protection run. ".
```

At the corresponding task (D's second backup task), the following message is displayed:

```
"cancelled the backup task of this object as it is already being backup up by another protection run."
```

Enabling the **Defer Incomplete Objects in Concurrent Runs** option helps you meet the defined SLAs, and improves the efficiency of the scheduled runs, thereby enhancing the backup performance.

> **Note: Defer Incomplete Objects in Concurrent Runs** can be enabled for runs with replication or archiving enabled.

15. **Leverage Storage Snapshots for Data Protection**: Toggle on to leverage storage array-based snapshots of supported Storage Snapshot Providers such as Cisco HyperFlex, HPE Nimble Storage, HPE Alletra 6000, NetApp ONTAP, Nutanix Prism element, Pure Flash Array, and IBM FlashSystem.

    Enabling this option greatly reduces the persistence time of VADP snapshots, thereby minimizing VM stuns and snapshot consolidation issues. The storage array-based snapshots can be leveraged only if the Storage Snapshot Provider is registered with the Cohesity cluster. For details on how to register, see Register a Storage Snapshot Provider.

    > **Note:**
    > - Disk exclusions in a VM is not supported if the **Leverage Storage Snapshots for Data Protection** option enabled.
    > - When **Leverage Storage Snapshots for Data Protection** option is enabled, **Leverage SAN Transport for Data Protection** option is not available for selection.
    > - The protection run may not be able to leverage storage snapshots for the backup of all VMs within the Protection Group; for example, when a VM has one or more Virtual RDM disks.
    > - The leverage storage snapshot feature doesn't work when multitenancy is enabled.

16. **Leverage SAN Transport for Data Protection**: Cohesity supports SAN transport for Fibre Channel (FC) and iSCSI-based storage arrays and datastores that contain such storage arrays. Turn on this option to enable the SAN transport mode over FC or iSCSI for backing up VMware VMs to the Cohesity cluster.

    To use the SAN transport mode over FC, ensure the following pre-requirements:

    - Fibre channel is installed and zoned appropriately between the Cohesity cluster and the primary storage.

    To use the SAN transport mode over FC and iSCSI, ensure the following pre-requirement:

- Volumes are mapped appropriately between the Cohesity cluster, ESXi hosts, and the primary storage.

> **Note:**
> - When **Leverage SAN Transport for Data Protection** option is enabled, **Leverage Storage Snapshots for Data Protection** option is not available for selection.
> - This option is valid for storage arrays supported by VMware. Cohesity has tested this with Pure Flash Array, HPE Nimble Storage Array, and IBM Storage FlashSystem configurations.
> - VMware SAN transport mode is not supported with Cohesity virtual nodes. To use SAN transport mode for backing up VMware VMs, zoning and granting datastores access to the Cohesity nodes has to be done manually.

You can turn on the **Allow NBDSSL Transport Fallback** toggle if you do not want to enforce strict SAN transport mode for backing up VMware to the Cohesity cluster. With toggle enabled, backup will not fail if the SAN transport is not available. If you disable this toggle, backup will fail if the SAN transport mode is unavailable.

> **Note:** NBDSSL is the default transport mode for On-Prem vCenter backups. The fallback to NBD if NBDSSL fails is no longer supported.

> **Note:** Ensure that the storage array connects to the primary environment (host) and the Cohesity node over the same protocol (FC on both ends or iSCSI on both ends). Cohesity does not support SAN transport for SAS (Serial Attached SCSI) based storage arrays.

17. **Exclude Disks**: (*Applies only to VMware VM*) By default, all the volumes on a selected server are protected. Toggle on to select disks to exclude for all VMs is this Protection Group. Provide the type of controller, bus number and unit number for each disk to exclude. Excluded disks are not backed up and are not recovered during VM recovery.

> **Note:** If you exclude a disk that is part of a striped volume, the Cohesity cluster does not index the volume even if **Indexing** is toggled on for this Protection Group. You will not be able to search for or recover individual files in that volume.

18. **Include Disks**: (*Applies only to VMware VM*) By default, all the volumes on a selected server are protected. Toggle on to select disks to exclude for all VMs is this Protection Group. Provide the type of controller, bus number, and unit number for each disk to include. Only the specified disks are backed up and recovered during VM recovery.
19. **Exclude Physical RDM Volumes**: (*Applies only to VMware VM*) VMware vSphere Storage APIs cannot be used to snapshot physical compatibility RDM disks. Enable this option to allow the backup of VMs by excluding physical compatibility RDM disks. If you disable this option, the backup of VMs with physical compatibility RDM disks will fail. Excluded disks are not backed up and are not recovered during VM recovery. Cohesity recommends you to use agent-based backups to protect data on the physical RDM disks.

> **Note:**
> - VMware does not support taking the snapshot or backup of the VMs having RDM (physical and virtual) on Shared Bus.
> - Backup of Virtual compatibility RDM disks is supported and does not depend on the **Exclude Physical RDM Volumes** toggle setting.

20. **Protection Type**: It provides the following options.
    1. **Auto**: Select this option to automatically identify the mode based on the version of virtual machines in the protection job. If there is a VM that does not support RCT, then the VSS backup approach is used for all the VMs in that protection job.
    2. **VSS**: Select this option to use Hyper-V VSS Writer as the backup approach.
    3. **RCT**: Select this option to use Hyper-V WMI Based Backup as the backup approach. Resilient Change Tracking (RCT) is available on Windows Server 2016 or later and the version of the virtual machine should be 8.0 or later.
21. **App Consistent Backup**: (*Applies only for VMware, Hyper-V, and AHV VMs*) Enable if you want the guest operating systems of all the VMs in the Protection Group to be quiesced before snapshots of these VMs are created. When this option is selected, the Cohesity cluster requests the VMware vSphere/AHV software to create a quiesced VM

snapshot by invoking VMware Tools/Nutanix Guest Tools (installed on the guest operating systems of the VM). The VMware Tools/Nutanix Guest Tools requests that applications on the guest OS quiesce their state so that application-consistent snapshots can be created. Quiescing of VMs before capturing snapshots ensures the integrity of the data saved in the snapshots.

If the Cohesity cluster fails to capture an app-consistent snapshot, you can use the **Take a Crash Consistent backup if unable to perform an App Consistent backup** option to capture a crash-consistent snapshot. The option is available only if the **App-Consistent backups option** is enabled.

For example, the Cohesity cluster may fail to perform an app-consistent backup for the following reasons: when VMware Tools/Nutanix Guest Tools is not installed on the VM, the VM (VMware) is powered off, or if the VM cannot be quiesced. If the **Take a Crash Consistent backup if unable to perform an App Consistent backup** is enabled in that Cohesity cluster, a crash-consistent snapshot is captured.

> **Note:** App-consistent backups apply to VMs only. For physical servers, Windows is app-consistent by default and Linux is crash-consistent. For more information, see Creating Application-Consistent Snapshots.

22. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.

23. **Cloud Migration:** Enable this option to support the cloud migration of VMs between hypervisors (such as vCenter) servers and cloud services for failover and failback. Cohesity agents must be installed on the Windows VMs prior to backing up the snapshots on the on-premises Cohesity cluster. For more information, see Disaster Recovery of VMs Using AWS Cohesity Cloud Edition.

    After enabling cloud migration, you can download the Cohesity Agent directly from the Cohesity Dashboard.

24. **CDP Synchronous Replication**: Displayed only if the selected policy has **Continuous Data Protection** and **Replication** enabled.

    Before you enable this field, you must enable **Encryption** in the remote Cohesity cluster to which the snapshots will be replicated. For details on how to enable **Continuous Data Protection** and **Encryption**, see the Create a Protection Policy section in Cohesity CDP for VMware.

    Toggle on to enable CDP synchronous replication and achieve near-zero RPO (Recovery Point Objective) on the DR target site (remote cluster). Once enabled, real-time I/Os of the VM are streamed to the remote Cohesity cluster and are available for

immediate recovery. Enabling Synchronous Replication minimizes the risk of data loss to near-zero for your mission-critical applications during an outage.

Synchronous replication is expensive in terms of storage because the protected data is replicated almost immediately to the remote cluster.

25. **Cancel Runs at Quiet Time Start**: Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
26. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.
27. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
28. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

    > **Note:**
    > The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

29. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.

- **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
- **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).

30. **Pause Future Runs**: Enable this option to pause future protection runs for the Protection Group.
31. **Description:** Optional. Click the plus sign and enter a description for the Protection Group.

> **Note:** For more information on AHV related options, see directly attached volume groups on Nutanix AHV.

1. Click **Protect**.

   The Protection page is displayed and includes the new Protection Group.

> **Note:**
> - If a server is being backed up and files will also be recovered to that same server, the file recovery task starts after the backup is completed. The converse is also true; if files are being recovered to a server and that same server will be backed up, the backup starts after the file recovery is completed.
> - For a File-Based backup, the "Protected" and "Protected Size" fields on the source page for physical sources are updated after the source refresh.

Protection for NVRAM files

The NVRAM (non-volatile random-access memory) file of a VMware VM stores the boot drive setting and, in the case of multiple disks, the boot order. As VMware recommends, Cohesity automatically backs up the NVRAM file when backing up its VM.

> **Note:** It is important to back up the NVRAM file if you created the VM with the Virtual Trusted Platform Module (vTPM).

Supported Workloads

Cohesity supports the backup and recovery of NVRAM files when backing up VMware VMs.

NVRAM backup and recovery is available for VMC, VCD, Storage Array Snapshot Arrays, Snapshot Integration, Archival, Replication, and Cloud Spins.

Troubleshooting with NVRAM Files

Errors

You might run into the following error scenarios when protecting VMware VMs:

- The backup of the VM(s) is successful, but the status of the protection run is set to "warning" and the VM fails to start.
- The VM(s) recovery is successful, but the VM fails to start.
- If the NVRAM file is missing for a recovered VM (which can happen if the backup was taken before upgrading the Cohesity cluster to 7.1), VMware will create a new NVRAM file with default settings. While this new file would suffice to start any VM, the VM might not start if the BIOS or UEFI settings were customized before recovery. For example, if you've modified the OS boot path, this setting will not be present in the newly created NVRAM file causing the VM to not start.

Solution

Follow these steps to start the VM by replacing its NVRAM file with the downloaded NVRAM file:

1. Navigate to **Data Protection** > **Recoveries**.
2. Select **Recover** and then choose **Virtual Machines** from the menu.
3. Select the VM whose NVRAM file you want to download.
4. Perform one of the following:
    1. If you've protected the NVRAM file, click the Actions Menu (⋮) of that VM and select **Download NVRAM File**. The NVRAM file is downloaded to your local system.
    2. If you have not protected the NVRAM file, download the NVRAM file from the original VM in the vCenter.
5. Ensure the name of the downloaded NVRAM file is the same as the NVRAM file in the recovered VM.
6. Replace the NVRAM file in the recovered VM with the NVRAM file downloaded in step 4.
7. Restart the VM.

Considerations

- If you are protecting a VM equipped with vTPM, enable the Manage KMS privilege in the vCenter for the user account registered with Cohesity.

- If you are recovering a VM equipped with vTPM using a different KMS or replacing the current KMS with a new one, the vTPM VMs may not start after the recovery process, even if the new KMS has the same name.
- The supported maximum size of the NVRAM file is 10 MB. Contact your Cohesity account team if your NVRAM files are larger than 10 MB.
- Restoring the NVRAM file from the encrypted VMs is not supported, except for vTPM encryption.

Protection for vSphere DataSets files

The VM configuration files and vSphere DataSets files are stored together, and a reference to the DataSet (.dsd) is written into the virtual machine configuration file (.vmx). The datasets of the VM are stored within the DataSets file. The DataSets files are denoted with a .dsd extension and are stored alongside the VM configuration files. The maximum size of the DataSets file is 300 MB.

Cohesity clusters now automatically back up the VM's DataSets file while protecting a VM and also allow you to recover the DataSets file.

DataSets file functionality is available in vSphere 8.0 and requires the VM to be running virtual hardware version 20.

Supported Workloads

Cohesity supports the backup and recovery of DataSets files when backing up VMware VMs.

vSphere DataSets file backup and recovery is available for VADP, Storage Array Snapshot Arrays, and HCI integration.

Troubleshooting with DataSets files

Errors

You might run into the following error scenarios when protecting VMware VMs:

- The backup task will raise a warning if the DataSets files are not backed up sccessfully.
- The restore task will raise a warning if the DataSets files are not restored successfully.

Considerations

- The maximum size of a DataSets file can be up to 300 MB. In most cases, the DataSets file size will be a few MBs. Contact your Cohesity account team if the files exceed 300 MB.
- Restoring the DataSets file from the encrypted VMs is not supported, except for vTPM encryption.

Nutanix AHV

Cohesity supports backup and restore of vTPM metadata. It also supports backup of directly attached Volume Groups for Nutanix AHV VMs.

AHV Virtual Trusted Platform Module

Directly Attached Volume Groups on Nutanix AHV

AHV Virtual Trusted Platform Module (vTPM)

You can now back up Virtual Trusted Platform Module (vTPM) metadata for AHV VMs. During backup, Cohesity records whether vTPM was enabled on the VM. vTPM configuration is backed up and restored only if it is enabled on the source VM at the time the backup is taken. This ensures that the original vTPM configuration is preserved.

vTPM metadata is backed up and restored to ensure the security configuration of VMs. vTPM metadata does not include any cryptographic keys or other data stored within the vTPM.

> **Note:** Ensure to manually back up the important cryptographic data stored in the vTPM. the cryptographic data stored in the vTPM module is important, it must be manually backed up by the user.

Directly Attached Volume Groups on Nutanix AHV

Nutanix AHV allows VMs to attach Volume Groups, which are collections of virtual disks exposed to the VM using the iSCSI protocol. Cohesity extends AHV VM protection to include these directly attached Volume Groups.

During protection, set the **Backup Directly Attached Volume Groups** to **Yes** under Additional Settings to back up the VM along with its directly attached volume groups.

On the **Activity page**, the volume groups that will be backed up are listed. If the option is disabled, the volume groups are displayed as skipped. If backup of directly attached volume groups is enabled but no volume groups are attached to the VM, this status is also reflected on the **Activity** page. It also displays the Nutanix UUID of the volume disk attached.

When enabled, snapshots capture:

- The attached Volume Groups' configuration.
- Disks belonging to the attached Volume Groups.

You can also configure disk inclusion and exclusion for volume groups. Disk selection is based on the disk address label.

When creating a protection group, select **Include or Exclude Disks** option, and then specify the disks that you want to include in or exclude from the backup.

On the **Activity** page, the volume groups and disks that are included or excluded from the backup are listed.

> **Note:**
> - Volume Group backups are crash-consistent, as Nutanix AHV does not support application-consistent snapshots for Volume Groups.
> - Enabling Backup Directly Attached Volume Groups may increase backup duration due to additional snapshot operations.
> - Inclusion or exclusion of Volume Groups is supported only via index-based selection.

Considerations

- Inclusion and exclusion of Volume Groups during backup is supported only through index-based selection. It uses the Nutanix assigned index for the VG attachment.
- Nutanix AHV does not support application-consistent snapshots for Volume Groups. All Volume Group backups are crash-consistent.
- Restored Volume Group attachments may not retain the original SCSI IDs used during backup.

Manage the VMware VM Backup Configuration

In Cohesity Data Cloud, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered VMware VMs and the assigned Storage Domain in the Protection Group settings.

Add VMware VMs to the Protection Group

You can add new VMware VMs to the existing Protection Group for backup. The newly added VMware VMs are backed up from the next protection run and it will be a full backup.

Perform the following steps to add VMware VMs to the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.

4. On the **Edit Protection** page, in the **Object** field, click the edit icon.
5. Select the VMware VM you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added VMware VMs are backed in the next scheduled protection run.

Remove VMware VMs from the Protection Group

From the Edit Protection page, you can remove VMware VMs from the Protection Group that you do not want to back up in the next protection run.

Perform the following steps to remove VMware VMs from the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the edit icon.
5. Unselect the VMware VM you want to remove from the Protection Group.
6. Click **Save Selection** > **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new VMware VMs that are added to the selected parent VMware host. However, you cannot enable auto-protect at the server-level containing both Read-Write volumes and Data-Protect volumes.

Perform the following steps to enable auto protect:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the edit icon.
5. Click the shield icon to enable Auto Protect.
6. Click **Save Selection** > **Save**.

When this option is enabled on the VMware parent entity(vCenter, Folder, Datacenter, Compute Resource, Host System), all the VMware VMs that are added to the parent VMware

entity in the future are automatically added to the Protection Group and are protected from the next protection run.

For example, you create a Protection Group that auto protects the VMware parent entity. The Protection Group runs for two weeks and then a new VM is added to the VMware parent entity. . The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list, and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Excluded objects that are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirement.

Perform the following steps to edit the policy settings:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Policy** field, click the **Edit Policy** icon .

   The **Edit Policy** page appears
5. Edit the settings and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to the Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

Perform the following steps to attach a new protection policy to the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.

4. On the **Edit Protection** page, from the **Policy** drop-down, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to back up VMware VMs. It provides you options to back up all the objects, only selected objects, or failed objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- Is configured with replication targets, you can select those targets.
- Is configured with archival targets, you can select those targets.

Perform the following steps to start an on-demand protection run:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. From the **Backup** drop-down list, select any one of the following:
        - **All objects in the Protection Group**.
        - **All objects that failed last run**.
        - **All objects that succeeded last run**.
        - **Selected objects in the protection run**.
    2. In the **Backup Type** drop-down list, select the backup type.
    3. Optional. In the **Replicate** section, select the replication target and the retention duration or click **Add** to add the replication target and the retention duration. Also, select **Replicate only fully successful runs** to replicate the data only from the completed protection runs.
    4. Optional. In the **Archive** section, click **Add** to add the external target to which you want to archive the data and the retention duration. Also, select **Archive only fully successful runs** checkbox to archive the data only from the completed protection runs.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

Perform the following steps to pause the future protection runs:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered

Perform the following steps to resume paused protection runs:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity Data Cloud provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Protection Group are not deleted.
- **Delete Protection Group and Snapshots**: Delete the Protection Group and all snapshots, any local snapshots captured by this Protection Group and stored on this cluster are deleted. In addition, any archived snapshots stored on external targets that were captured by this Protection Group are also deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

Perform the following steps to delete the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for VMware VM.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Run.

If the Protection Group is replicating to a remote cluster, the failover job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the failover job.

Creating Application-Consistent Snapshots

Cohesity supports application-consistent backup for VMware, Nutanix Acropolis, and Hyper-V. By enabling application-consistent backups for a Protection Group, the guest-operating systems of all the VMs in the Protection Group will be quiesced before snapshots of these VMs are created.

Application-Consistent Snapshots for VMware VMs

Toggle **App-Consistent backups** for a Protection Group if you want the guest operating systems of all the VMs in the job to be quiesced before snapshots of these VMs are created. If this option is selected, the Cohesity cluster makes a request to the VMware vSphere software to create a quiesced VM snapshot by invoking VMware Tools (installed on the guest operating systems of the VM). The VMware Tools requests that applications on the guest OS quiesce their state so application-consistent snapshots can be created. This quiescing of VMs prior to capturing snapshots ensures the integrity of the data saved in the snapshots.

App-consistent backups apply to VMs only. For physical servers, Windows is app-consistent by default and Linux is crash-consistent.

Snapshots created without quiescing are known as Crash-Consistent Snapshots.

When the **App-Consistent backups** toggle is on, all the guest Operating Systems of the protected Objects (VMs) in this Protection Group must have VMware Tools installed. The Cohesity Dashboard prevents you from saving the Protection Group if any selected VMs do not have VMware Tools installed.

If the **App-Consistent backups** toggle is on, the **Take a Crash Consistent backup if unable to perform an App Consistent backup** toggle is available. Toggle it on if you want the Cohesity cluster to capture a crash-consistent snapshot if the Cohesity cluster fails to capture an app-consistent snapshot. For example, the Cluster may be unable to perform an app-consistent backup when VMware Tools is not installed on the VM, the VM is powered off, or the VM cannot be quiesced. If the **Take a Crash Consistent backup if unable to perform an App Consistent backup** toggle is off and the Cohesity cluster is unable to perform an app-consistent backup of a VM, a snapshot is not captured.

If the **App-Consistent backups** option is selected for a Protection Group and the Job is protecting Linux VMs, the Cohesity cluster invokes the VMware vSphere software to create the Snapshot. For a Linux VM, the VMware vSphere software *always* reports that the VM was not quiesced because Linux does not provide a system quiescing mechanism like Windows VSS. Because of this VMware behavior, the Cohesity cluster does not check and report on the quiescing return status for Linux VMs. However, other status are checked and reported, for example if Snapshots of VMs are not created, the protection run fails.

If the **App-Consistent backups** option is selected for the Job, when the Cohesity cluster calls VMware vSphere software to quiesce Linux or Windows VMs, if quiesce custom pre-freeze and post-thaw scripts can be found in the correct locations, these scripts are run as part of the quiescing and create Snapshot process. For more information, see the appropriate documentation for the Operating System type:

- Running custom quiescing scripts inside Windows virtual machine with ESX 3.5 Update 2 and later (1006671) Knowledge Base Article
- Configuring Custom Quiesce Scripts for Linux Virtual Machines

Application-Consistent Snapshots for AHV VMs

Toggle App-Consistent backups for a Protection Group if you want the guest operating systems of all the AHV VMs in the Protection Group to be quiesced before snapshots of these VMs are created. If this option is selected, the Cohesity cluster makes a request to the AHV software to create a quiesced VM snapshot by invoking the Nutanix Guest Tools (installed on the guest operating systems of the VM). When you invoke the Nutanix Guest Tools, the VSS and application-consistent snapshot applications are automatically enabled. With these applications enabled, the application-consistent snapshots for all the VMs that support VSS are taken using the Nutanix native in-guest VmQuiesced Snapshot Service (VSS) agent.

Snapshots created without quiescing are known as Crash-Consistent Snapshots.

When the **App-Consistent backups** toggle is on, all the guest Operating Systems of the protected Objects (VMs) in this Protection Group must have Nutanix Guest Tools installed. The Cohesity Dashboard prevents you from saving the Protection Group if any selected VMs do not have Nutanix Guest Tools installed.

If the App-Consistent backups toggle is on, the **Take a Crash Consistent backup if unable to perform an App Consistent backup** toggle is available. Toggle it on if you want the Cohesity cluster to capture a crash-consistent snapshot if the Cohesity cluster fails to capture an app-consistent snapshot. For example, the Cluster may be unable to perform an app-consistent backup when Nutanix Guest Tools is not installed or running on the VM, or the VM cannot be quiesced. If the **Take a Crash Consistent backup if unable to perform an App Consistent backup** toggle is off and the Cohesity cluster is unable to perform an app-consistent backup of a VM, a snapshot is not captured.

If the **App-Consistent backups** option is selected for the Protection Group, when the Cohesity cluster calls AHV software to quiesce Linux or Windows VMs, if quiesce custom pre-freeze and post-thaw scripts can be found in the correct locations, these scripts are run as part of the quiescing and create Snapshot process. For more information, see the Nutanix Documentation.

Application-Consistent Snapshots for Hyper-V VMs

Enable **App-Consistent backups** for a Protection Group if you want the guest operating systems of all the Hyper-V VMs in the Protection Group to be quiesced before snapshots of these VMs are created. If this option is selected, the Cohesity cluster makes a request to the Hyper-V software to create Application Consistent checkpoints.

When the **App-Consistent backups** toggle is ON, all the guest Operating Systems of the protected Objects (VMs) in this Protection Group must have:

- Hyper-V Integration Services installed. For more information, see Hyper-V Integration Services.
- Production Checkpoints enabled. For more information, see Enable Production Checkpoints.

If the Cohesity cluster fails to capture an app-consistent snapshot, you can use the **Take a Crash Consistent backup if unable to perform an App Consistent backup toggle to capture a crash-consistent snapshot** option. This option is available only if the **App-Consistent backups** toggle is ON. For example, the Cohesity cluster may be unable to perform an app-consistent backup for the following reasons: when Hyper-V Integration Services is not installed or running on the VM, or the VM cannot be quiesced. If the **Take a Crash Consistent backup if unable to perform an App Consistent backup** toggle is ON for that Cohesity cluster, a crash-consistent snapshot is captured.

Cancel Protection Run at Object Level

You can now cancel the protection run of individual objects within a Protection Group created for **VMware vCenter**.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Protection Group States at Granular Level

A Protection Group run can be in the following states:

- Protection Group is in progress; where the backup of the objects is in progress.
- Protection Group is progress but the backup of a few objects is in failed status.
- Protection Group is completed; where the backup of a few objects is failed and a few are success and is waiting in the retry interval time.

  > **Note:** Retry is the time you set for the Cohesity cluster to attempt backup before the Protection Group fails. Retry interval time is defined in the Protection Policy along with retry attempts.

- Retry attempts are over and the Protection Group is completed with Success or Failed status.

The protection run cancellation at the object level is applicable for the first three states mentioned above. You cannot cancel the protection run at the object level if the Protection Group run is complete and the retry attempts are over.

Cancel Protection Run at Object Level

To cancel the protection run of individual objects within a Protection Group, perform the following steps:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group.

   The Protection page displays the protection runs of the Protection Group.
3. Select the in-progress protection run from which you want to cancel the protection run of the individual objects.

   The Run details page displays the list of objects that are protected in the protection run.
4. Perform one of the following actions based on your requirement:

- To cancel the protection run of individual objects, click the actions menu (⋮) next to the object and then click **Cancel VM**.
- To cancel the protection run of multiple objects, select the checkboxes next to the objects and click on the cross icon that appears on the column heading of the Run details page.

Considerations

Consider the following points when canceling the protection run of the objects when the protection run is in progress:

- If you are canceling the protection run of the objects that are being backed up, then the protection run of these objects will get canceled immediately.
- If you are canceling the protection run of the objects whose backups are failed, then the protection run of these objects will get canceled in the next retry attempt. The next retry attempt will happen once the current protection run is complete and retry interval is over.
- If you are canceling the protection run of the failed objects during the retry interval time (while the protection run is waiting for the next retry), then the protection run of these objects will get canceled immediately.

  For example, consider the following scenario:

  A Protection Group is created for backing up 10 objects. The retry time set is 5 minutes. During the protection run,

  - 5 objects (V1, V2, V3, V4, V5) are already backed up and status is success.
  - 3 objects (V6, V7, V8) are being backed up and the status is in progress.
  - 2 objects (V9, V10) are failed in backing up and status is failed.

  Now, if you cancel the protection run of V6, V7, or V8 objects then the protection run of these objects will get canceled immediately.

  If you cancel the protection run of V9 or V10 objects then the protection run of these objects will get canceled in the next retry attempt. The next retry attempt will happen 5 minutes after the current protection run is complete.

  If the protection run of V6, V7, and V8 is complete and you cancel the protection run of V9 and V10 during the retry interval time of 5 minutes, then the protection run of these objects would get canceled immediately and the next retry attempt will happen immediately for V9 and V10.

Migrate VMware VMs Using Cohesity

All the VMs that need to be migrated must be protected using one or more Protection Groups in Cohesity. When the initial backup run is started for the Protection Group, Cohesity

first performs a full backup of the source VMs. Once the initial backup is completed, you can initiate the migration for the source VMware VMs. As a part of the migration, Cohesity creates a standby VM on the target vCenter for each source VM being migrated and restores the backed-up data to these standby VMs. Now, whenever there is an incremental backup, Cohesity migrates only the incremental backed-up data to standby VMs. When you decide to complete the migration, Cohesity will migrate the latest incremental backed-up data to standby VMs and then will stop syncing the data. You can power on the standby VMs once the migration is complete. This reduces the downtime of the VM migration as the cutover time (time between powering off the source VMs and powering on the standby VMs) is minimal.

Perform Migration for VMware VMs

To perform migration for VMware VMs:

1. Register both the source VMware vCenter (where the source VMs reside) and the target VMware vCenter (where the VMs need to be migrated) with the Cohesity cluster. For more information, see Register or Edit a Hypervisor Source.
2. Create a Protection Group for backing up the source VMware VMs. For more information, see Add or Edit a Protection Group for Virtual Servers.
3. Initiate VMware VM Migration.
4. Finalize VM Migration.

Initiate VMware VM Migration

When you initiate the VMware VM migration, Cohesity will first create standby VMs and then migrate or restore the backed-up data to the standby VMs. After that, whenever there is an incremental backup, Cohesity will sync this incremental backed-up data with the standby VM.

To initiate VM migration:

1. In the Cohesity Dashboard, select **Data Protection** > **Recovery** > **Virtual Machines** > **Migrate Virtual Machine (VMware only)**.
2. Search for Protected Object (VM) or Protection Group by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs is displayed.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
3. From the list of search results, select one or more VMs that need to be migrated. You can also select Protection Groups.

> **Note:**
> - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for migration and the option to select those VMs individually will be greyed out in the search result list.
> - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

4. Click **Next: Migration Options**.
5. Provide the following **Migration Target** options:
    1. **Register Source**: Select a registered source to migrate the VMs to.
    2. **Resource Pool**: Select a Resource Pool to store the migrated VMs.

        > **Note:** When protecting a VMware source, if you select a Resource Pool that is not part of the ESX cluster, all datastores in the vCenter are displayed. Ensure you select a datastore that is part of the selected Resource Pool.

    3. **Data Stores**: Select the target datastore. The VM files (such as VMDK files) extracted from the VMs are stored in the selected datastore.
    4. **VM Folder**: Optionally select a logical VM folder in the data center in which to create the VMs.
6. Provide details for the following **Options**:
    1. **Network**: From the drop-down menu, select a network to attach the virtual Network Interface Card (vNIC) to a new network. The new network might be a standard port group or a distributed port group. For VMware VMs, the selected Network port group will be attached to all the network interfaces of the VMs being migrated.
    2. **Rename**: Add prefix and/or suffix strings to the names of the VMs that will be migrated by this task. For example, you could specify Prefix Test_ and Suffix _QA to the VM name WinOS and when the VM is migrated, the new VM that is created is called Test_WinOS_QA.

3. **Cluster Interface**: By default, the **Auto Select** option is enabled and the migration task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
4. **Task Name**: Change the default name of the migration task.

7. Click **Start Migrating**.

The migration task is initiated. To complete the migration, see Finalize VM Migration.

Finalize VM Migration

You can finish the migration at any point in time by selecting the **Finalize Migration** option on the recoveries page. When you finalize the migration, Cohesity will migrate the latest incremental backed-up data to the standby VMs and then will stop syncing the data further. You can optionally choose to power on the standby VMs once the migration is complete.

To complete the migration, perform the following steps:

1. On the recoveries page, select the migration task you initiated.

   The Migrate VM page appears. The page displays the last synced snapshot details and the sync status.
2. Click **Finalize Migration** on the top right of the page to complete the migration task.
3. In the Finalize Migration screen, you can choose to **power on VMs after migration**.
4. Click **Finalize Migration**.

Sync Objects

When you initiate VM migration, Cohesity will automatically sync the data from the source VMs to the standby VMs. You can also manually sync the data from the source VMs to the standby VMs if required.

To manually sync the data:

1. On the recoveries page, select the migration task you initiated.
   - To manually sync all the objects (VMs) selected for migration:
     a. Click the actions menu on the top-right of the page.
     b. Select **Sync All Objects**.
   - To manually sync individual objects selected for migration:
     a. Click the sync icon available for the individual VM.

Cancel VM Migration

You can cancel the migration task that you initiated earlier.

To cancel the migration task, on the recoveries page, select the migration task you initiated earlier.

- To cancel the migration of all the objects selected for migration:
  a. Click the actions menu on the top-right of the page.
  b. Select **Cancel Migration**.
  c. Confirm your action by clicking **Cancel Migration**.
- To cancel the migration of individual objects selected for migration:
  a. Click the delete icon available for the individual VM.
  b. Confirm your action by clicking **Cancel Migration**.

> **Note:** Canceling the VM migration task will not delete the standby VMs that are already created.

Consideration

- The migration target has to be a VMware vCenter.
- During Replication, an inactive copy of the Protection Group is created on the remote Cohesity cluster. Migration of VMware VMs protected by such inactive Protection Groups is not supported. For details on Replication, see About Replication.

Protect Migrated VMs or vCD

The VMs of a VMware vCenter Server that are being protected by a Cohesity Protection Group might migrate to another VMware vCenter or vCD for any of the following reasons:

- Upgrade or maintenance activity.
- Inplace vCenter upgrade.
- Inventory deletion and re-import into same or another vCenter
- Failover and Failback and so on.
- Migration of workloads from vCenter to VCD.

If the VMware vCenter to which the VMs are migrated is already registered with the Cohesity cluster, then by default, the Cohesity cluster discovers these VMs as new VMs on the VMware vCenter Server. However, if you enable the **Detect VM migration across vCenter to Preserve backup chain** option while registering or in the registered VMware vCenter source (from which the VM is migrated), then:

- Cohesity discovers these VMs as migrated VMs on the VMware vCenter to which the VMs are migrated.
- When a Protection Group is run on the VMware vCenter to which the VMs are migrated, Cohesity eliminates the need of taking a full backup of the migrated VMs if they were already protected before the migration, and continues to take an incremental backup.
- On the Cohesity cluster, the migrated VMs are displayed with a tag with the name, **vCenter Migrated**.

For information on enabling **Detect VM migration across vCenter to Preserve backup chain** option, see Register or Edit a Hypervisor Source.

If the VMs of a VMware vCenter protected by the Cohesity cluster are migrated to a vCD, which is already registered with the Cohesity cluster, then Cohesity DataPlatform preserves the backup history of these migrated VMs. This feature eliminates the need to take a full backup of the migrated VMs when a Protection Group is run on the vCD to which the VMs are migrated. For information on enabling **Detect VM migration across vCenter to Preserve backup chain** option, see Detect VM migration across vCenter to Preserve backup chain.

View Migrated VMs Original Source

From VMware vCenter to another VMware vCenter

To view the migrated VMs original source:

1. In the Cohesity cluster, select **Data Protection** > **Sources**.
2. Select the VMware vCenter to which the VMs are migrated and then select the **All Objects** tab.

   The objects of the VMware vCenter appear in a Hierarchical view. The VMs that are migrated from another VMware vCenter are displayed with the tag name, **vCenter Migrated** displayed.
3. Click the eye icon displayed next to the VM. The **Virtual Machine Details** pop-up window displays the original VMware vCenter source name and the migration timestamp.

From VMware vCenter to vCD

To view the migrated VMs original source:

1. In the Cohesity cluster, select **Data Protection** > **Sources**.
2. Select the vCD to which the VMs are migrated and then select the **All Objects** tab.

   The objects of the vCD appear in a Hierarchical view. The VMs that are migrated from another VMware vCenter are displayed with the tag name **vCenter Migrated** displayed.

3. Click the eye icon displayed next to the VM. The **Virtual Machine Details** pop-up window displays the original VMware vCenter source name and the migration timestamp.

Consideration

Considerations when protecting migrated VMs

From VMware vCenter to another VMware vCenter:

- The backups of the first Protection Group created to protect the migrated VMs are always incremental. Backups of subsequent Protection Groups created to protect the migrated VMs are incremental by default. Contact your Cohesity account team if you want full backups from the subsequent Protection Groups.

  For example, Protection Group J1 is the first Protection Group created to protect the migrated VMs A and B. Following this, Protection Group J2 is also created to protect VMs A and B. J1 backups are always incremental. J2 backups are incremental by default and can be converted to full backups by contacting your Cohesity account team.
- If you want to recover the snapshot of the VM that was taken before migration, then while recovering, you must select the VM that is protected by the Protection Group created for the VMware vCenter source from which the VMs were migrated.
- If you want to recover the snapshot of the VM that was taken after migration, then while recovering, you must select the VM that is protected by the Protection Group created for the VMware vCenter source to which the VMs are migrated.
- You can enable the **Detect VM migration across vCenter to Preserve backup chain** option for an existing registered source by editing the source. If you are enabling this option to an existing VMware vCenter source, then you must refresh (**Data Protection** > **Sources** > <VMware source> **Action menu**>**Refresh**) the source for this feature to function. Ensure that the source VMware vCenter is reachable till you refresh the target VMware vCenter.
- If your Cohesity cluster does not identify all the migrated VMs or incorrectly identifies VMs as migrated, contact your Cohesity account team.
- To perform incremental backup for any other scenario, contact your Cohesity account team.

From VMware vCenter to another vCD:

- The migration target has to be a vCD.
- During Replication, an inactive copy of the Protection Group is created on the remote Cohesity cluster. Migration of VMware VMs protected by such inactive Protection Groups is not supported. For details on Replication, see About Replication.
- Considerations when protecting migrated VMs:

- The backups of the first Protection Group created to protect the migrated VMs are always incremental. Backups of subsequent Protection Groups created to protect the migrated VMs are incremental by default. Contact your Cohesity account team if you want full backups from the subsequent Protection Groups.

  For example, Protection Group J1 is the first Protection Group created to protect the migrated VMs A and B. Following this, Protection Group J2 is also created to protect VMs A and B. J1 backups are always incremental. J2 backups are incremental by default and can be converted to full backups by contacting your Cohesity account team.
- If you want to recover the snapshot of the VM that was taken before migration, then while recovering, you must select the VM that is protected by the Protection Group created for the VMware vCenter source from which the VMs were migrated.
- If you want to recover the snapshot of the VM that was taken after migration, then while recovering, you must select the VM that is protected by the Protection Group created for the vCD source to which the VMs are migrated.
- You can enable the **Detect VM migration across vCenter to Preserve backup chain** option for an existing registered source by editing source. If you are enabling this option to an existing VCD source, then you must refresh (**Data Protection** > **Sources** > <**VMware source**> **Action menu**>**Refresh**) the source for this feature to function. Ensure that the source VMware vCenter is reachable till you refresh the target VCD.
- If your Cohesity cluster does not identify all the migrated VMs or incorrectly identifies VMs as migrated, contact your Cohesity account team.
- To perform incremental backup for any other scenario, contact your Cohesity account team.

Cohesity CDP for VMware

To protect your VMware VMs, you can optionally choose to leverage the Continuous Data Protection (CDP) option available on the Cohesity Data Cloud. Cohesity's CDP provides near-zero RPOs for your applications.

Every change to the VM (IOs) is captured and replicated to the Cohesity cluster. This provides a continuous window of protection allowing you to restore a VM to a selected point-in-time within the configured CDP window.

Cohesity uses the VMware API for I/O Filtering (VAIO) based solution for providing Continuous Data Protection. In this process, Cohesity deploys an I/O Filter in a vSphere environment. The scope of the filter installation is at a cluster level (also known as ClusterComputeResouce) within the vCenter. Once the filter installation is complete, an I/O Filter is attached to the VM marked for protection. This filter intercepts all the real-time IOs of the VM and sends the IOs to the Filter Daemon running on the ESXi host. The Filter

Daemon continuously replicates the IOs to the Cohesity cluster (IO Journal) based on the VM Storage Policy attached to the VM.

In releases earlier than 7.0, the Cohesity cluster automatically deploys the VAIO filter on the ESXi host clusters. From release 7.0 onwards, Cohesity provides an additional option to download the VAIO from the Cohesity cluster and upload it to an HTTP server to deploy it manually on the ESXi host clusters using the vSphere IO filter installation APIs. For details, see Deploy the VAIO Filter Manually.

You can either manually create or update an existing VM Storage Policy and then attach it to the VMs or Cohesity will automatically create and attach the VM Storage Policy to the VMs you want to perform continuous data protection.

If you are manually attaching the VM Storage Policy, then ensure to:

- Attach a Storage Policy to the VMs before a CDP-enabled Protection Group is created for these VMs.
- While creating the VM Storage Policy, select the **Enable host based rules** checkbox in the **Policy structure** tab and create rules for Replication data service under **Host based services** with the following settings:
  - **Provider**: cohfilter
  - **vCenterUUID**: <vCenterUUID>
  - **CohesityID**: <ClusterId:IncarnationId>
  - **CohesityClusterEndpoint**: <endpoint>:20004
  - **HeartbeatInterval**: 30

> **Note:** You can view the default Storage Policy created by Cohesity on the VMware vCenter as a reference for the above settings.

During the backup process, all the IOs of a workload will be saved in a write-ahead-log (WAL), and is used to recover the applications at any point-in-time as long as all the previous IOs are available. If any IOs are missed during the CDP, then a VADP backup is triggered, and the CDP backup process is re-initialized. For better RTO, the hydrated snapshots of the VMs are also captured to the Cohesity cluster every 15 minutes.

> **Note:** When using VMware CDP, disk exclusion settings do not prevent data capture on excluded disks. CDP operates at the I/O level and captures changes across all disks in the VM. CDP backups may fail when disk exclusion is used.

Core Components

This section lists the core components of the Cohesity CDP.

- **I/O Filter**: An I/O Filter is run per VM on the ESXi host as part of the Virtual Machine Executable (VMX) process. It intercepts the IOs of a VM synchronously in real-time and replicates those IOs to the Cohesity cluster.
- **I/O Filter Daemon**: An I/O Filter Daemon, also implemented using VAIO, provides more resources to perform replication, caching, batching, and other operations. It runs in the ESXi host environment and communicates with individual VM's I/O Filters over a Unix socket. An I/O Filter Daemon is ideal for replicating collected IOs from VMs over the network asynchronously.
- **Cohesity Data Cloud**: The CDP service starts on all the Cohesity cluster nodes and receives the Filter Daemon requests.

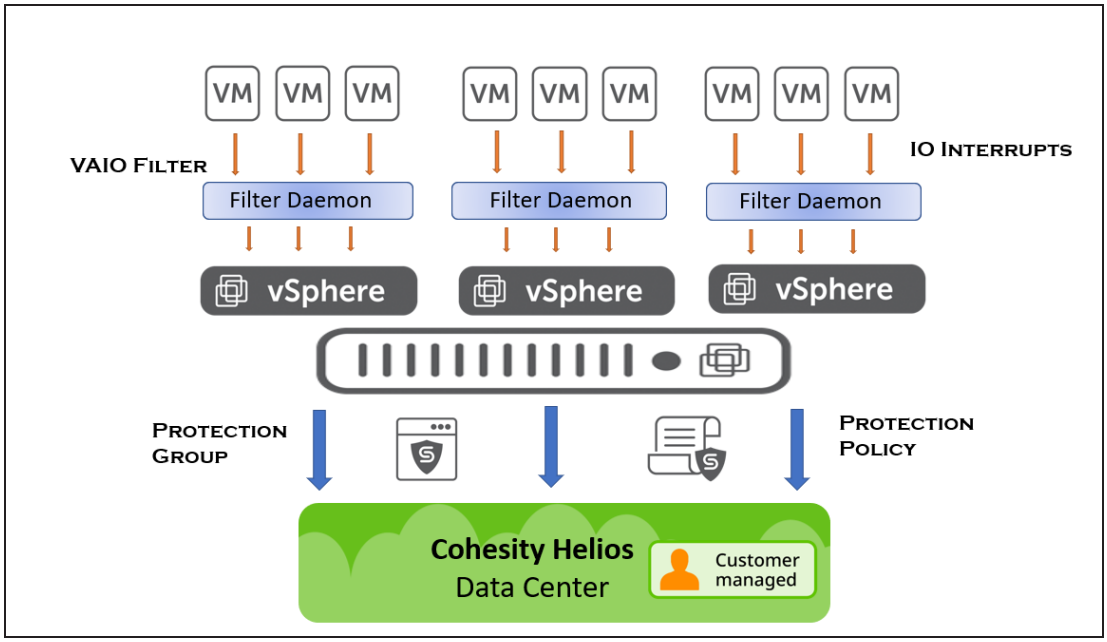

CDP Requirements for VMware vCenter

If you are enabling CDP for near-zero RPO, ensure that the VMware vCenter meets these additional requirements:

Requirements	Details
ESXi version	VMware ESXi 8.0.3, 8.0.1, 7.0 U3, 7.0 U2, 7.0 U1, 7.0, 6.7 U3, 6.7 U2, 6.7 U1, 6.7, 6.5 U3, 6.5 U2
Memory required per VM disk	1 GB (from the heap)

Requirements	Details
Memory required by Daemon	4 GB
ESXi Hosts	All the ESXi hosts within the Cluster Compute Resource should be on the same VLAN.
Upgrade filter	For automatic upgrades or uninstallation of filters, a VMware cluster should have Distributed Resource Scheduler (DRS) enabled with shared storage.

CDP Related Permissions

For CDP, in addition to the required VMware vCenter Related Permissions, you must also require the following VMware privileges:

Privilege Level		Permissions Required	Description
Host	Configuration	Maintenance	To move hosts automatically to maintenance mode to uninstall CDP filter.
		Query patch	To deploy CDP Filter to the host.
Profile-driven Storage		Profile-driven storage update	To create and update Storage profiles to enable CDP for virtual machines.
		Profile-driven storage view	To view defined storage capabilities and storage profiles in order to manage CDP.

Enable Cohesity CDP on VMs

To enable CDP on VMs, create a new protection policy and attach the Continuous Data Protection option to the policy. Within the CDP option specify the retain point in time configuration, and associate this policy with the VM’s Protection Group.

After associating the protection policy with the Protection Group, the protection run will first initiate a full VADP-based backup and, in the process, attach the I/O Filter to the ESXi host. An incremental protection run follows the initial full backup, and the VM's continuous protection will begin from this point on. You can view all the continuously protected VMs from the same Dashboard page.

Steps

Perform these tasks to enable CDP for VMs:

1. Create a Protection Policy with the **Continuous Data Protection** option.
2. Associate this policy with the VMs' Protection group that you plan to protect using CDP. See Add or Edit a Protection Group for Virtual Servers.

Create a Protection Policy

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. Also, for Disaster Recovery, you can add CDP and replication for the Protection Group using the Protection Policy.

To create a Protection Policy with CDP & Replication:

1. Navigate to **Data Protection** > **Policies**.
2. Click **Create Policy** located at the top right of the page.

   The Create Protection Policy page appears.
3. In the **Policy Name** field, specify a protection policy name. The name can contain alphanumerics, underscores, hyphens, periods, and spaces. This field is required and can be changed later.
4. In the **Incremental Backup** field, specify when backups are captured, and how long they are retained.
5. Click **More Options**.
6. Click the **Continuous Data Protection** option from the floating menu. In the **Retain Point in Time for** field, enter a number and select Minute, Hour, or Day from the drop-down. This value defines the period the Cohesity cluster will retain the captured CDP logs. You can use this option to restore the VM applications to a point-in-time as opposed to periodic snapshots available with regular protection runs.

   When CDP is enabled for a protection policy, VM IOs are captured continuously and streamed in real time to the local cluster. As a result, data protected by CDP-enabled policies are available on the local cluster with near-zero RPO.
7. Click the **Replication** icon to replicate snapshots with CDP to another Cohesity cluster. The copied snapshots are stored in the storage domain of the remote cluster. (The storage domain, however, is specified in the Protection Group.) The replication schedule for copying Snapshots cannot be more often than the protection schedule for capturing Snapshots. For example if Snapshots are captured hourly, the Cohesity cluster cannot replicate them every 30 minutes but it can replicate them hourly or any value over one hour.

   When **Continuous Data Protection** and **Replication** are enabled, VM IOs are captured real-time and replicated periodically (every 15 minutes by default) on the remote cluster.

a. **Replicate to**: Select either a Remote Cluster or AWS.

   From the drop-down, select either a cluster source or an AWS S3 storage source to replicate snapshots. If you select an S3 storage object, you also need to select a region.

   You can also register a new cluster to replicate to by selecting **Register Remote Cluster**.

   From Cohesity release 7.0 onwards, you can enable **CDP Synchronous Replication** for a Protection Group. Once enabled, VM IOs of every VM in that Protection Group are streamed directly to the remote cluster (continuous streaming, not periodic).

   As a result, Synch Replication makes data available for immediate (near-zero RPO) recovery. For details on enabling **CDP Synchronous Replication**, see Add or Edit a Protection Group for Virtual Servers section.

   > **Note:** Irrespective of whether you use an existing remote cluster or register a new one, you must enable the **Encryption** toggle in the remote cluster if you have enabled the **CDP Synchronous Replication** field in the Protection job to replicate the real-time IOs and logs to the remote cluster synchronously. For details on the **CDP Synchronous Replication** field, see Add or Edit a Protection Group for Virtual Servers. For detailed instructions on registering a remote cluster and enabling encryption, see Create a Connection to the Remote Cohesity Cluster.

b. **Retain for xx days**: Specify how long the snapshots are retained on the remote cluster.

To add additional replication schedules, repeat these steps. To replicate to multiple Cohesity clusters, create multiple replication schedules with different replication targets.

If you specify multiple replication schedules with the same replication target that replicate the same set of snapshots but with different retention periods, the snapshots are retained for the longest time period. For example if you define two replication schedules with the same replication target that retain the same set of snapshots but one schedule specifies a 90-day retention period and the other schedule specifies a 180 day retention period, the set of snapshots is retained for 180 days.

8. Click **Save**.

The policy is immediately available to use in Protection Groups. If you want to configure additional policy settings, see Create or Edit a Standard Policy.

> **Note:** If replication is enabled for CDP, then the blackout window option in the policy would not be honored.

Recovery

You can perform a point-in-time recovery of VMs, files and folders, Virtual Disks, and IVM of a CDP-based Protection Group s by selecting the points in the timeline bar on the recovery page. This timeline will also show when the VADP backups were taken. So, the user can optionally recover from the VADP snapshots as well. For detailed instructions on recovering VMs, files and folders, Virtual Disks, and IVM, see Recover Virtual Machines.

Deploy the VAIO Filter Manually

From release 7.0 onwards, you can deploy the VAIO filter manually on the ESXi host cluster. To deploy the VAIO filter, use the VMware vSphere IOFilterManager APIs that can be accessed through the vCenter MOB interface.

The vibURL required to deploy the VAIO filter is generated by downloading the VAIO installer from the Cohesity cluster and directly hosting the VAIO installer on an HTTP server accessible from the vCenter server and the ESXi hosts where the VAIO filter needs to be installed.

Prerequisites

- Ensure the VMware vCenter meets the CDP requirements.
- Ensure the VMware vCenter meets the conditions and permissions required to install the I/O filers in its ESXi host clusters.
- Uninstall existing I/O Filters from the ESXi hosts and the ESXi host clusters.

> **Note:** The VAIO filter can only be installed on the ESXi host cluster and not on the ESXi host.

Download the VAIO Filter from the Cohesity cluster

To download the VAIO filter from the Cohesity cluster:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Cohesity Agent Download**.
3. In the **CDP VMware IO Filter** field of the **Download Agents** dialog, perform the following:

1. Use the drop-down that displays the default ESXi version, **6.7**, to pick the ESXi version of the host in which you want VAIO deployed. Clicking on the drop-down displays both **6.7** and **7.0**. Select **6.7** or **7.0** based on the following:
   - If all the ESXi hosts in the cluster are of versions earlier than 7.0, choose **6.7**.
   - If all the ESXi hosts in the cluster are of versions later than 7.0, choose **7.0**.

     > **Note:** If your cluster has 6.7 and 7.0 ESXi hosts, contact your Cohesity account team to enable this option.

2. Click the **Download** icon ( ).

The VAIO filter's installer (the VIB-vSphere Installation Bundle) is downloaded to your local system. Generate the vibURL required to deploy the VAIO filter in the vCenter MOB interface by directly hosting the downloaded installer on an HTTP server accessible from the vCenter server and the ESXi hosts where the VAIO filter needs to be installed. Your Cohesity cluster can also act as the HTTP web server.

Generate the URL from the Cohesity cluster

To generate the URL for the VAIO filter:

1. Construct the URL as follows:

   Format

   ```
   <URL of the HTTP server or the Cohesity cluster>/files/bin/installers/
   ```

   Example

   ```
   http://10.2.111.11/files/bin/installers/
   ```

   When you enter the constructed URL in the address bar of the browser, all the I/O filters available on the Cohesity cluster are displayed.

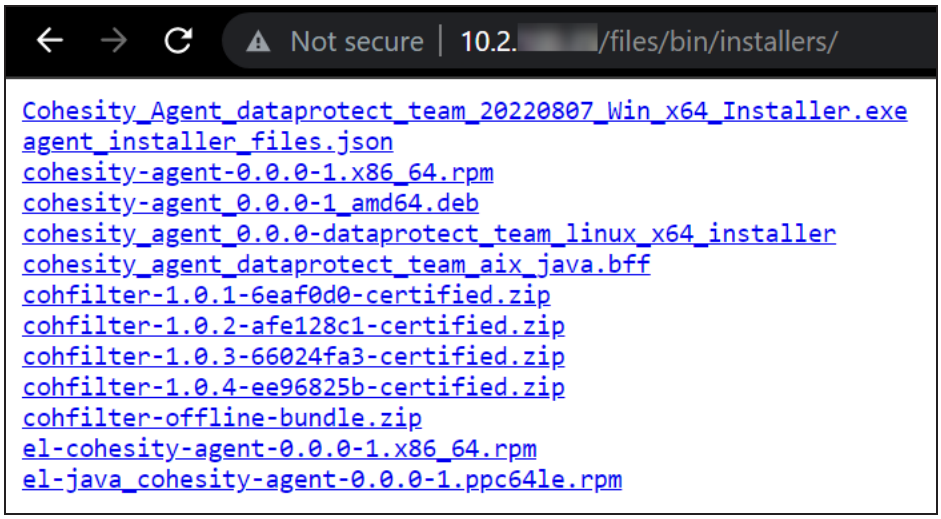

2. From the displayed list, decide which I/O filter to download based on the ESXi version of the host in which you want VAIO deployed:
   - If all the ESXi hosts in the cluster are of versions earlier than 7.0, choose *cohfilter-1.0.4-ee96825b-certified.zip*
   - If all the ESXi hosts in the cluster are of versions 7.0 or later, choose *VMW-esx-7.0.0-coh-cohfilter-1.0.3-44d0c2a2-certified.zip*
   - The same driver version is applicable for both ESXi 7.0 and 8.0.x.
3. To copy the link to the clipboard, right-click the I/O filter and select **Copy Link Address**.

   You use the copied link address to install the VAIO filter in the VMware MOB interface.

Install Cohesity CDP I/O Filter on VMware vSphere Manually

Prerequisites

- **Cohesity I/O Filter Bundle**
  To obtain the HTTP URL I/O filter bundle, follow the steps outlined in Generate the URL from the Cohesity cluster.
- **Cluster MOID**
  Identify the Managed Object ID (MOID) of the target vSphere cluster where the I/O filter will be installed.

To install the Cohesity CDP I/O Filter on VMware vSphere manually:

1. Access the VMware MOB.
   a. Open a web browser and navigate to the vCenter MOB interface using the following URL format *https://<vcenter-url>/mob*.
   b. Log in with a vCenter user account that has the necessary privileges.

2. Navigate to **IOFilterManager**.
    a. On the MOB home page, click on the value associated with the content property.
    b. On the next page, click the link for **IOFilterManager**.
3. Initiate I/O filter installation.
    a. On the **IOFilterManager** page, locate and click the **InstallIoFilter_Task** method.
    b. A new window will appear when the method is invoked. Enter the following parameters:
        i. **vibURL (string)**: Enter the full HTTP URL of the staged Cohesity I/O filter bundle .zip file.
        ii. **cluster (ManagedObjectReference:ClusterComputeResource)**: Enter the clusterMOID of the target cluster (for example, domain-c10).
    c. Click the **Invoke Method** to start the installation.
4. Monitor the Installation.
    a. After invoking the method, the MOB displays a task ID.
    b. Click the returned task ID to view task details. Monitor the status field to track progress.
    c. The installation completes successfully when the task status changes to success.

Steps to Verify Installation

To confirm installation at the ESXi cluster level and driver installation on each host:

1. Access the cluster user interface.
2. Navigate to the ESXi cluster where the I/O filter has been installed.
3. Click on the **Configure** tab.
4. Select **IO Filters**.
5. **Verification**: If you see an entry that begins with *cohfilter*, the installation was successful at the ESXi cluster level.

Upgrade CDP I/O Filter

To upgrade the Cohesity I/O Filter from the Cohesity cluster:

1. Navigate to **Data Protection** > **Sources**.

    The Sources page displays a list of sources registered with the Cohesity cluster.
2. Select the VMware vCenter Server that has CDP enabled for its protected objects.

    The Protected Objects tab displays the list of protected objects.

3. Select the All Objects tab and select the Physical Hierarchy View.
4. Navigate to a CDP-protected ClusterComputeResource. Optionally you can use the filter option available on the page to view the protected ClusterComputeResource.

   An upgrade icon ( ) is displayed next to the CDP-protected object if an I/O Filter upgrade is available.

5. Click the upgrade icon ( ) next to the protected object, and then select **Upgrade CDP I/O Filter**.

   The Upgrade CDP I/O Filter page appears.
6. Click **Continue**.

> **Note:**
> - Cohesity will upgrade the I/O Filters on all the ESXi hosts and VMs belonging to the same VMware ESXi cluster of the selected object.
> - Cohesity will perform the upgrade of the I/O Filter on the ESXi hosts only if the ESXi hosts in the cluster are in maintenance mode. Cohesity will withhold the upgrade till the ESXi hosts go to maintenance mode.

Uninstall CDP I/O Filter

Ensure these prerequisites are met before you uninstall the Cohesity I/O Filters attached to VMs on the ESXi hosts:

- If the VMs are not migrated to another cluster that has a VM Storage Policy with Cohesity I/O Filter, then ensure that the Cohesity I/O Filter is not attached to any VMs.
- This procedure requires the ESXi hosts to be in maintenance mode. You should shut down the VMs to put the ESXi hosts in maintenance mode.
- If the ESXi hosts are set up with DRS mode, then moving an ESXi host in maintenance mode migrates the VMs to another ClusterComputeResource which has the same filter installed. You can perform a rolling uninstallation to remove the I/O Filter associated with all the hosts to remove the filter definition at the cluster level.

After detaching the I/O Filter from all the VMs, it is recommended to remove the VM Storage Policy from the vSphere created by the Cohesity cluster.

To uninstall the Cohesity I/O Filter from the Cohesity cluster:

1. Navigate to **Data Protection > Sources**.

   The Sources page displays a list of sources registered with the Cohesity cluster.

2. Select the VMware Cluster that has CDP enabled for its protected objects.

   The Protected Objects tab displays the list of protected objects.
3. Select the **All Objects** tab, and select the Physical Hierarchy View.
4. Navigate to a CDP-protected object. Optionally you can use the filter option available on the page to view the protected object.

   A filter icon ( ) with the following message displayed next to the CDP-protected object if an I/O Filter upgrade is available.

   ```
   Uninstall CDP IO Filter
   ```
5. Click the icon ( ) next to the protected object, and then select **Uninstall CDP IO Filter**.

   The uninstall CDP IO Filter page appears.
6. Click **Continue**.

> **Note:** Once initiated, the Cohesity VMware I/O filter installation will continue to completion until CDP is removed from a protection policy.

Considerations

- If a new host is added to the Cluster Compute Resource of the VMware vCenter, CDP will be auto-deployed on the newly added host and the I/O Filter Daemon will start running on it.
- The Cohesity cluster will continue to protect the CDP-enabled VMs that are moved (or migrated) to another ESXi host using VMware Storage vMotion.
- You can perform a point-in-time recovery of VMs, files and folders, Virtual Disks, and IVM of a CDP-based Protection Groups by selecting the points in the timeline bar on the recovery page. This timeline will also display when the VADP backups were taken and you can optionally recover from the VADP snapshots. For detailed instructions on recovering VMs, files and folders, Virtual Disks, and IVM, see Recovery.
- If you are manually attaching the VM Storage Policy, then ensure to:
  - Attach a Storage Policy to the VMs before a CDP-enabled Protection Group is created for these VMs.
  - While creating the VM Storage Policy, select the **Enable host based rules** checkbox in the **Policy structure** tab and create rules for Replication data service under **Host based services** with the following settings:

  - **Provider**: cohfilter
  - **vCenterUUID**: <vCenterUUID>
  - **CohesityID**: <ClusterId:IncarnationId>
  - **CohesityClusterEndpoint**: <endpoint>:20004
  - **HeartbeatInterval**: 30

> **Note:** You can view the default Storage Policy created by Cohesity on the VMware vCenter as a reference for the above settings.

- If there is no VM Storage Policy attached to the VMs, then when the Protection Group is run, Cohesity will automatically create and attach the VM Storage Policy required for CDP.
- A CDP-enabled Protection Group supports a maximum of 25 VMs per node and 40 VMs per ESXi host, in each Cohesity cluster. Attempts to create Protection Groups that exceed this limit fail with the following error messages:

```
The limit on the number of CDP-protected VMs from an ESXi host is being exceeded on the host.
```

```
The maximum number of VMware VMs supported on this cluster for CDP is <number of nodes in the Cohesity cluster times 25>
```

- If the Storage Policy attached to the VM does not have the correct CDP attributes configured, then a corresponding error message is displayed for that VM on the CDP status page (**Protection> CDP** > **Protected Objects**).
- Do not make any updates to the Cohesity-managed I/O Filter or the VM Storage Policy directly from vCenter as this will cause the CDP to fail.
- CDP will stop if the FQDN of the Cohesity cluster is changed after the CDP protection runs are configured for a VMware VM. To reenable CDP, update the new FQDN in the CDP Storage policy and reapply the policy on all the VMs associated with that policy.
- If a Protection Group is deleted or paused, then you must either manually detach the VM Storage Policies that were manually attached or update the Storage Policies to remove the CDP attributes. Once you resume the Protection Group, you can reattach the VM Storage Policies with CDP attributes.
- If the time stamp on the Cohesity cluster and the workload VM are not in sync, then the CDP protection window displayed in the Cohesity cluster will be inaccurate.
- You can select point-in-time restore only for individual VMs and not for a Protection Group.
- The performance of your vCenter server might be impacted if many VMs are added to a CDP policy at once.

- Cohesity clusters monitor the I/O throughput of CDP-enabled objects continuously. The guardrails help you track the resource utilization, performance, and stability of the Cohesity cluster.

  When the I/O throughputs exceed the pre-defined thresholds, appropriate alert messages are displayed under the Message column of the Protection Run details page.

  - **CDP disabled for a set duration**. If the I/O throughput of an object sporadically increases above the threshold, causing a resource scarcity on the Cohesity cluster, CDP is disabled for a duration determined by Cohesity.

    Alert Message

    CDP was temporarily disabled at <hh:mm> because of High IO Throughput. It will be enabled automatically after <default period set by Cohesity>.

    CDP is automatically re-enabled once the time duration set by Cohesity expires. To modify the time duration set by Cohesity, contact your Cohesity account team.

  - **CDP disabled temporarily**: If the I/O throughput of a CDP-enabled object exceeds the threshold for a prolonged period, CDP is disabled until you enable it again.

    Alert Message

    CDP disabled at <hh:mm> PM because of continuous high I/O throughput. Please review the I/O throughput of an entity.

    Disabling CDP ensures the sustainability of the Cohesity cluster and prevents the objects that exceed the thresholds from affecting the other objects in the Cohesity cluster.

    When CDP is disabled with the preceding message, it is not enabled unless you manually enable it. You can review the I/O throughput of such entities and re-enable CDP if you like. You may also contact your Cohesity account team for more details.

- CDP protection is not supported for the VMware VMs residing in the vVol datastores.
- For CDP-enabled VMware VMs that use vSAN-based datastores, you must attach the storage policy manually. For details, see CDP Offline Storage Attach for vSAN Datastores.
- With ESXi 8.0.x versions, the automatic installation of the filter driver does not work as expected.

Failover Protection Group for VMs

> **Note:** Failover of Protection Groups is only supported for VMware VMs.

When a Protection Group with a replication schedule starts replicating snapshots from a primary Cohesity cluster to a remote Cohesity cluster, a Protection Group in Failover Ready state is created on the remote Cohesity cluster. If the original capturing Cohesity cluster and/or the original Source are down because there is a planned maintenance outage or a disaster destroys them, you can failover the Protection Group in Failover Ready state and deactivate the original task on the capturing Cohesity cluster. As part of the Failover, you are prompted to optionally recover the VMs from the replicated snapshots stored on the remote Cohesity cluster.

> **Note:** When you create a Protection Group on a Cohesity cluster with a replication schedule after the first snapshot is replicated, a Failover Ready copy of the Protection Group is created on the remote Cohesity cluster that is receiving the snapshot. For more information on Failover Ready, see About Failover Ready. You can failover the Protection Group of VMs. You cannot failover the Protection Group of Views or SQL Servers. You can, however, clone replicated Views and restore SQL Servers. For more information, see Failover a Protection Group.

During Failover, the following actions occur:

1. **Disaster at Source** - On the primary Cohesity cluster, the backup job fails to capture the snapshots of the source.
2. **Replication Fails** - As the source is down or unreachable, the replication job fails to replicate the snapshots onto the remote Cohesity cluster. The Failover Ready for VMs on the remote Cohesity cluster is activated and causes the rejection of the incoming replicated snapshots created by the associated job on the primary Cohesity cluster.
3. **Failover VMs** - You are prompted to optionally recover the VMs from the replicated snapshots on the remote Cohesity cluster to a new source.

The state of a Protection Group may also become "Failover Ready" when an active Protection Group is deactivated. For more information, see About Failover Ready.

For more information about what happens during a Failover, see Failover from Replicated Copies.

For more information on disaster recovery of VMs using the Cohesity Cloud Edition, see Back Up VMs to On-Prem, Replicate and Failover to AWS Cloud.

To Failover a Protection Group on the remote Cluster:

1. Navigate to **Data Protection > Protection**.
2. Find the Protection Group in the Failover Ready state in the list, click on the actions menu ⋮, and select **Failover**.

3. Select a source to place the recovered VMs and click **Failover Job and Continue to Recovery**.

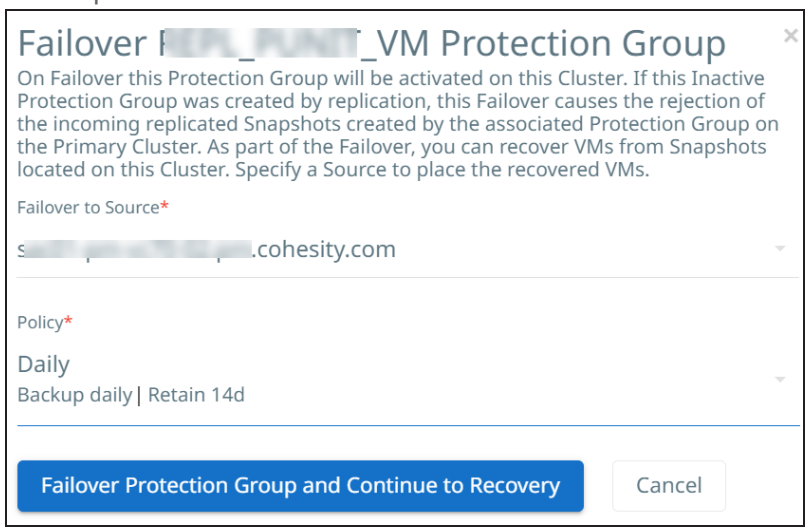

4. Set Recover Options for VM Failover.
5. View VM Failover Confirmation.
6. Deactivate Original VM Protection Group. The Failover of the Protection Group in the Cohesity Dashboard does not automatically deactivate the original task on the original cluster that captured the snapshots. To deactivate, you must follow the instructions in Deactivate Original VM Protection Group
7. Restore VM Protection Group to Running. The Failover of the Protection Group in the Cohesity Dashboard does not automatically restore the task to running on the new cluster. To restore the task to a running state, you must follow the instructions in Restore VM Protection Group to Running.

Set Recover Options for VM Failover

In this section, you select the recover options while recovering VMs as part of the Failover of an Inactive Protection Group.

To set recover options:

1. Toggle **Rename Recovered VMs** on if you want to add Prefix and/or Suffix strings to the VM names of the new VMs created by this task. For example, you could specify Prefix `Test_` and Suffix `_QA` to the VM name `WinOS` and when the VM is recovered, the new VM that is created is called `Test_WinOS_QA`.
2. Select the appropriate **Recovery Location** option
   - **Recover back to original location**—Select this option to recover the VM files (such as the VMDK files) to their original datastores and create new instances of the VMs in the original location in the original source. For more information, see Recover to the Original Location.

- **Recover to a new location**—Select this option to recover the VM files (such as the VMDK files) to an alternate datastore and create new instances of the VMs in a new location of a registered Source. For more information, see Recover to a New Location. For details about virtualization or cloud platform-specific options, see the respective section.
  - Acropolis
    a. **Source:** Select a registered Source to recover the VMs to.
    b. **Storage Container:** Select a storage container to store the recovered VMs.
  - AWS
    a. **Source:** Select a registered Source to recover the VMs to.
    b. **Region:** Select a region in the project to recover the VMs.
    c. **VPC:** Select an Amazon Virtual Private Cloud (VPC) to store the recovered VMs.
    d. **Subnet:** Select a subnet in the Amazon VPC to store the recovered VMs.
    e. **Key Pair:** Select an Amazon key pair that has permission to log in to the recovered VMs.
    f. **Network Security Group:** Select the associated network security group.
  - Azure
    a. **Source:** Select a registered Source to recover the VMs to.
    b. **Resource Group, Storage Account:** Select the corresponding Resource Group and Storage Account to which the recovery has to be done.
    c. **Storage Container:** Select the storage container in which page blobs will be created.
    d. **Virtual Network:** Select the virtual network where the VM must be placed.
    e. **Subnets:** Subnets in the virtual network where VMs must be placed.
  - GCP
    a. **Source:** Select a registered Source to recover the VMs to.
    b. **Project:** Select a Google Cloud Platform (GCP) project to recover the VMs to.
    c. **Region:** Select a region in the project to recover the VMs to.

d. **Availability Zone:** Select an availability zone in GCP to recover the VMs.

e. **VPC:** Select a Virtual Private Cloud (VPC) to store the recovered VMs.

- Hyper-V
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Compute:** Select the compute host for the recovered VMs.
  c. **Volume:** Select a volume to store the recovered VMs.

  When recovering to a standalone Hyper-V host as a new location, SMB volumes are not supported locations.
- RHV
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Data Center:** Select a data center to store the recovered VMs.
  c. **Cluster:** Select an RHV/OLVM to store the recovered VMs.
  d. **Storage Domain:** Select an RHV/OLVM storage domain to store the recovered VMs.
- VMware
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Resource Pool:** Select a Resource Pool to store the recovered VMs.

  > **Note:** When protecting a VMware Source, if you select a Resource Pool that is not part of the ESX cluster, all datastores in the vCenter are displayed. Ensure you select a datastore that is part of the selected Resource Pool. If you select the wrong datastore, storage migration fails, but the VM remains attached to the ESX Host and runs from the Cohesity datastore. You can manually migrate the VM to the desired datastore at a later time.

  c. **Datastore:** Select one or more available target datastores. The VM files (such as VMDK files) extracted from the Snapshots are stored in the selected datastores. The number of target datastores you select must be equal to or less than the number of VMs selected for recovery.

Cohesity also displays each datastore's amount of free space and total capacity. For more information about recovering VMs to multiple datastores, see Recover VMs to Multiple Datastores.

> **Note:** Not applicable when recovering to a standalone ESX host. After recovery, VMs continue to run on the Cohesity datastore because Storage vMotion is not supported.

d. **VM Folder** Optionally select a logical VM folder in the datacenter in which to create the VMs.

> **Note:** Not applicable when recovering to a standalone ESX host.

- vCloud Director
  a. **Source:** Select a registered vCloud Director source to which you plan to recover the VMs or vApp.
  b. **Organization VDC:** Select the target organization VDC.
  c. **vApp:** Optional. Select the target vApp. If not specified, VMs are recovered as standalone VMs inside the target organization VDC.
  d. **Storage Profile:** Select the target Storage Profile. The Cohesity cluster chooses any compliant datastore associated with the Storage Profile to store the VM files. If you are recovering multiple VMs, the Cohesity cluster chooses one or more datastores associated with the Storage Profile.

3. Select the appropriate **Recovery Type** option.
   - **Instant Recovery**: if you select this option, the VM(s) will be usable instantly in the target environment and will be moved to target storage later. The VM can be used in production during Instant Recovery when data is copied to the VM from the Cohesity cluster. In such scenarios, the VM performance may be low as the Storage vMotion is still in progress.

     When recovering VMware VMs, by default, the selected VM is always instantly available after recovery.

     > **Note:** For instant recovery to Standalone ESXi, the VM(s) will be available immediately in the target environment and will have to be moved to the target storage manually.

- **Copy Recovery**: You can select this recovery type for **VMware** sources only. If you select this option, the recovered VMs will be usable in the target environment only after all the data is copied from Cohesity to the target storage. Once the restore is complete, the VM will be available to use and the performance of the VM will be normal.

  > **Note:** You should have `VirtualMachine.Config.Rename` permission on your user account to use this recovery option.

  Considerations when using the Copy Recovery option to recover VMware sources

  - This is the only option to recover data from VMware cloud solutions.
  - Copy recovery is not supported on Hybrid Extender deployments. Only vCenter and stand-alone ESXi hosts are supported.

4. Specify the networking options.
   The network options displayed depend on the Recovery Location option that was selected.

   - **Keep original**—For each recovered VM, keep the original virtual Network Interface Cards (vNICs) and attach them to the original network connections. This option is only supported when VMs are recovered back to their original location.
   - **Start Connected**—For each recovered VM, connect to the original or new network when the VM reboots. If this option is not selected, the VMs are not connected to any network on reboot.
   - **Detach network**—For each recovered VM, the virtual Network Interface Card (vNIC) is removed from the VM.
   - **Attach to a new network**—For each recovered VM, attach the virtual Network Interface Card (vNIC) to a new network. The new network might be a standard port group or a distributed port group.

     When recovering VMware VMs, the selected Network port group will be attached to all the network interfaces of the VMs being recovered.

     When recovering Hyper-V VMs, the selected virtual switch will be attached to all the network interfaces of the VMs being recovered.

> **Note:**
> - Recovering a network attached CentOS VM on VMware or Hyper-V does not attach the VM to the network. Attach the VM to the network manually.
> - If the VMware Resource Pool is not part of the ESX cluster, all networks available to the Datacenter are displayed. Ensure you select a network available to the selected Resource Pool. If you select the wrong network, VM creation and recovery will fail.

5. **Auto Select:** When enabled, the recovery task automatically selects the correct VLAN.
6. Select **Instant Recovery** if you want to change the Hyper-V recovery mechanism so that VMs become available instantaneously after recovery. Applicable to Hyper-V only.

   Select **Preserve VMs UUID** if you want to preserve the UUID for the recovered VMs. Applicable to Hyper-V only.
7. Select **Leave recovered VMs powered off** if you want the recovered VMs to remain powered off after they are created.

   If you recover VMs from a Storage Domain that has Cloud Tier enabled, Cohesity recommends leaving the VMs powered off.
8. Select **Continue recovery even if errors occur when recovering VMs** if you want the Recover task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cluster will still attempt to create the other VMs. If this is unselected, then the Recover task is discontinued only if VM creation fails, so if VM creation succeeds, but VM storage migration fails, the task continues.
9. In the Selected Objects section, click **Show VMs** to view the VMs associated with the Failover Job. By default, the VM Snapshots of the last Run of the Job are recovered.

   Click Edit ( ) to select Snapshots captured by a different protection run.
10. Click **Finish**.

The next step is View VM Failover Confirmation.

View VM Failover Confirmation

In this step, you can view a summary of the options for the current Failover task and a list of objects that will be recovered by the task.

The next step is to Deactivate Original VM Protection Group

Deactivate Original VM Protection Group

Only Protection Groups that are protecting VMs can be deactivated. Protection Groups that are protecting a view cannot be deactivated.

> **Note:** If possible, deactivate the original Protection Group on the original capturing Cohesity cluster. If the original Cohesity cluster that captured the Snapshots is down or unreachable, you may be unable to deactivate the Protection Group.

The example below shows the deactivation of a Protection Group protecting VMs. Deactivating a Protection Group protecting SQL Servers is similar but not all details are the same. For example, the transaction logs are not described or pictured.

During deactivation, the following actions occur:

1. Snapshots are no longer captured by the original Protection Group on the original capturing Cohesity cluster. However, the existing snapshots captured by the Protection Group remain in the storage domain. They are not deleted.
2. Snapshots created by the Protection Group are no longer replicated to the receiving (remote) Cohesity cluster. (If you have already initiated the failover procedure, the replication has already stopped.)
3. You can optionally choose that the VMs on the original source be powered off. The VMs and the VM files, however, are not deleted. (The Cohesity cluster may be unable to power off the VMs if the original source is down or unreachable.)

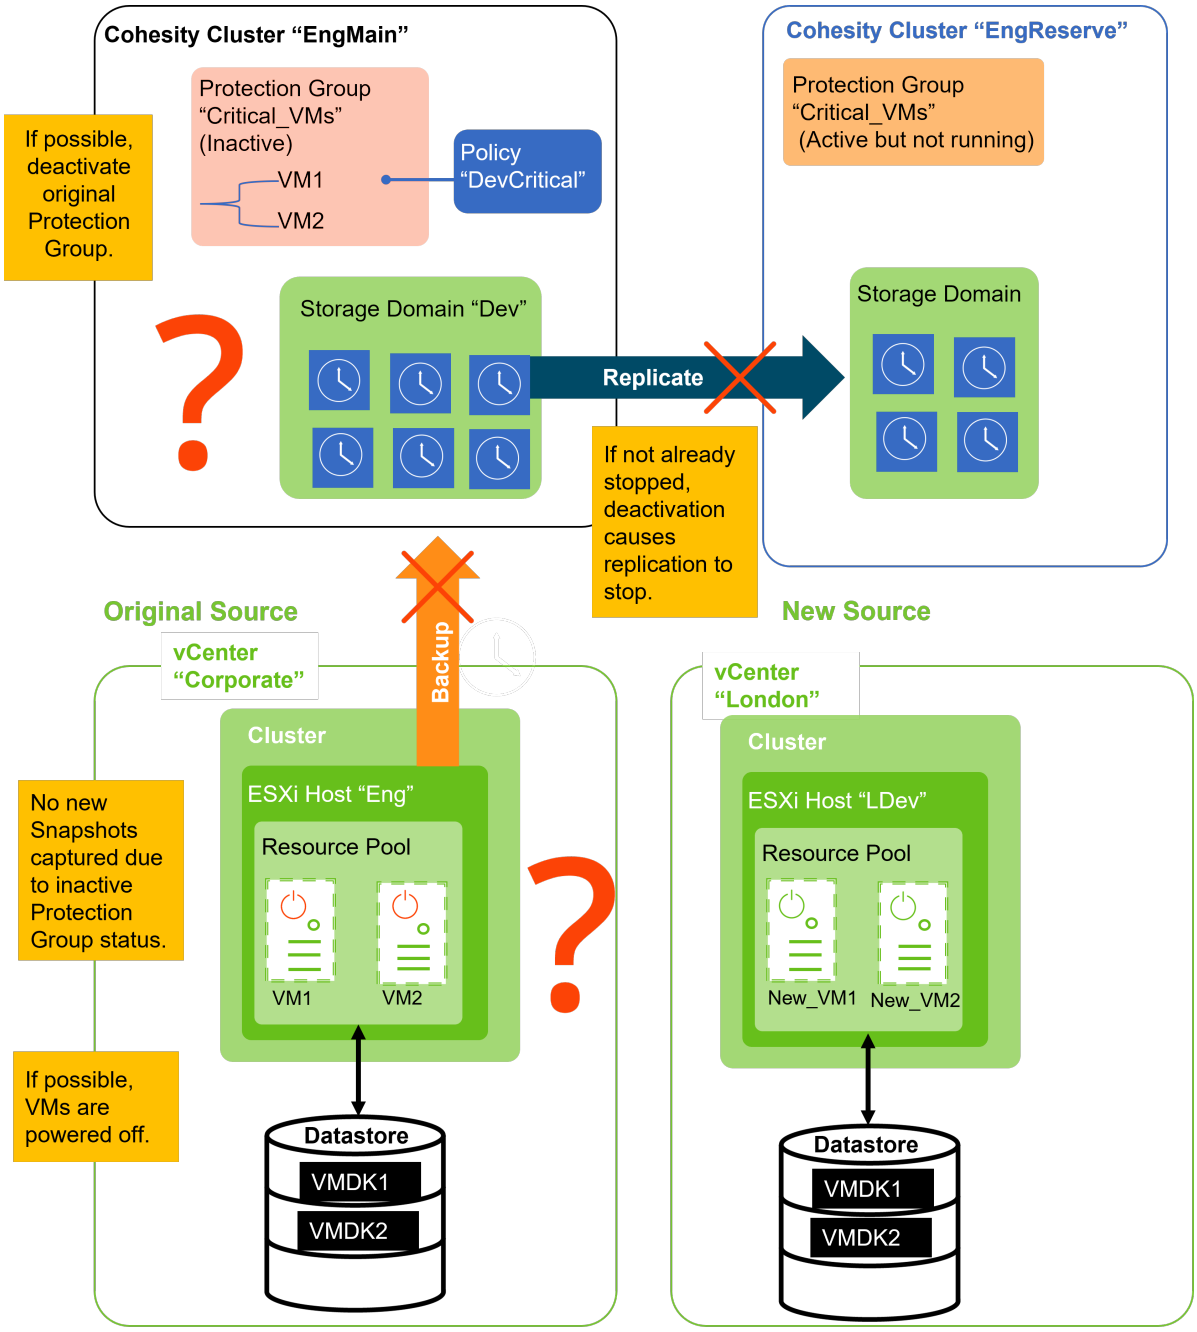

To deactivate the original Protection Group:

1. Log in to the Cohesity cluster where the original Protection Group was running.
2. Navigate to **Data Protection > Protection**.
3. Find the original Protection Group in the list, click on the actions menu ⋮ and select **Deactivate**.
4. Click **Deactivate and Power off VMs**, or click **Deactivate Only** if you do not want

to power off the VMs.

5. Optionally delete the VMs and VM files using the VMware vSphere client.

Restore VM Protection Group to Running

You can restore the Protection Group activated by the Failover, by assigning a new policy to the Protection Group and selecting Objects (VMs) to back up. These actions restore the Protection Group to a running state on the new cluster. To learn more, see Restore Activated Protection Group.

To restore the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, find the inactive Protection Group (created by replication).
3. From the Actions menu (⋮) of that Protection Group, select **Resume**.

This is the last step in the Failover process.

Viewing VM Snapshot Types

Cohesity records Snapshot consistency information (Application-Consistent, Crash-Consistent, Not Powered On) for each VM that the Cluster backs up. You can view this information on the Run Details page. For information about creating App-Consistent backups, see Creating Application-Consistent Snapshots.

To view a VM's Snapshot type:

1. On the Protection Group page, click the name of the Job that captured the Snapshot.
2. On the Job Details page, click the protection run that you want details for.
3. On the Run Details page, mouse over a VM to see its Snapshot type. A message may also provide additional details about the Snapshot type.

When you create a Recover or Clone task, a Snapshot's type is viewable when selecting the Snapshot to Recover or Clone. To view a VM's consistency information, click Edit. If one type displays, this means all VMs in the snapshot are that type. If the Mixed type displays, this means the snapshot contains VMs of different types, such as crash-consistent and not powered on. You can display individual VMs to see each VM's type.

Note the following about Application-Consistent Backups:

Enabling or disabling the Application-Consistent option for an Incremental (CBT) Protection Group results in the creation of full snapshots (instead of incremental snapshots) during the next Protection Group run.

Truncate Logs in Microsoft Exchange Servers on VMs

Cohesity supports the ability to truncate the Microsoft Exchange email server transaction logs after a successful backup of a Windows server that hosts the Exchange email server. After the Cohesity cluster successfully captures a snapshot during a Protection Group, the Cohesity cluster truncates the Exchange transaction logs on the server.

Cohesity has qualified truncating logs for the following Microsoft Exchange servers running on virtual servers:

- Microsoft Exchange 2019 running on supported Windows servers (Standalone and DAG)
- Microsoft Exchange 2016 running on supported Windows servers (Standalone and DAG)
- Microsoft Exchange 2013 running on supported Windows servers (Standalone and DAG)
- Microsoft Exchange 2010 SP3 running on supported Windows servers (Standalone and DAG)

While selecting VMs to protect, you can toggle **Exchange Log Truncation** on for individual VMs where Exchange is registered with the Cohesity cluster.

> **Note:** Cohesity supports log truncation of Microsoft Exchange servers on VMware VMs only.

If **Exchange Log Truncation** is on and the Windows server has an Exchange server but the log truncation fails, an error is reported. However, the backup can successfully complete.

If **App-Consistent Backups** is on for the Protection Group and **Exchange Log Truncation** is on for some VMs in the job, the backup of these VMs is crash-consistent. However, these VMs contain VSS snapshots that were backed up while the Exchange server was quiesced (in an app-consistent state). You can restore this app-consistent VSS snapshot using the Microsoft VSS tools. Note that other applications in these VMs are not quiesced when VSS snapshots are backed up. These other applications were captured in the crash-consistent state in the VSS snapshot.

> **Note:**
> - Exchange log truncation does not require an additional adapter, but Exchange must be registered on the VM you want to protect.
> - **App-Consistent Backups** must be toggled on if you want to protect an Exchange VM.

Register Exchange on a VM

1. Navigate to **Data Protection > Sources**.
2. Navigate the hierarchy to the VM that is running Exchange.
3. Mouse over the VM's actions menu and click **Register Applications**.
4. Select the Cohesity agent to use, either **Installed Agent** or **Auto-Deployed Agent**.
5. Select MS Exchange.
6. Click **Register**.

   The VM's application and agent information displays when you click the icon.

Enable or Disable Exchange Log Truncation

To enable or disable Exchange Log Truncation, please perform the steps below:

1. Navigate to **Data Protection** > **Protection**.
2. Click the Protection Group (VM that is running Exchange).
3. Click the **Settings** tab.
4. Click **Virtual Machines** and then click **Edit** from the list of objects.
5. Toggle on or off **Exchange Log Truncation**.
6. Click **Save**.
7. Click **Protect**.

Configuring Custom Quiesce Scripts for Linux Virtual Machines

If the **Create App-Consistent backups** option is selected for a Protection Group that are protecting Linux Virtual Machines and quiesce custom pre-freeze and post-thaw scripts can be found on Linux VMs in the locations documented below, these scripts are run during the quiesce and create Snapshot process.

> **Warning:** The quiesce scripts must be able to be run by VMware Tools.

Script Name and Directory Location	Description	Example
`/usr/sbin/pre-freeze-script`	Script is run before the Snapshot is created.	Create a custom script that quiesces applications running on the Linux Guest OS before the Snapshot of the VM is created.
`/usr/sbin/post-thaw-script`	Script is run after the Snapshot has been created.	Create a custom script that restarts applications on the Linux Guest OS after the Snapshot of the VM is created. The script restarts applications that were quiesced by the pre-freeze script.

> **Note:** These are the script locations for Linux VMs. The script locations for Windows are documented in Running custom quiescing scripts inside Windows virtual machine with ESX 3.5 Update 2 and later (1006671) Knowledge Base Article.

Recover Virtual Machines

The VMs that are backed up onto the Cohesity cluster can be recovered as VMs, virtual disks, files and folders. You can also perform instant mount volume of the VMs.

Recover Files or Folders

Cohesity provides the ability to recover files and folders from a snapshot created earlier by a Protection Group. Files and folders can be recovered to their original location (limitations apply to physical servers and specific operating systems) or to a newly specified location, which can be within the original source or a different one. You can choose to retain the recovered files' and folders' original (at the time of the backup) permissions and attributes. You can also download files and folders from selected snapshots that were created by a Cohesity Protection Group.

To learn more, see About File Recovery.

For more information about the recovery options:

- Recover Files or Folders to the Original Location
- Recover Files or Folders to a New Location

Windows VMs

- If the Storage Replica feature is enabled on a Windows server with GPT disks, then file recovery (using the Cohesity agent) or IVM recovery to the original location is not supported. Cohesity recommends the following alternatives:
  - For file-level recovery, consider one of the following:
    - Recover to an alternate location.
    - Recover using VMware tools.
    - Use the Download Files option.
  - For IVM recovery, recover to an alternate location instead.
- After making system configuration changes to a Windows 8 or Windows 2012 System VM, such as renaming an existing drive letter or adding a new disk, these changes may not immediately take effect due to a Windows registry refresh issue. To force the drive letters to be updated on the VM, reboot the system in the VM. This issue affects how files are indexed by the Cohesity cluster and displayed while browsing the contents of the VM in the Cohesity Dashboard.
- Considerations when recovering to physical servers that run:
  - Windows 2012 or later - None

  If the OS does not support your recovery, you must recover to an alternate physical server running Windows 2012 Server or later, or use downloads.
- Windows behavior (giving each user its own set of mapped networked drives) prevents recovering to an SMB share that was mounted by a different user.
- Recovering files/folders with names longer than 200 characters may return an error. This limitation is due to Windows behavior when handling files/folders with long names.
- File-based recovery to Windows VMs does not support hardlinks and alternate data streams.

Linux VMs

- For Linux VMs, Cohesity supports file recovery from LVM volumes. One LVM volume can consume more than one loop back device, so Linux VMs may support fewer than 8 volumes when configured with the default number of loop devices.
- For Linux Logical Volume Manager (LVM), if all the disks for a volume group are not found by the Cohesity cluster, the Cohesity cluster will not process that volume group. As a result of that, no volumes of this volume group will be recognized or indexed by the Cohesity cluster.
- When recovering a Linux file, the Cohesity Linux Agent runs the following commands in sudo:

  - mount
  - umount
  - findmnt
  - timeout
  - blkid
  - lsof
  - ls
  - rsync
  - losetup
  - dmsetup
  - lvs
  - vgs
  - lvcreate
  - lvremove
  - lvchange
- Recovering files to Linux VMs is not supported in the following cases:
  - When run as a non-root user that does not have sudo access.
  - If `ALL=(ALL) NOPASSWD:ALL` is not set for the recover user in the /etc/sudoers file.
  - Recovering to Linux VMs requires requiretty to be disabled for the recover user in the /etc/sudoers file, otherwise recovery will fail.To disable requiretty for a recover user, add the following line in the /etc/sudoers file, where <USERNAME> is the name of the recover user with sudo access:

    ```
    Defaults:<USERNAME> !requiretty
    ```
  - The recovery directory path length is greater than 4096 characters.
  - There is not enough space in /tmp for Cohesity to push linux_agent.
- For RHEL7, if Open VM Tools is installed instead of VMware Tools, TMPDIR may not point to /tmp. When recovering to location "/tmp/<SOME_DIRECTORY>", files may be recovered to a different location.

  Example: If the recovery location is '/tmp/DIR1', files are recovered to a different location, such as '/tmp/systemd-private-c74aea179e9a43c789a19306d880274f-vmtoolsd.service-9GhOBD/tmp/DIR1'

VMware

- File level recovery for VMware ESXi environments does not support RAID-5 volumes on dynamic disks. Simple, striped, spanned and mirrored volumes on dynamic disks

are supported.

VMware Tools

If you use VMware Tools for recovery, ensure the destination VM meets the following conditions:

- The VM is powered on.
- The most recent version of VMware Tools is installed and running on the VM.
- open-vm-tools can be used on the supported Linux releases; open-vm-tools, if used, must be installed and running.
- The user account used to access the VM must be a local user with write permissions for that VM.
- Ensure you run VMware tools in sync with the vCenter and ESXi host versions. For compatibility of VMware Tools with guest operating systems, see the VMware article.
- The use of VMware tools for file-level recovery from a deduplicated Windows volume is not supported.

Considerations

If you recover files and folders of VMware VMs using VMware Tools, then consider the following:

- A VMware Tools service restart during a Recovery operation may disrupt Recovery. If the VMware Tools service restarts during a Recovery operation, the following error message is returned: The guest operations agent could not be contacted. After multiple retries to contact the guest operations agent, an error message stating that it started the copy but it could not get the status is returned. Go to the recovery location to verify whether the operation succeeded.
- The maximum recovery speed is 1-2 Mbps due to a known VMware limitation. For more details, see the VMware article.
- Ensure port 443 is open on the ESXi host that hosts the target VM. For more details, see Manage Firewall Ports.
- Cohesity recommends that you *not* use VMware Tools if the size of the recovered files or folders exceeds 10 GB or 10,000 files.
- For Windows VMs:
  - ACLs cannot be restored even if you have enabled the **Preserve File/Folder Attributes** option. Only the VM data will be restored.
  - Recovery of reparse points (such as shared folders, mount points, or junction points) cannot be restored.
  - When using the VMware Tools option for FLR on Windows VMs, some ACLs cannot be recovered due to limitations in the VMware API.

    - If a Windows VM includes volumes created from a storage pool (Microsoft Storage spaces), FLR is not supported.
- For Linux VMs:
    - The **Preserve File/Folder Attributes** option restores only basic user, group, world permissions, and timestamps. Advanced permissions and file attributes such as ACLs are not restored.
    - Permissions for guest files and folders are retained only when the user running the restore operation has permission to change the group ownership on the restored files and folders. If the user does not have change group ownership permissions, the restore operation will fail. In this case, retry the operation without enabling the **Preserve File/Folder Attributes** option.
    - VMware APIs do not support the creation of files in the root folder.
    - VMware APIs cannot restore symbolic links.
    - Restoring files to a folder is not supported if the user account doesn't have write permission to that folder (even if the user has sudo permissions). Cohesity recommends using a root user or a user account with write privileges to the target folder.
    - If you restore files and folders using a guest user account that does not have permission to change file/folder ownership, the restored files are owned by the guest user account used to perform the restore. If you restore files and folders using a guest user account that has permission to change file/folder ownership, the restored files are owned by (or the ownership is retained by) the guest user account used to create the files and folders.
    - When a parent folder is restored all of its sub folders are restored also; Empty folders can be restored only when the parent folder is restored.
    - Restoring hard links is supported; hard links restored with their source files use the same index node (inode).
    - Agent-based File-Level Recovery (FLR) can fail on Btrfs file systems running Ubuntu 22.x when recovering files to the same location due to a filesystem UUID conflict with the existing mounted filesystem. As a worksround, use tool-based FLR instead of agent-based FLR in this scenario.

Hyper-V

File level recovery for Hyper-V requires the following:

- For Windows target VMs:
    - The VM must be part of an Active Directory (AD). The Hyper-V host must also be in the same AD as the target VM.
    - Remote powershell is enabled from Hyper-V host to VM.

    - The Hyper-V Guest Service Interface is enabled on the target VM: Enable-VMIntegrationService -Name 'Guest Service Interface.
    - The Windows Remote Management (WinRM) service is enabled on the target VM. Users must execute the "winrm quickconfig" command to enable WinRM on the target VM.
- For Linux target VMs:
    - Cohesity Linux agent is installed on the target VM.

Nutanix AHV

File or folder level recovery of VMs in a Nutanix Acropolis Hypervisor requires a Linux or Windows Cohesity agent installed on the target VM. For details on how to install a Cohesity agent on a Nutanix VM, see Download and Install the Agent on Nutanix AHV.

Troubleshooting

- **Error**: When recovering files or folders, the virtual disks are part of the target VM. These virtual disks are attached as SCSI disks that can be any of the supported adapter types: LSI Logic Parallel, LSI Logic SAS or VMware Paravirtual. During this step, you may encounter the following error: "Disk adapter with required slots - <n> is not available. Try creating a new adapter". Here, <n> is the number of virtual disks that are being attached. This can occur if the VM's disk adapter does not have the required number of slots (one SCSI adapter can support 15 virtual disks).

  Solution:Attempt the operation *after* creating a new SCSI adapter. Additionally, the number of virtual disks where files and folders are being recovered from is limited to 15 at a time. Remove some files (or folders) and retry the recovery.
- **Error**: If LVM is not configured, running a protection job fails with the “No LVM volume found in idbmantasdev” error.

  Solution: For details on how to resolve this issue, see the Registering or running a Linux protection job fails when no LVM volumes are detected KB article.

Limitations

- You can download files and folders from selected snapshots that were created by a CohesityProtection Group. However, only items that were indexed when the snapshot was created can be downloaded. Recovering and searching for files and folders in snapshots created by other backup software is not supported.
- Indexing, file recovery and browsing files and folders on VMs are not supported for drives with disk-level encryption (such as BitLocker). On physical Servers, however, these workflows are supported.
- Encrypted folders that have been renamed or deleted cannot be recovered.
- Encrypted VMs are not indexed.
- Recovering files to a VM where vMotion is in process is not supported.

- File recovery is not supported for ReFS volumes in these environments: physical, VMware, Hyper-V and AHV.
- Downloading files and folders from tape archive locations is not supported.
- Recovering files and folders from VMs to physical servers and from physical servers to VMs is not supported.
- The downloadable zip file can contain regular files and folders only; symlinks are not supported. When unzipping the downloaded files/folders, use a zip utility that supports the ZIP64 format.
- If you perform point-in-time recovery of a file that is deleted within the CDP restore window, then the file-level recovery fails with the following error message:

```
Error while obtaining volume information for the file <File_path> -
[kInvalid]: Failed to get the local path for file: <file_path>
```

Recover Files or Folders to the Original Location

To recover files and folders to the original server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select...	Then...
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. > **Note:** When entering the path of a Windows VM, be sure to reformat as shown in the example: Instead of *C:\Windows\System32\winevt* use *C:/Windows/System32/winevt/*. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. 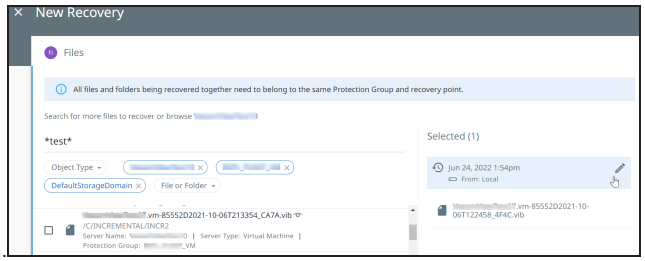 The Edit recovery point page appears. 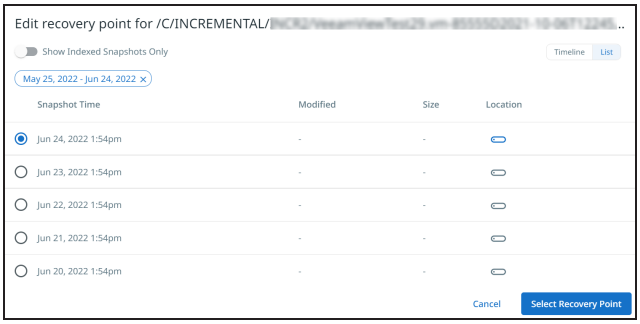 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots**

If you select...	Then...
	Only option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select...	Then...
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select...	Then...
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:

- **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.
- **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

  Character replacement explanation for : and \ in file and directory names in zip files

  If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
  Example:
  - The original name is "a:b" then the archived name is "a_COLON_b"
  - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

  If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
  Example:
  - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
  - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **Original Server** to recover the files or folders to the original server.
6. (Applies only if you are recovering files and folders of a VMware VM). Select one of the following **Restore Method** options and provide the required information:
   - **Auto Deploy Cohesity Agent**: Select this option to automatically deploy the Cohesity Agent on the target VM. This allows you to use local or domain credentials to access the target VM. For Windows VMs, the following considerations are applicable:
     - If you use local user credentials, that user must be a local administrator on the server.
     - If you use domain user credentials, ensure that the username is entered in the format <domain/username>. The user must either be a domain administrator or a member of the Domain Administrators group on the server. Additionally, make sure that User Account Control (UAC) is disabled for the user.
   - **Use Existing Cohesity Agent**: Select this option to use the existing Cohesity Agent deployed on the target VM.

> **Note:** To recover files and folders of VMware VMs using Cohesity Agent, VMware tools must be installed and running on the target VM.

- **Use VMware Tools**: Select this option to use the VMware Tools for recovery. You have to enter the credentials required to access the target VM. If you are using domain credentials to access a target VM, then you should provide the username in the format <domain/username>. Ensure to install the VMware Tools on the target VM.

7. To recover to a different location in the original server, disable the **Recover to Original Path** option, and then provide the location to which the files or folders are to be recovered in the **Recover To** field. By default, the files and folders will be recovered to the original location.

> **Note:** For Hyper-V FLR, you must provide guest OS credentials on the user interface during the recovery workflow.

8. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

     - C:\restore\1.txt
     - C:\restore\cohesity<taskid>_0\1.txt
     - C:\restore\bar
     - C:\restore\cohesity<taskid>_1\bar

     If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
   - **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only, files, then files will receive

the new timestamps.

- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

9. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Files or Folders to a New Location

To recover files and folders to a new server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select...	Then...
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. > **Note:** When entering the path of a Windows VM, be sure to reformat as shown in the example: Instead of *C:\Windows\System32\winevt* use *C:/Windows/System32/winevt/*. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up

If you select...	Then...
	to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select...	Then...
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up the browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select...	Then...
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:

- **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.
- **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

  Character replacement explanation for : and \ in file and directory names in zip files

  If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
  Example:
  - The original name is "a:b" then the archived name is "a_COLON_b"
  - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

  If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
  Example:
  - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
  - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **New Server** to recover the files or folders to a new location.
6. From the **Registered Source** drop-down menu, select a registered server to which you want to recover the files and folders. You can also register a new source by clicking **Register Source** from the **Source** drop-down menu.
7. From the **Target** (for VM and physical server) or **Volume** (for NAS) drop-down menu, select a VM or volume to which the files and files are to be recovered.

   > **Note:** For recovering files and folders of a VMware VM, ensure that VMware Tools are installed and running on the target VM. The target list displayed includes only VMs with VMware Tools installed and running.

8. (Applies only if you are recovering files and folders of a VMware VM). Select one of the following **Restore Method** options and provide the required information:
   - **Auto Deploy Cohesity Agent**: Select this option and provide the username and password of the target VM to automatically deploy the Cohesity Agent on the target VM.

- **Use Existing Cohesity Agent**: Select this option to use the existing Cohesity Agent deployed on the target VM.

  > **Note:** To recover files and folders of VMware VMs using Cohesity Agent, VMware tools must be installed and running on the target VM.

- **Use VMware Tools**: Select this option to use the VMware Tools for recovery. You have to enter the credentials required to access the target VM. If you are using domain credentials to access a target VM, then you should provide the username in the format <username@domain.com>. Ensure to install the VMware Tools on the target VM.

9. In the **Recover To** field, change the default location to which the files or folders are recovered. By default, the files and folders are recovered to the **/tmp/** directory for VMs.

   For NAS, an example path on an alternate location could be /tmp/recover. You can recover encrypted files/folders to their original location only.

   Example for Windows alternate location:

   - D:\cohesity_restore\1.txt
   - D:\cohesity_restore\bar

   These requirements may need to be met to recover files to a Windows VM.

   If recovering to a standalone Hyper-V host or cluster, you must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.

   Files to be recovered: "C\foo\1.txt" and "C:\1stfloor\bar". Specified alternate location: "D:\cohesity_restore". At the end of the copy, the target folder looks like this:

   > **Note:** For Hyper-V FLR, you must provide guest OS credentials on the user interface during the recovery workflow.

10. Optional. Edit the default settings of the following **Recovery options**:

    - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

      Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

      Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

- C:\restore\1.txt
- C:\restore\cohesity<taskid>_0\1.txt
- C:\restore\bar
- C:\restore\cohesity<taskid>_1\bar

If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

- **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only, files, then files will receive the new timestamps.
- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

11. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover VMs

Cohesity provides the ability to recover Protected Objects (such as VMs) from a Snapshot created earlier by a Protection Group. You can recover VMs from a Cohesity cluster or a currently registered archive. You can recover VMs to their original location or a new location.

Follow these procedures to recover VMs from a Cohesity cluster:

- Recover VMs to the Original Location
- Recover VMs to a New Location

> **Tip:** You can also watch the VMware VM Recovery video to learn how to recover a VMware VM to an original or new location.

To learn more about recovering VMs, see About Recovering VMs.

- Only copy recovery from the VMware source to AVS is supported.
- From AVS to VMware source, both copy recovery and instant recovery are supported.
- Recovering a VM from an AVS source to any other heterogeneous source is not supported.
- Only copy recovery is supported for a VM backed up from the AVS environment.

> **Note:** Recoveries between VMware source environments and AVS may fail due to potential incompatibilities between the two platforms.

- Test and dev clone of VM backed up from the AVS environment is not supported.
- VMDK recovery and IVM are not supported for a VM backed up from the AVS environment.

Recover VMs to the Original Location

You recover the data stored in the VMs and the configuration files of the VMs (for example, VMDK files which store the configuration of VMware VMs) to their original datastores and create new instances of the VMs in the original location in the original source.

Before you can recover VMs to its original location, snapshots of those VMs must exist on the Cohesity cluster. To create Snapshots, create and run a Protection Group.

To recover VMs to its original location:

1. Navigate to **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > VMs**.

   The Virtual Machines page appears.
3. Search for a Protected Object (VM) or Protection Group that contains the snapshots to recover from by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more VMs or Protection Groups that contain the snapshots to recover.

> **Note:**
> - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
> - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

5. In the **VM Protection Group Selection** dialog box, review the list of Virtual Machines in the Protection Group and click **Confirm**. This pop-up is displayed for vCenter, Standalone ESXi Host, and vCloud Director.
6. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a...	Then...
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the Protection Run drop-down menu, select the snapshot from which you want to recover the VMs. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a...	Then...
VM	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud (

To choose a different snapshot for a...	Then...
	) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

7. Click **Next: Recover Options**.
8. From the **Recover To** options, select **Original Location** to recover the VM to the original location.
9. (*Applies only to VMware and Acropolis*) Select the **Recovery Type**:

> **Note:**
> - You should have the `VirtualMachine.Config.Rename` permission on your user account to use this recovery option.
> - For VMware VMs, Virtual RDM disks are recovered as flat disks.
> - For more information about AHV recovery options, see Recover a VM with Directly Attached Volume Groups.

- **Instant Recovery**: The VMs will be instantly available after recovery in the target location. The data will be moved to the target storage later. The VM can be used in production during Instant Recovery when data is copied to the VM from the Cohesity cluster. In such scenarios the VM performance may be low as the Storage vMotion is still in progress.

  If you choose to keep the VM in a powered-off state when initiating instant recovery, wait for the Storage vMotion to complete before powering the VM on. Cohesity recommends that you choose Copy Recovery for powered-off VMs for better performance.

  > **Note:** For instant recovery to Standalone ESXi, the VM(s) will be available immediately in the target environment and must moved to the target storage manually.

- **Copy Recovery**: The VMs will be available in the target location only after all the data is copied to the target storage from the source location (Cohesity

cluster or cloud). Once the restore is complete, the VM will be available to use.

> **Note:** Copy Recovery option is the only option to recover VMs from VMware Cloud solutions.

> **Note:**
> - You can instantly recover vCloud Director VMs to their original location if all the vCloud Director VMs of the Protection Group belong to the same 'tag based placement rule' Storage profile.
> - VMs can also be backed up and recovered using the vCD plugin installed on the VMware vCloud Director web console. For more details on the installation and usage of the vCD plugin, see Plugin Installation.

10. **Existing VM Handling** options list (*Applies to VMware, Hyper-V, and AHV*)
    - **VMware**: From the **Existing VM Handling** options list, choose how you want to handle the original VM when recovering to the original location:
      - **None**: Recovers the VM as a new VM and retains the original VM. In the **Rename** field available in the **Recovery Options**, you can specify the name of the recovered VM by adding prefix and/or suffix strings to the name of the original VM. By default, the recovered VM is named as `copy-<original_VM_name>`.
      - **Overwrite Existing VM**: Recovers the VM by deleting the original VM. The recovered VM will have the original VM name. The original virtual machine is deleted prior to the recovery process. Therefore, a recovery failure or cancelation of recovery will also lead to the loss of the original VM.
        Once you select this option, a pop-up dialog box will be displayed. Review the message, type **YES**, and click **Confirm**.
      - **Attempt Differential Recovery**: Attempts to recover the VM by overwriting only the difference between the original VM and the snapshot selected for recovery. Any newly added data in the original VM is deleted.
        - This option is available only if you have selected **Overwrite Existing VM** and the **Recovery Type** option as **Copy Recovery**.
        - This option is disabled if you have selected **Leverage SAN Transport** under Recovery Options.

- You can choose this option if there are no hardware configuration changes involved in the original VM.
- If you want to reclaim free space for thin-provisioned disks, then Cohesity recommends not to attempt differential recovery and only perform **Overwrite Existing VM** recovery.
- Differential recovery reduces the amount of data transfer in a recovery process. In the pulse log, you can view the amount of data transfer saved by selecting differential recovery.
- If the original VM is not present or if the attempt to differential recovery fails, then Cohesity will perform **Overwrite Existing VM** recovery.
- In the original VM, if there are any newly added disks or any disks that are excluded during backup, then the recovered VM will not have these newly added disks and the disks excluded during backup. If you want to retain the newly added disks and the excluded disks in the original VM, contact Cohesity Support. However, Cohesity recommends you to choose the **Keep Existing VM** option if you want to retain the newly added disks and the excluded disks in the original VM while recovering.
- All the snapshots present on the original VM are consolidated and removed as part of differential recovery.
- **Keep Existing VM**: Recover the VM as a new VM with the original VM name. The original VM is retained powered off. The original VM will be renamed to `obsolete_<time_stamp><Orginal_VM_Name>`.

- **Hyper-V**: From the Existing VM Handling options list, choose how you want to handle the original VM when recovering to the original location:
  - **None**: Recovers the VM as a new VM and retains the original VM. In the Rename field available in the Recovery Options, you can specify the name of the recovered VM by adding prefix and/or suffix strings to the name of the original VM. By default, the recovered VM is named as `copy-<original_VM_name>`.
  - **Overwrite Existing VM**: Recovers the VM by deleting the original VM. The recovered VM will have the original VM name. The original virtual machine is deleted prior to the recovery process. Therefore, a recovery failure or cancellation of recovery will also lead to the loss of the original VM. Once you select this option, a pop-up dialog box will be displayed. Review the message, type YES, and click Confirm.

    The VM will initially be located on the Hyper-V host based on the original VM's UUID.

In highly available VMs that migrate to another node in the Hyper-V cluster after a backup, recovery will fail because deleting VMs from a non-original location is currently unsupported.

In the case of Hyper-V standalone clusters or clusters managed by SCVMM, restoring a VM will fail if a VM with the same UUID is present on multiple nodes. A VM can only be successfully restored if its original UUID is found on the original node.

If the VM exists on the host, it will be deleted first. If the VM is not present, the recovery process will proceed directly. Once it is confirmed that the VM is no longer on the host, it will be restored from the backup snapshot to its original location, using the same UUID. The recovery can be either instant or a copy recovery, depending on the options selected by the user.

VM recovery in SCVMM will fail if multiple VMs with the same original UUID are detected under SCVMM-managed hosts or clusters.

> **Note:** This feature is only available if the cluster and agents are on the required version.

When you choose to overwrite an existing VM, the current VM will be deleted before the recovery process begins. If the recovery fails or is canceled, the original VM will be permanently lost. To recover the VM, you must initiate a new restore from a specific point in time, which is determined by the backup snapshot you select.

- **AHV**: From the **Existing VM Handling** options list, choose how you want to handle the original VM when recovering to the original location. This option is available in both **Copy Recovery** and **Instant Recovery** options:
  - **None**: Recovers the VM as a new VM and retains the original VM. In the Rename field available in the Recovery Options, you can specify the name of the recovered VM by adding prefix and/or suffix strings to the name of the original VM. By default, the recovered VM is named as `copy-<original_VM_name>`.
  - **Overwrite Existing VM**:Recovers the VM by deleting the existing VM that has the same UUID. The recovered VM is created with the original UUID and retains the original identity and location. If you select **Overwrite Existing VM**, the system performs the following actions before initiating recovery:
    - Removes the VM's volume group attachments
    - Deletes the existing VM

The existing VM is deleted before the recovery process begins. Therefore, a recovery failure or cancellation can result in the loss of the original VM. When you select this option, a confirmation dialog box is displayed. Review the message, type **YES**, and click **Confirm**.

The recovered VM is instantiated using the original UUID and is placed in the original location on the target cluster.

- **Keep Existing VM**: Recovers the VM as a new VM while retaining the original VM. The system removes the existing VM's volume group attachments and creates a clone of the existing VM. The cloned VM is assigned a new UUID and it retains the name of the original VM, `obsolete_<time_stamp><Orginal_VM_Name>`, after which the volume group attachments are reattached to the cloned VM.

  The recovered VM is created with the original UUID and name of the recovered VM is changed based on the rename option, ensuring both the cloned VM and the recovered VM coexist in the environment.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

11. Optional. Edit the default settings of the following **Recovery Options**:

Recovery Options	Description
Create Empty Disks for Excluded Disks	(Applicable only to AHV and Hyper-V Protection Groups) Enable this option if you want Cohesity to create blank disks for every disk you excluded from the backup. The metadata of the excluded disks will be used to create the blank disks. These blank disks are available post-recovery and only have the metadata of the original disk but no data. You can configure and run databases and applications on these disks, saving you from the effort of creating new disks. This feature is especially helpful if you have many VMs with excluded disks.

Recovery Options	Description
Exclude Disk Handling	(Applicable only to VMware) • **Create Excluded Disk as an Empty Disk**: Enable this option if you want Cohesity to create blank disks for every disk you excluded from the backup. The metadata of the excluded disks will be used to create the blank disks. These blank disks are available post-recovery and only contain the metadata of the original disk, but no data. You can configure and run databases and applications on these disks, saving you from the effort of creating new disks. This feature is especially helpful if you have many VMs with excluded disks. > **Note:** The Create Excluded Disk as an Empty Disk option is not supported with Differential Copy Recovery. • **Create Independent Disk as an Empty Disk**: By default, independent disks are not backed up. Enable this option if you want Cohesity to create blank disks for such independent disks. These blank disks are available post-recovery and only contain the metadata of the original disk, but no data. > **Note:** Excluded physical RDM disks cannot be created as blank disks, even when disk exclusion handling is enabled.

Recovery Options	Description
Network	By default, the VMs that are to be recovered do not have a virtual Network Interface Card (vNIC) attached. Enable the **Attach** option to attach a virtual Network Interface Card (vNIC) to each VM that is to be recovered. With the **Attach** option enabled, each VM will connect to its original network and the following option is displayed: • **Start Connected**– Enable this option to connect to the original network when the VMs reboots. If this option is disabled, the VMs are not connected to any network on reboot. While recovering VMs to the original location, there is no explicit control to preserve the MAC address. If the **Attach** option is enabled and the VM is configured to have a manual MAC address assigned at the time of backup, the same address will be recovered. **Note:** When recovering AHV VMs to the original location, the recovered VM will have the same number of NICs as the source VM. If the VM is configured to have an automatic MAC address at the time of backup, the same setting will be applied to the recovered VM. For information on preserving MAC address when recovering VMs to a new location, see Recover VMs to a New Location.
NFS Protocol	(*Applies for VM recovery*) NFS Protocol version (NFS v3 or NFS v4) Select the NFS protocol version to expose Cohesity storage as an NFS share on the primary source environment.
Rename	(*Applies for VM recovery*) Add Prefix and/or Suffix strings to the names of the new VMs created by this task. For example you could specify Prefix `Test_ and Suffix _QA` to the VM name WinOS and when the VM is recovered, the new VM that is created is called `Test_WinOS_QA`.
Rename vApp/vApp Template	(*Applies for vApp or vApp template recovery*) Add prefix and/or suffix strings to the names of the new vApp or vApp template created by this task. For example, you could specify prefix `Test_ and Suffix _QA` to the vApp template name WinOS and when the vApp template is recovered, the new vApp template that is created is called `Test_WinOS_QA`.

Recovery Options	Description
Rename VMs within vApp	(*Applies vApp or vApp template recovery*) Add prefix and/or suffix strings to the names of the new VMs within vApp created by this task. For example, you could specify prefix `Test_` and Suffix `_QA` to the VM name WinOS and when the VM within the vApp is recovered, the new VM that is created is called `Test_WinOS_QA`.
Leverage SAN Transport for Copy Recovery	**Leverage SAN Transport** SAN transport will work only with the Copy Recovery option. You can turn on the Leverage SAN Transport toggle when you want to perform the copy recovery of VMware objects to vCenter target environments. You can now perform copy recovery of vCloud Director (vCD) objects, such as VMs and vApps, to vCD target environments, You can also leverage the SAN transport mode while you perform the copy recovery of the vCD objects. Once SAN Transport is enabled for copy recovery, the disk provisioning strategy selected during the Cohesity cluster setup will be reset to “Thick Provisioned Eager Zeroed” in adherence to the recommendation from VMware. For more details, see the KB article. Cohesity supports SAN transport for Fibre Channel (FC) and iSCSI-based storage arrays and datastores that contain such storage arrays. Turn on this option to enable the SAN transport mode over FC or iSCSI for backing up VMware VMs to Cohesity cluster. To use the SAN transport mode over FC, ensure the following pre-requirement: • Fibre channel is installed and zoned appropriately between the Cohesity cluster and the primary storage (Pure Flash Array and HPE Nimble Storage Array). To use the SAN transport mode over FC and iSCSI , ensure the following pre-requirement: • Volumes are mapped appropriately between the Cohesity cluster, ESXi hosts, and the primary storage.
	Allow NB DSSL Transport Fallback You can turn on the Allow NBDSSL Transport Fallback toggle if you do not want to enforce strict SAN transport mode for backing up VMware to the Cohesity cluster. With the toggle enabled, backup doesn't fail if the SAN transport is not available. If you disable the toggle, the backup fails if SAN transport mode is unavailable. NBDSSL encrypts all data passed over the TCP connection.
Power State	Disable **Power On** if you want the recovered VMs to remain powered off after they are created. If you recover VMs from a Storage Domain that has Cloud Tier enabled, Cohesity recommends leaving the VMs powered off. By default, the VMs are powered on after they are created.

Recovery Options	Description
Continue on Error	Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cohesity cluster will still attempt to create the other VMs. If this option is disabled, then the recovery task is discontinued if VM creation fails. However, if the VM creation succeeds, but VM storage migration fails, the task continues. By default this option is disabled.
Instant Recovery	(*Applies only to Hyper-V*) The **Instant Recovery** option allows you to access the recovered VMs almost instantly while the data is migrated to the VM's primary storage in the background. This is the default recovery type in the recovery workflow. **Note:** • You must turn off the toggle (from the **Recovery Options** list) to disable the Instant Recovery if you want to perform a Copy Recovery of the VM. • For Instant Recovery, the VM returns to its normal state once the storage migration is complete, as it no longer relies on the SMB view for storage. Prior to the storage migration, in SCVMM scenarios, the VM could be reported as being in an unsupported state.
Preserve VMs UUID	Select **Preserve VMs UUID** if you want to preserve the UUID for the recovered VMs. This is applicable only to Microsoft Hyper-V VMs. **Note:** If the UUID of the VM and the restored VM are same, a new UUID will be used by the restored VM.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down. **Note:** The Cluster Interface option is not supported for VMware VM Copy Recovery.
Task Name	Change the default name of the recovery task.

12. Click **Recover**. Review the Recovery Summary and click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover VMs to a New Location

You recover the data stored in the VMs and the configuration files of the VMs (for example, VMDK files which store the configuration of VMware VMs) to their original data stores and create new instances of the VMs in the original location in the original source. Before you can recover VMs to their original location, snapshots of those VMs must exist on the Cohesity cluster.

Before you can recover VMs to their new location, Snapshots of those volumes must exist on the Cohesity cluster. To create Snapshots, create and run a Protection Group.

To recover VMs to a new location:

1. Navigate to **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > VMs**.

   The Virtual Machines page appears.
3. Search for Protected Object (VM) or Protection Group that contains the snapshots to recover from by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more VMs or Protection Groups that contain the snapshots to recover.

   > **Note:**
   > - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
   > - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.
5. In the **VM Protection Group Selection** dialog box, review the list of Virtual Machines in the Protection Group and click **Confirm**. This pop-up is displayed for

vCenter, Standalone ESXi Host, and vCloud Director.

6. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a...	Then...
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the Protection Run drop-down menu, select the snapshot from which you want to recover the VMs. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a...	Then...
VM	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived in the cloud by clicking the cloud (

To choose a different snapshot for a...	Then...
	) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

7. Click **Next: Recover Options**.
8. From the **Recover To** options, select **New Location** to recover the VM files (such as the VMDK files) to an alternate datastore and create new instances of the VMs in a new location of a registered Source. For more information, see Recover to a New Location. For details about virtualization or cloud platform-specific options, see the respective section.
    - Acropolis
        a. **Source:** Select a registered Source to recover the VMs to.
        b. **Storage Container:** Select a storage container to store the recovered VMs.
    - AWS
        a. **Source:** Select a registered Source to recover the VMs to.
        b. **Region:** Select a region in the project to recover the VMs.
        c. **VPC:** Select an Amazon Virtual Private Cloud (VPC) to store the recovered VMs.
        d. **Subnet:** Select a subnet in the Amazon VPC to store the recovered VMs.
        e. **Key Pair:** Select an Amazon key pair that has permission to log in to the recovered VMs.
        f. **Network Security Group:** Select the associated network security group.
    - Azure
        a. **Source:** Select a registered Source to recover the VMs to.
        b. **Resource Group, Storage Account:** Select the corresponding Resource Group and Storage Account to which the recovery has to be done.
        c. **Storage Container:** Select the storage container in which page blobs will be created.
        d. **Compute:** Select the VM size type such as **Standard_D1** to use when recovering the VMs.

e. **Availability Set:** Optionally, select the availability set to which the VM has to be recovered. For more information, see Azure Availability Sets.
f. **Virtual Network:** Select the virtual network where the VM must be placed.
g. **Subnets:** Subnets in the virtual network where VMs must be placed.

- GCP
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Project:** Select a Google Cloud Platform (GCP) project to recover the VMs to.
  c. **Region:** Select a region in the project to recover the VMs to.
  d. **Availability Zone:** Select an availability zone in GCP to recover the VMs.
  e. **VPC:** Select a Virtual Private Cloud (VPC) to store the recovered VMs.
- Hyper-V
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Compute:** Select the compute host for the recovered VMs.
  c. **Volume:** Select a volume to store the recovered VMs.

  When recovering to a standalone Hyper-V host as a new location, SMB volumes are not supported locations.
- RHV
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Data Center:** Select a data center to store the recovered VMs.
  c. **Cluster:** Select an RHV/OLVM to store the recovered VMs.
  d. **Storage Domain:** Select an RHV/OLVM storage domain to store the recovered VMs.
- VMware
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Resource Pool:** Select a Resource Pool to store the recovered VMs.

> **Note:** When protecting a VMware Source, if you select a Resource Pool that is not part of the ESX cluster, all datastores in the vCenter are displayed. Ensure you select a datastore that is part of the selected Resource Pool. If you select the wrong datastore, storage migration fails, but the VM remains attached to the ESX Host and runs from the Cohesity datastore. You can manually migrate the VM to the desired datastore at a later time.

c. **Datastore:** Select one or more available target datastores. The VM files (such as VMDK files) extracted from the Snapshots are stored in the selected datastores. The number of target datastores you select must be equal to or less than the number of VMs selected for recovery.

   Cohesity also displays each datastore's amount of free space and total capacity. For more information about recovering VMs to multiple datastores, see Recover VMs to Multiple Datastores.

   > **Note:** Not applicable when recovering to a standalone ESX host. After recovery, VMs continue to run on the Cohesity datastore because Storage vMotion is not supported.

d. **VM Folder** Optionally select a logical VM folder in the data center in which to create the VMs.

   Cohesity recommends that you select the "VM" folder to recover to the root folder or select a folder from the displayed list of folders. If not, the Cohesity cluster recovers to a newly created folder named as follows:

   cohesity_vm_folder_<short_cluster_id>_<task_id>

   > **Note:** Not applicable when recovering to a standalone ESX host.

- vCloud Director

  a. **Source:** Select a registered vCloud Director source to which you plan to recover the VMs or vApp.

  b. **Organization VDC:** Select the target organization VDC.

  c. **vApp:** Optional. Select the target vApp. If not specified, VMs are recovered as standalone VMs inside the target organization VDC.

  d. **Storage Profile:** Select the target Storage Profile. The Cohesity cluster chooses any compliant datastore associated with the Storage Profile to store the VM files. If you are recovering multiple VMs, the Cohesity cluster chooses one or more datastores associated with the Storage Profile.

  e. **Catalog**: (*Appplies only for vApp template recovery*) Select the catalog for the vApp template recovery.

9. (*Applies only to VMware and Acropolis*) Select the **Recovery Type**:

   > **Note:** For more information about AHV recovery options, see Recover a VM with Directly Attached Volume Groups.

- **Instant Recovery**: The VMs will be instantly available after recovery in the target location. The data will be moved to the target storage later. The VM can be used in production during Instant Recovery when data is copied to the VM from the Cohesity cluster. In such scenarios the VM performance may be low as the Storage vMotion is still in progress.

  If you choose to keep the VM in a powered-off state when initiating instant recovery, wait for the Storage vMotion to complete before powering the VM on. Cohesity recommends that you choose Copy Recovery for powered-off VMs for better performance.

  > **Note:** For instant recovery to Standalone ESXi, the VM(s) will be available immediately in the target environment and must be moved to the target storage manually.

- **Copy Recovery**: The VMs will be available in the target location only after all the data is copied to the target storage from the source location (Cohesity cluster or cloud). Once the restore is complete, the VM will be available to use.

  > **Note:**
  > - You can instantly recover vCloud Director VMs to their original location if all the vCloud Director VMs of the Protection Group belong to the same 'tag based placement rule' Storage profile.
  > - VMs can also be backed up and recovered using the vCD plugin installed on the VMware vCloud Director web console. For more details on the installation and usage of the vCD plugin, see Plugin Installation.

  > **Note:** Copy recovery is the only option to recover VMs from VMware Cloud solutions.

10. (*Applies only for vApp template recovery*) From the **Recover As** options, select how you want to recover the vApp template:
    - **vApp Template**: Select this option if you want to recover the vApp template as a vApp template itself.
    - **vApp**: Select this option if you want to recover the vApp template as a vApp.
11. Optional. Edit the default settings of the following **Recovery Options**:

Recovery Options	Description
Create Empty Disks for Excluded Disks	(Applicable only to AHV and Hyper-V Protection Groups) Enable this option if you want Cohesity to create blank disks for every disk you excluded from the backup. The metadata of the excluded disks will be used to create the blank disks. These blank disks are available post-recovery and only have the metadata of the original disk but no data. You can configure and run databases and applications on these disks, saving you from the effort of creating new disks. This feature is especially helpful if you have many VMs with excluded disks.
Exclude Disk Handling	(Applicable only to VMware) • **Create Excluded Disk as an Empty Disk**: Enable this option if you want Cohesity to create blank disks for every disk you excluded from the backup. The metadata of the excluded disks will be used to create the blank disks. These blank disks are available post-recovery and only contain the metadata of the original disk, but no data. You can configure and run databases and applications on these disks, saving you from the effort of creating new disks. This feature is especially helpful if you have many VMs with excluded disks. > **Note:** The Create Excluded Disk as an Empty Disk option is not supported with Differential Copy Recovery. • **Create Independent Disk as an Empty Disk**: By default, independent disks are not backed up. Enable this option if you want Cohesity to create blank disks for such independent disks. These blank disks are available post-recovery and only contain the metadata of the original disk, but no data. > **Note:** Excluded physical RDM disks cannot be created as blank disks, even when disk exclusion handling is enabled.

Recovery Options	Description

Network	By default, the VMs that are to be recovered do not have a virtual Network Interface Card (vNIC) attached. Enable the **Attach** option to attach a virtual Network Interface Card (vNIC) to each VM that is to be recovered.
	
	With the **Attach** option enabled, the following option is displayed:
	
	• **Network**: From the drop-down menu, select a network to attach the virtual Network Interface Card (vNIC) to a new network. The new network might be a standard port group or a distributed port group.
	  • When recovering VMware VMs, the selected Network port group will be attached to all the network interfaces of the VMs being recovered.
	  • When recovering the VMware VMs to vCloud Director, the Network drop-down menu displays the organization VDC networks attached to the Organization VDC on the vCloud Director.
	  • When recovering Hyper-V VMs, the selected virtual switch will be attached to all the network interfaces of the VMs being recovered.
	  • Recovering a network attached CentOS VM on VMware or Hyper-V, does not attach the VM to the network. Attach the VM to the network manually.
	  • If the VMware Resource Pool is not part of the ESX cluster, all networks available to the Datacenter are displayed. Ensure you select a network available to the selected Resource Pool. If you select the wrong network, VM creation and recovery will fail.
	  > **Note:** When recovering AHV VMs to the original location, the recovered VM will have the same number of NICs as the source VM.
	
	• **Start Connected**: Enable this option to connect to the new network when the VM reboots for each recovered VM. If this option is not selected, the VMs are not connected to any network on reboot.
	• **Preserve MAC Address**: Enable this option to preserve the MAC address when recovering to an alternate location.
	  > **Note:** If the VM is configured to have an automatic MAC address at the time of backup, a manual MAC address will be set on the recovered VM to preserve the MAC address.

Recovery Options	Description
	• **Network Override**: The option allows you to override the common network settings selected for the recovery job for VMware VMs. You can choose the VM for which the virtual Network Interface Card (vNIC) of the selected snapshot will be listed. For each vNIC, you can select the target network along with the "Start Connected" and "Preserve Mac Address" options. You can select the VM for which the virtual Network Interface Card (vNIC) of the chosen snapshot will be displayed.
NFS Protocol	(*Applies for VM recovery*) NFS Protocol version (NFS v3 or NFS v4) Select the NFS protocol version to expose Cohesity storage as an NFS share on the primary source environment.
Rename	(*Applies for VM recovery*) Add prefix and/or suffix strings to the names of the new VMs created by this task. For example, you could specify Prefix `Test_ and Suffix _QA` to the VM name WinOS and when the VM is recovered, the new VM that is created is called `Test_WinOS_QA`.
Rename vApp/vApp Template	(*Applies only when recovering vApp or vApp template*) Add prefix and/or suffix strings to the names of the new vApp or vApp template created by this task. For example, you could specify prefix `Test_ and Suffix _QA` to the vApp template name WinOS and when the vApp template is recovered, the new vApp template that is created is called `Test_WinOS_QA`.
Rename VMs within vApp	(*Applies only when recovering vApp or vApp template*) Add prefix and/or suffix strings to the names of the new VMs within vApp created by this task. For example, you could specify prefix `Test_ and Suffix _QA` to the VM name WinOS and when the VM within the vApp is recovered, the new VM that is created is called `Test_WinOS_QA`.
Power State	Disable **Power On** if you want the recovered VMs to remain powered off after they are created. If you recover VMs from a Storage Domain that has Cloud Tier enabled, Cohesity recommends leaving the VMs powered off. By default, the VMs are powered on after they are created.

Recovery Options	Description
Continue on Error	Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cohesity cluster will still attempt to create the other VMs. If this option is disabled, then the recovery task is discontinued if VM creation fails. However, if the VM creation succeeds, but VM storage migration fails, the task continues. By default this option is disabled.
Instant Recovery	(*Applies only to Hyper-V*) The **Instant Recovery** option allows you to access the recovered VMs almost instantly while the data is migrated to the VM's primary storage in the background. This is the default recovery type in the recovery workflow. **Note:** • You must turn off the toggle (from the **Recovery Options** list) to disable the Instant Recovery if you want to perform a Copy Recovery of the VM. • For Instant Recovery, the VM returns to its normal state once the storage migration is complete, as it no longer relies on the SMB view for storage. Prior to the storage migration, in SCVMM scenarios, the VM could be reported as being in an unsupported state.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down. **Note:** The Cluster Interface option is not supported for VMware VM Copy Recovery.
Task Name	Change the default name of the recovery task.

12. Click **Recover**. Review the Recovery Summary and click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Teardown of Failed VM Recovery Tasks

Cohesity supports the teardown of failed VM recovery tasks, allowing you to clean up residual resources left behind when these tasks do not complete successfully.

To teardown a failed VM recovery task that has left behind certain resources:

1. Navigate to **Data Protection > Recoveries**.
2. Locate the instant VM recovery task you wish to terminate.

   To teardown a recovery task, Click the **Actions** menu (⋮) of that task and select **Teardown**.

   > **Note:** The Teardown option is available only if there are objects to destroy.

3. When prompted to confirm the teardown, click **Teardown**.

   A confirmation dialog appears with a warning stating that the restored VMs with data on the Cohesity mounted NFS datastore will be removed.

4. Click **Confirm**.

   The teardown process is initiated, and the Cohesity cluster performs the following steps:

   1. The VM running from the Cohesity datastore is removed.
   2. The NFS datastores are unmounted from the ESXi host.
   3. Cohesity cluster resources are released.
   4. The user interface for Cohesity will show the current status of the teardown task.
   5. The teardown status of each VM is displayed under the "Instant Recovery" column.
   6. "Destroyed" status indicates successful completion of the teardown process.

   > **Note:**
   > - The VMs that have been manually migrated to the target datastore will not be considered for teardown.
   > - The teardown task can be retried in case of any issues.

Recover Nutanix AHV

Cohesity supports recovering of the directly attached Volume Groups for Nutanix AHV VMs to the original and new locations.

Recover a VM with Directly Attached Volume Groups

Recover a VM with Directly Attached Volume Groups

During recovery, you can choose whether to restore the Volume Groups that were attached to the VM at the time of protection.

Select the protection group that includes attached Volume Groups, and in the **Recovery Options** section, set **Restore Directly Attached Volume Groups** to **Yes**, then click **Recover**.

When you enable the **Restore Directly Attached Volume Groups** option:

- Volume Groups attached to the VM at backup time are recreated.
- Newly created Volume Groups are attached to the recovered VM.
- Volume Group restore is supported only during copy recovery and not supported for instant recovery.

The **Activity** page displays the UUIDs created for the restored volume groups. It also lists the directly attached volume groups that were successfully restored.

FileLevel Recovery for Volume Group Disks

FileLevel Recovery (FLR) is supported for disks belonging to directly attached Volume Groups if:

- The disks were mounted in the guest OS at backup time (for example, listed in /etc/fstab on Linux).
- The mounted disks are supported by the Cohesity agent.

Considerations

- If multiple VMs share the same Volume Group and are restored together, each restored VM receives its own independent copy of the Volume Group; a single shared Volume Group is not recreated across multiple VMs.
- Volume Group categories are not restored during recovery.
- Instant recovery is not supported for Volume Group restores.

Recover Virtual Disks

You can recover VMware VM virtual disks to a point in time from a Snapshot created earlier by a Protection Group. You can recover a virtual disk to its original VM or a different VM in the same vCenter. If recovering to the original VM, you can overwrite the disk in its original datastore, or attach it as a new disk. If recovering to a different VM, the disk is attached as a new disk in the datastore you select. For more information, see:

- Recover Virtual Disks to the Original Location
- Recover Virtual Disks to a New Location

Prerequisite

VMware licensing is required for all VMware products to facilitate processes such as backup or recovery with Cohesity. For more information on the VMware license, see the following VMware articles:

- VMware vSphere Storage APIs
- Storage vMotion Requirements and Limitations

Process Overview

The recover task completes the following steps:

1. A new physical instance (VMDK) is extracted from the Snapshot and stored temporarily in the View (datastore) on the Cohesity cluster.
2. The new VMDK is migrated from the View to the specified datastore in the vCenter using Cohesity's implementation of vSphere Storage vMotion.
3. The temporary View (datastore) is deleted automatically once the Storage vMotion migration is completed.

After a successful recovery, the disks are attached to the VM. Depending on your environment, if the virtual disks were mapped to volumes, additional steps may be required to bring them online.

Note the following about recovering virtual disks:

- Support is currently limited to recovering VMware virtual disks.
- If the **Overwrite Original Disk** option is selected, the original disk is removed from the VM and deleted, and the backed-up Cohesity disk is attached. If the recovery fails during the overwrite, the original disk will not be available, but you can retry the recovery to restore the disks.
- If a virtual disk gets deleted after it is backed up, the disk must be recovered as a new disk.
- If one or more disks will be overwritten during recovery, the target VM will automatically be powered off prior to the recovery and can optionally be powered on automatically after the recovery. Cohesity recommends suspending activity on the VM prior to performing the disk recovery.
- If a Windows VM includes volumes created from a storage pool (Microsoft Storage spaces), VMDK recovery is not supported.

Recover Virtual Disks to the Original Location

If you are recovering the virtual disk to the original VM, you can overwrite the disk in its original datastore or attach it as a new disk.

To recover virtual disks to the original server:

1. Navigate to **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > Virtual Disks**.

   The Virtual Disks page appears.
3. Search for the VM whose virtual disks you want to recover by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, server type, or date.
4. From the list of search results, select the VM whose virtual disks you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

     b. Select the snapshot to be recovered by clicking the corresponding blue dot.

        For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar.

     c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     d. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the **List** mode:

a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page.

   By default, the page displays the snapshots available in the last one month.

b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

c. Select the snapshots to be recovered.

d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the virtual disks to the original location.
8. Provide the following information as required for **Disk Selection**:

Field	Description	Action to performed
Source Disks	List of virtual disks to recover	From the list, select the virtual disks to recover.
Mapped Volumes(s)	Indicates whether the virtual disks are mapped to the volumes.	None
Size	Indicates the size of the virtual disks.	None

Field	Description	Action to performed
Recovery Type	Option to recover the virtual disks.	Select one of the following options. • **Overwrite Original Disk**–Specify this option to recover the virtual disks to their original datastores using the same disk ID and controller information. The original disk is overwritten. The original virtual disk is deleted prior to the recovery process. Therefore, a recovery failure or cancelation of recovery will also lead to the loss of the original virtual disk. • **Recover as a New Disk**–Specify this option to recover the virtual disk as a new disk (without overwriting the original disk).
Datastore	The datastore to store the virtual disks extracted from the Snapshots.	Select a datastore to store the virtual disks extracted from the snapshots.

9. Optional. Edit the default settings of the following **Recovery options**:
   - **Power State**: Enable the **Power off VMs before restore** option to ensure that the VM is powered off before the virtual disks are recovered. Enable the **Power on VMs after restore** option to power on the VM after the virtual disks are recovered.
   - **Task Name**: Change the default name of the recovery task.
10. Click **Start Recovery**.

The recovery process is initiated. You can monitor the recovery task from the Recoveries page.

Recover Virtual Disks to a New Location

If recovering the virtual disk to a different VM, the disk is attached as a new disk in the datastore you select.

To recover virtual disks to a new server:

1. Navigate to **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > Virtual Disks**.

   The Virtual Disks page appears.

3. Search for the VM whose virtual disks you want to recover by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, server type, or date.

4. From the list of search results, select the VM whose virtual disks you want to recover.

   By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

     b. Select the snapshot to be recovered by clicking the corresponding blue dot.

        For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar.

     c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     d. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page.

        By default, the page displays the snapshots available in the last one month.

b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

c. Select the snapshots to be recovered.

d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** to recover the virtual disks to an alternate location.
8. From the **Source** drop-down menu, select a registered source having the VM on to which you want to recover the virtual disks. You can also register a new source by clicking **Register Source** from the **Source** drop-down menu.
9. From the **Target** drop-down menu, select a VM on which you want to recover the disk.
10. Provide the following information as required for **Disk Selection**:

Field	Description	Action to performed
Source Disks	List of virtual disks to recover	From the list, select the virtual disks to recover.
Mapped Volumes(s)	Indicates whether the virtual disks are mapped to the volumes.	None
Size	Indicates the size of the virtual disks.	None
Recovery Type	Option to recover the virtual disks.	Select **Recover as a New Disk** from the drop-down list and to recover the virtual disk as a new disk.
Datastore	The datastore to store the virtual disks extracted from the Snapshots.	Select a datastore to store the virtual disks extracted from the snapshots.

11. Optional. Edit the default settings of the following **Recovery options**:
    - **Power State**: Enable the **Power off VMs before restore** option to ensure that the VM is powered off before the virtual disks are recovered. Enable the **Power on VMs after restore** option to power on the VM after the virtual disks

are recovered.

- **Task Name**: Change the default name of the recovery task.

12. Click **Recover**.

The recovery process is initiated. You can monitor the recovery task from the Recoveries page.

Teardown of Failed VMDK Recovery Tasks

Cohesity supports the teardown of failed VMDK recovery tasks, allowing you to clean up residual resources left behind when these tasks do not complete successfully.

> **Note:** Tearing down a successful recovery is not allowed.

To teardown a failed VMDK recovery task that has left behind certain resources:

1. Navigate to **Data Protection > Recoveries**.
2. Locate the VMDK Recovery task you wish to terminate.

   To teardown a recovery task, Click the **Actions** menu (⋮) of that task and select **Teardown**.

   > **Note:** The Teardown option is available only if there are objects to destroy.

3. When prompted to confirm the teardown, click **Teardown**.

   A confirmation dialog appears with a warning stating that the restored VMs with data on the Cohesity mounted NFS datastore will be removed.

4. Click **Confirm**.

   The teardown process is initiated, and the Cohesity cluster performs the following steps:

   1. VMDKs hosted on Cohesity datastore are detached from the target VM.
   2. The Cohesity datastore is unmounted from the ESXi host.
   3. The underlying restore view is deleted.
   4. The user interface for Cohesity will show the current status of the teardown task.
   5. The teardown status of each VMDK is displayed under the "Instant Recovery" column.
   6. "Destroyed" status indicates successful completion of the teardown process.

> **Note:**
> - VMDKs that have been manually migrated to the target datastore will not be considered for teardown.
> - The teardown task can be retried in case of any issues.

Instant Volume Mount

Instant volume mounting is supported for virtual and physical environments. This feature allows the selected backup volumes to be made available at the target location where you can then complete the desired operations. Instant Volume Mount is available for backed-up volumes, whether they are stored locally on the Cohesity cluster or archived to a registered external target. Use cases include granular recovery of Microsoft Exchange, SQL and SharePoint data.

You can mount the volume to its original location or to a new location. For more information, see:

- Instant Volume Mount to the Original Location
- Instant Volume Mount to a New Location

The following table lists the sources and the guest operating systems on which instant volume mounting is supported.

Source	Support for IVM on guest OS	
	Windows	Linux
AHV	No	No
VMware	Yes	Yes
HyperV	Yes	No
Physical*	Yes	Yes

*Supported only for block-based backups.

Prerequisites

- VMware licensing is required for all VMware products to facilitate processes such as backup or recovery with Cohesity. For more information on the VMware license, see

the following VMware articles:

- VMware vSphere Compute Virtualization
- Storage vMotion Requirements and Limitations

- The default path of the recovered file or instantly mounted volume is *<Cluster_IP_address>/<Cloned view name>\<PATH>*.

  To access this location, both the target and the Cohesity cluster must be in the same AD domain.

Considerations

- When recovering a file or instantly mounting a volume from a Windows VM or Physical Server Backup Source that has Windows deduplication installed and enabled for one or more volumes, you must choose a target machine that also has Windows deduplication installed (it does not have to be enabled for any volume). (However, this rule does not apply to Nutanix AHV VMs. If AHV VMs are enabled with Windows deduplication, the only supported recovery option is full VM recovery.)

  If the target does not have Windows deduplication installed:

  - File level recovery will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."
  - Instant volume mounting will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."

  To determine if Windows deduplication is installed on the Source or target machine, follow the steps given below:

  1. Open **Server Manager**.
  2. Select **Roles and Features > File and Storage Services > File and iSCSI Services**.
  3. Select the **Data Deduplication** check box, if necessary.

4. Click **Next** until the **Install** button is enabled and then click **Install**.

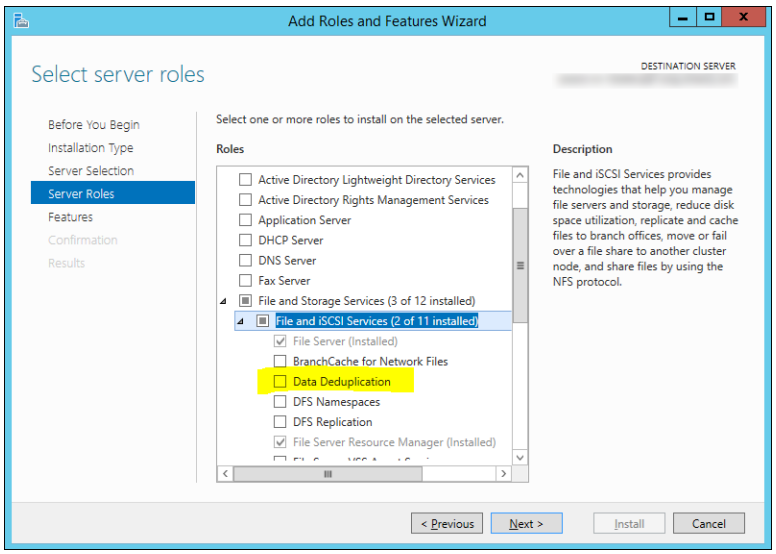

- Tearing down a cloned database or instant volume mount deletes the mounted volumes. Any new or modified data on these volumes will be deleted along with the volumes, so ensure you back up any important data before teardown.
- Review the following considerations when performing instant volume mount to Hyper-V VMs:
  - Only Windows VMs are supported.
  - Dynamic Disks (LDM and LVM) are not supported.
  - The Bring Disks Online option requires the following:
    - VM must be part of Active Directory, the VM and the Hyper-V host must be in the same AD.
    - Users must execute "winrm quickconfig" to enable winrm on the target VM and remote powershell must be enabled from Hyper-V host to VM.
  - Instant volume mount and file level recovery from Gen 1 to Gen 2 type VMs is not supported.
  - If SCVMM is unregistered from the Cohesity cluster, ensure you tear down all instant volume mounts. Not tearing them down can prevent the VM from being backed up when the source is registered with a different Cohesity cluster.
  - Instant volume mounting Hyper-V 2012 R2 VMs without a SCSI controller is not supported. This is because Hyper-V disallows dynamically adding a SCSI controller, which is required to add the virtual disks.
  - On 2012 R2 VMs, if an instant volume mount disk is attached to a VM during a Protection Group run, that snapshot cannot be application-consistent. If this occurs, the event viewer may contain a VSS-catastrophic error or similar message.

- Instant Volume Mount for NetApp stub file is not supported.
- You cannot instantly mount a volume from a VM to a physical server, and vice-versa.
- Before you can instantly mount a volume, a Snapshot of the Object you want to recover must exist.
- Cohesity does not currently provide support for protecting the vSphere Clustering Service (vCLS) using temporary NFS datastores. For optimal results, we recommend configuring a VMware vSphere® Distributed Resource Scheduler (DRS) rule to ensure that vCLS is provisioned only on the datastores that you have added.
- If a Windows VM includes volumes created from a storage pool (Microsoft Storage spaces), Instant Volume Mount is not supported.

Troubleshooting

This section provides troubleshooting information for issues that might occur during IVM recovery.

Authentication Failures

Error Message

The IVM recovery fails due to an authentication failure and displays the following error:

"The specified network password is not correct".

Possible Root Cause

The cluster joined the AD but the Windows host has not joined the AD.

Solution

- Create a local user on the Cohesity cluster and use the local user to access the SMB share. To create a local user on the Cohesity cluster, contact Cohesity Support.

or

- Disable SMB Authentication on the Cohesity cluster. Perform the following steps to disable SMB Authentication:
  1. Log in to the Cohesity cluster.
  2. Select **SmartFiles** > **Views** and click the **Global Settings** tab.
  3. Toggle off the **SMB Authentication** setting.

Instant Volume Mount to the Original Location

To perform instant volume mounting to the original server:

1. Navigate to **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > Instant Volume Mount**.

   The Instant Volume Mount page appears.

3. Search for the server that contains the data you want to mount by entering characters of the server name or by specifying the wildcard character *.

   A list of servers displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, server type, or date.

4. From the list of search results, select the server that contains the data you want to mount.

   By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

     b. Select the snapshot to be recovered by clicking the corresponding blue dot.

        For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar.

     c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     d. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page.

By default, the page displays the snapshots available in the last one month.

b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

c. Select the snapshots to be recovered.

d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Select Volume(s)** option, perform one of the following actions:
   - Enable **Select all Volumes** to recover all the volumes of the selected server
   - From the **Select one or more options** drop-down list, select the individual volumes you want to recover.

   After volumes are mounted, they may be displayed as volumes different from their source volume. For example, D: volume may be mounted as X: volume.
8. From the **Recover To** options, select **Original Location** to recover the volume to the original location.
9. (Applies only for VMs) With the **Ensure Disks are Online** option enabled, select one of the following options to deploy Cohesity Agent on the target VM.

   > **Note:** For Linux target VMs, you cannot disable this option. For Window target VMs, you can disable this option.

   - **Auto Deploy Cohesity Agent:** Select this option and provide the username and password of the target VM to automatically deploy the Cohesity Agent on the target VM.
   - **Use Existing Cohesity Agent**: Select this option to use the existing Cohesity Agent deployed on the target VM

     The following table provides information on how the volume is mounted when the **Ensure Disks are Online** option is enabled and disabled.

Environment	Ensure Disks are Online option is Turned On	Ensure Disks are Online option is Turned Off
VMware	After the VMDKs are attached to the VM, Cohesity cluster explicitly mounts the volume. To mount the volume to the VM, Cohesity cluster requires the built-in administrator credentials, such as the local system administrator of the VM. It uses the credentials provided in the **Mount Target Credentials** fields for mounting the volume. If the UAC property "Admin Approval mode" is disabled, then you can specify any user who is part of the system's Local Administrator group. • **Note:** VMware Tools are required to bring the disks online, but mounting is still supported without VMware Tools.	The VMDKs are just attached to the VM and the VM may bring the disks online. Cohesity cluster does not explicitly bring the volumes online within the VM. In this case, the Target Volume displays as "n/a" after mounting the volume succeeds.
Standalone Hyper-V host or cluster	You must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.	Not Applicable

10. Optional. Edit the default settings of the following **Recovery Options**:
    - **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
    - **Task Name**: Change the default name of the recover task.
11. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

> **Note:** If a drive letter was provided as the file path for a file-based backup of NTFS volumes, the default attributes for drive letters were preserved, so hidden drives remain hidden. They are still recovered and can be accessed by providing the full path in the search bar.

Instant Volume Mount to a New Location

To perform instant volume mounting to a location:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > Instant Volume Mount**.

   The Instant Volume Mount page appears.
3. Search for the server that contains the data you want to mount by entering characters of the server name or by specifying the wildcard character *.

   A list of servers displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, server type, or date.
4. From the list of search results, select the server that contains the data you want to mount.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

     b. Select the snapshot to be recovered by clicking the corresponding blue dot.

        For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar.

     c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud (

) icon if available in the **Location** option.

d. Click **Select Recovery Point** and then go to step 6.

- To select a snapshot from the **List** mode:

a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page.

By default, the page displays the snapshots available in the last one month.

b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

c. Select the snapshots to be recovered.

d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud (

) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Select Volume(s)** option, perform one of the following actions:
   - Enable **Select all Volumes** to recover all the volumes of the selected server
   - From the **Select one or more options** drop-down list, select the individual volumes you want to recover.

   After volumes are mounted, they may be displayed as volumes different from their source volume. For example, D: volume may be mounted as X: volume.
8. From the **Recover To** options, select **New Location** to recover the volume to a new location.
9. From the **Source** drop-down menu, select a registered source having the server to which you want to recover the volume. You can also register a new source by clicking **Register Source** from the **Registered Source** drop-down menu.
10. From the **Target** drop-down menu, select a physical server or VM to which the volume will be mounted.

> **Note:** To mount to a physical server, the server must already have the Cohesity agent installed and registered with the Cohesity cluster. If not, the server will not be available as a mount target.

11. (Applies only for VMs) With the **Ensure Disks are Online** option enabled, select one of the following options to deploy Cohesity Agent on the target VM.

> **Note:** For Linux target VMs, you cannot disable this option. For Window target VMs, you can disable this option.

- **Auto Deploy Cohesity Agent:** Select this option and provide the username and password of the target VM to automatically deploy the Cohesity Agent on the target VM.
- **Use Existing Cohesity Agent**: Select this option to use the existing Cohesity Agent deployed on the target VM

The following table provides information on how the volume is mounted when the **Ensure Disks are Online** option is enabled and disabled.

Environment	Ensure Disks are Online option is Turned On	Ensure Disks are Online option is Turned Off
VMware	After the VMDKs are attached to the VM, Cohesity cluster explicitly mounts the volume. To mount the volume to the VM, Cohesity cluster requires the built-in administrator credentials, such as the local system administrator of the VM. It uses the credentials provided in the **Mount Target Credentials** fields for mounting the volume. If the UAC property "Admin Approval mode" is disabled, then you can specify any user who is part of the system's Local Administrator group. • **Note:** VMware Tools are required to bring the disks online, but mounting is still supported without VMware Tools.	The VMDKs are just attached to the VM and the VM may bring the disks online. Cohesity cluster does not explicitly bring the volumes online within the VM. In this case, the Target Volume displays as "n/a" after mounting the volume succeeds.
Standalone Hyper-V host or cluster	You must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.	Not Applicable

12. Optional. Edit the default settings of the following **Recovery Options**:

- **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
- **Task Name**: Change the default name of the recover task.

13. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

> **Note:** If a drive letter was provided as the file path for a file-based backup of NTFS volumes, the default attributes for drive letters were preserved, so hidden drives remain hidden. They are still recovered and can be accessed by providing the full path in the search bar.

Teardown a Mounted Volume

Teardown a mounted volume when you are finished working with it.

To teardown a mounted volume:

1. Navigate to **Data Protection > Recoveries**.
2. To teardown a recovery task, click the Actions menu (⋮) of that task and select **Teardown**.

> **Note:** The Teardown option is available only if there are objects to destroy.

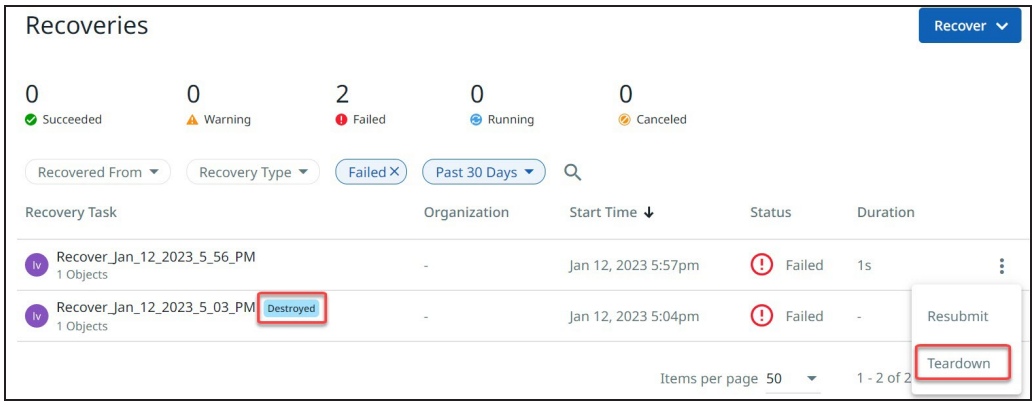

3. When prompted to confirm the teardown, click **Teardown**.

The teardown process is initiated, and the Cohesity cluster performs the following steps:

1. The virtual disk files (such as VMDK or VHD) are detached from the server.
2. The virtual disk files stored on the view are deleted.
3. The view is unmounted from the physical server or hypervisor host and then removed from the Cohesity cluster.

A "Destroyed" tag for a task on the Recoveries page indicates the completion of the teardown process.

Download a VMX File of a Virtual Machine

Cohesity supports the ability to download the VMX file of a Virtual Machine from the latest Snapshot of a Protection Group. A VMX file is the primary configuration file for a VMware Virtual Machine. Advanced VMware administrators may want to download and view the settings in a VMX file. In addition, they may want to download, modify and then upload a VMX file back to the original datastore to implement workarounds or global changes.

> **Warning:** Downloading a VMX file and uploading it directly back into your datastore without modifications can result in unusable and inaccessible VMs. This feature is for advanced VMware administrators who are experienced in directly editing VMX files and manually configuring VMs. Use at your own risk.

Before you download the VMX file of a Virtual Machine from the latest Snapshot of a Protection Group, you must create and run a Protection Group.

To download a VMX file:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover > Virtual Machines > VMs**.
3. Search for the Virtual Machine with the associated VMX file that you want to download. Enter characters of the VM name and a list of items displays. You can also optionally narrow the search results by specifying filter criteria, for example you can filter the search results by a specific Protection Group. You can optionally specify the wildcard character * in the search.
4. In the list, find the VM with the associated VMX file that you want to download, click the Actions Menu (⋮) and select **Download VMX file for MyVM**, where **MyVM** is the name of your VM.

If the Cohesity cluster locates the VMX file in the last Snapshot of the Protection Group, the VMX file is downloaded. (In some cases the VMX file may be unreachable when the Snapshot was originally created and therefore not saved and cannot be downloaded later. For example the Cohesity cluster cannot save a VMX file if the HTTPS port to the ESXi host was not open when the Snapshot was created.)

In addition, if the last Run of the Protection Group failed to back up the specified VM, the VMX file for that VM cannot be downloaded.

If you require downloading a VMX file from Snapshot besides the last Snapshot, contact Cohesity Support.

> **Warning:** Downloading a VMX file and uploading it directly back into your datastore without modifications can result in unusable and inaccessible VMs. This feature is for advanced VMware administrators who are experienced in directly editing VMX files and manually configuring VMs. Use at your own risk.

Azure VMware Solution (AVS)

Cohesity makes it simple and affordable to back up, recover, and manage data in Azure VMware Solution (AVS). AVS combines VMware Software Defined Data Center (SDDC) with Microsoft Azure Cloud bare-metal infrastructure.

See the following topics for details:

1. Plan and Prepare for AVS Protection
2. Register your AVS (SDDC)
3. Protect your AVS source
4. Recover VMs, files and folders

Plan and Prepare for AVS Protection

Before registering your AVS source, ensure it meets user privileges and firewall port requirements.

User Privileges for VMware

To protect your AVS solution software-defined data center (SDDC), you need a user account with the vCenter Server Privileges.

Firewall Ports

Ensure the firewall ports are open to allow communication between the Cohesity clusters and the AVS environment.

Considerations

File-level recovery from AVS VMs that reside on volumes with Windows Data Deduplication enabled is not supported.

Register AVS Source

To start protecting your vCenter in AVS, register it as a source on the Cohesity cluster.

Before you Begin

Ensure that the requirements for the AVS source, such as user account role privileges and firewall ports, are met.

To begin with, register your AVS source with the Cohesity cluster.

To register AVS Solution with the Cohesity cluster:

1. Navigate to **Data Protection** > **Sources**.
2. Select **Register** > **Virtual Machines**.
3. Select a **Hypervisor Source Type** from the drop-down.
4. Enter the information for your source type.

   Click **Register**.

   The AVS source has been successfully registered and is now visible on the Sources page. For information on managing your AVS source, see Manage AVS Source.

Manage AVS source

View the Registered AVS Source

To view a registered AVS Source:

1. Navigate to **Data Protection** > **Source**.

   The Sources page displays all the sources registered with that Cohesity cluster. Click the Actions menu (⋮) of an AVS source to Edit, Refresh, or Unregister.

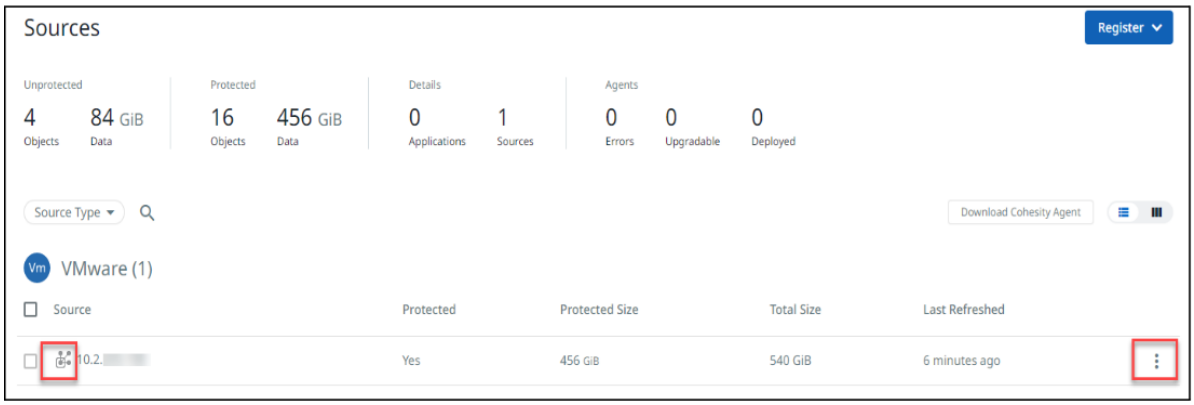

2. Click on a source to view the hierarchy of the data centers, clusters, ESXi hosts, folders, and VMs in that AVS source.

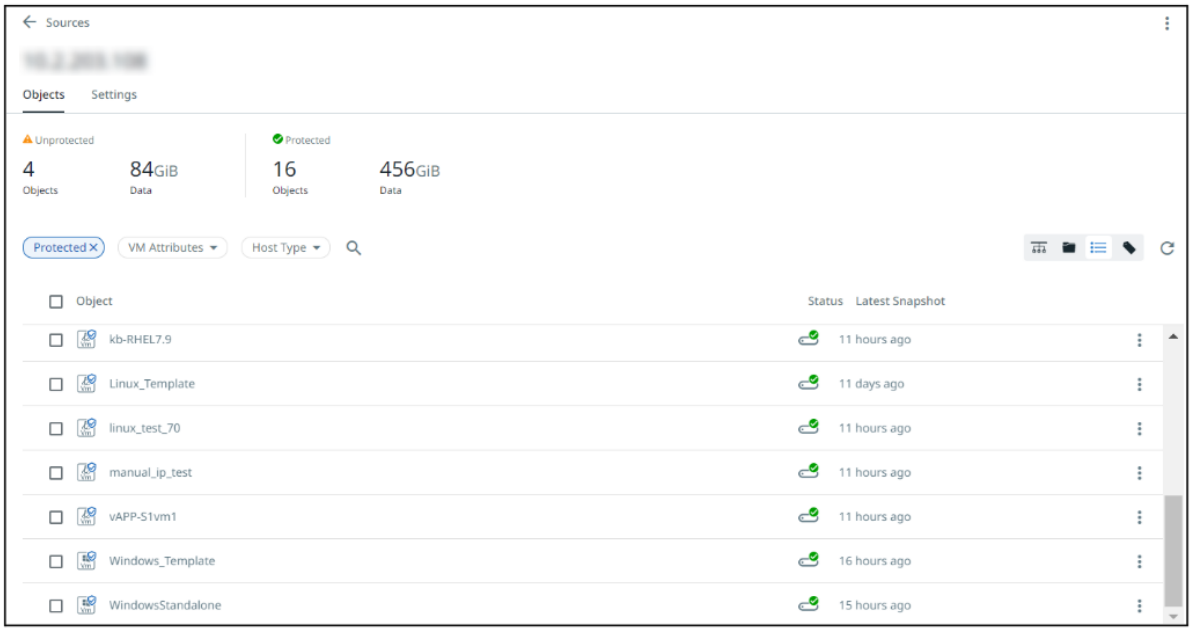

a. A selected check box ( ) indicates the object is selected to be backed up.

b. - Indicates a Linux VM.

c. - Indicates a Windows VM.

d. - Indicates a Windows or Linux template.

e. - Indicates a VM.

f. - A blue shield indicates a protected VM.

g. - Indicates an excluded object that is not protected. An ancestor of the object is auto-protected, but this object is excluded through inheritance (an ancestor is explicitly excluded).

h. - Click to view the general information—type, server size, and version of the selected object.

3. Filters on the **Objects** tab help you display only the objects you want. The filtering options are:
   - **Protection Status**. Protected or Unprotected.
   - **VM Attributes**. VMware Tools Found, MS SQL Servers, vCenter Migrated.

Cohesity Agent Installed or Cohesity Agent Not Installed.

Host Type. Linux or Windows.

4. Enter the name of an object in the **Search** box. You can even enter a partial name or an asterisk * wildcard. As you type, the objects that contain your search term appear.

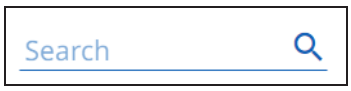

Display modes can show the objects in a hierarchy or a flat list. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select Data Protection > Sources, find the source in the list, and click the icon.

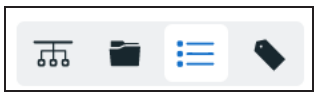

Update the AVS Source Configuration

You can update the AVS configuration you provided during the registration process with the latest configuration.

Refresh the AVS Source

You can refresh the AVS source configuration to fetch the latest changes. To refresh the AVS source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AVS source and select **Refresh**.

Unregister AVS Source

To stop backing up your AVS source, unregister it from the Cohesity cluster. To unregister the AVS source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the AVS source and select **Unregister**.

Protect AVS Source

After registering your AVS source, you can use Cohesity to protect the VMs, files and folders in that source.

Backup AVS Source

Create a Protection Job for AVS objects. For more information, see Add or Edit a Protection Group for Virtual Servers.

Recover AVS Source

After you protect your AVS source, you can recover VMs and files from your backups to their original or new location.

Recover AVS VMs to a New or Original Location

Recover AVS Files or Folders to a New or Original Location

Recover AVS VMs to a New or Original Location

You can recover the AVS VMs to the original (same location as the source) or a new location. After selecting the restore point, you can modify the storage domain and network. While recovering, you can either auto-deploy a Cohesity Agent, use an existing agent, or use VMware tools.

Recover AVS Files or Folders to a New or Original Location

Cohesity provides the ability to recover AVS files and folders from a snapshot created earlier by a Protection Group. Files and folders can be recovered to their original location (limitations apply to physical servers and specific operating systems) or to a newly specified location, which can be within the original source or a different one. You can choose to retain the recovered files' and folders' original (at the time of the backup) permissions and attributes. You can also download files and folders from selected snapshots that a Cohesity Protection Group created. For more information on recovering AVS files or folders to a new or original location, see Recover Files or Folders.

VMware vCloud Director (vCD)

VMware vCloud Director (vCD) allows Cloud Service providers to manage multiple tenants with data centers in different geographic locations and data stores that require different performance levels. Cloud Service Providers can use vCD to create organizations that are virtual administration units. Each organization can streamline services to multiple tenants.

Cohesity offers integration with vCD in two ways:

- Protect vCD resources from the Cohesity cluster
- Protect vCD resources from the vCD web console

Considerations

- Not all vCD storage policy types are supported. During instant recovery, Cohesity uses a tag-based storage policy for the temporary Cohesity datastore. After the storage vMotion completes, the VM is reconfigured to the target storage policy. During copy recovery, you can select a target storage policy at recovery time.

Protect vCD resources from the Cohesity cluster

The Cohesity Data Cloud enables you to configure multi-tenant environments that can be securely isolated from each other. A single Cohesity cluster can service backups and workloads of multiple tenants with secure namespace isolation and per-tenant encryption.

For more information, see Protect vCD Resources from the Cohesity Cluster.

Protect vCD resources from the vCD web console

By installing the Cohesity vCD plugin, you can integrate Cohesity protection services with the vCD web console. When vCD is combined with Cohesity Data Cloud, you can provide backup and recovery for your vCD tenants.

Cohesity offers granular data protection by enabling you to protect the entire vCloud Director or just the organizations, VMs, and vApps within it. In addition to allowing cloud service providers or administrators to manage the organizations, Cohesity Data Cloud enables tenants to leverage data protection as self-service, lowering operating costs and strengthening tenant efficacy.

For more information, see Protect vCD resources using the Cohesity vCD Plugin.

Elastic Org VDC

Seamless use of Elastic Org VDCs with multiple resource pools is now supported for backup and recovery operations.

Key Concepts for vCD Protection

vCD

vCD allows Cloud Service providers to support multiple tenants with data centers across different geographic locations or data stores requiring different performance levels.

vCD Web Console

The vCD Web Console allows Service Providers to create organizations that are virtual administration units.

Organization

An organization is a collection of users, groups, and compute resources. System administrators create organizations while the organization administrations manage the users and groups in an organization. An organization is mapped to a vSphere data pool and can contain multiple data stores.

vApp

vApp is a deployment unit containing one or more VMs and their network configuration.

Protect vCD Resources from the Cohesity Cluster

Protecting the vCD resources from the Cohesity cluster involves the following steps:

1. Before protecting the vCD resources, ensure that you meet the requirements.
2. Setup vCD Protection for Tenants.
3. Backup the vCD objects
4. Recover the vCD sources
    - Recover VMs and vApps
    - Recover Files and Folders

Plan and Prepare for vCD Protection

Prerequisites

Cohesity cluster supports data protection and recovery of supported vCD resources. For supported versions of vCD, see the Support Matrix and Supported Software sections.

Minimum Permissions

Ensure that the vCD user account you use to manage and protect the vCD resources has adequate permissions. For details, see Minimum Permissions.

Considerations

- If a VMware vCloud Director (vCD) source sub-object is assigned to a tenant, the recovery of VMs and vApps to an alternate location will fail in the 6.2 release. When an entire vCD is registered within a tenant, then recovery to both the original location and alternate location is supported.
- If multiple VMs with different storage profiles are restored to the original location, the default storage policy of the organization VDC will be used for restoration purposes.
- Updating a vCD source to add a new vCenter instance requires entering credentials for all new and existing vCenter instances.

Backup and Recover Metadata for VMware vCD Objects

Cohesity supports backup and recovery of the metadata for VMware vCD objects such as vApps, VMs, and vApp templates. Metadata includes user-defined properties of the vCD objects whose values can be text, date, time, boolean, number, strings, badges, etc. You can recover metadata to the original location or a new location. The read-only metadata is not recovered while recovering vApps, VMs, and vApp templates. Cohesity also supports backup and recovery of metadata for Test & Dev clones.

Setup vCD Protection for Tenants

As a service provider, perform the following steps to set up data protection of vCD sources for your tenants:

1. **Register your vCD source**

   Register your vCD as a source with the Cohesity cluster. For details on how to do this, see Register and Manage a vCD Source.

2. **Enable Multi-Tenancy**

   Enable Organizations (Multitenancy) on Cluster by enabling the Organizations feature.

3. **Create and Map Organizations**

   To allow a vCD Organization user to protect their Organization's resources, create Cohesity Organizations.

   From the Cohesity cluster, you can also perform the following tasks to configure the organization:

   a. Manage Access for Organizations
   b. Assign Users and Groups for Organizations
   c. Assign objects
   d. Impersonate an Organization

Register and Manage a vCD Source

Register vCD as a source

Once registered, its objects can be protected or backed up on the Cohesity cluster.

To register a source to back up:

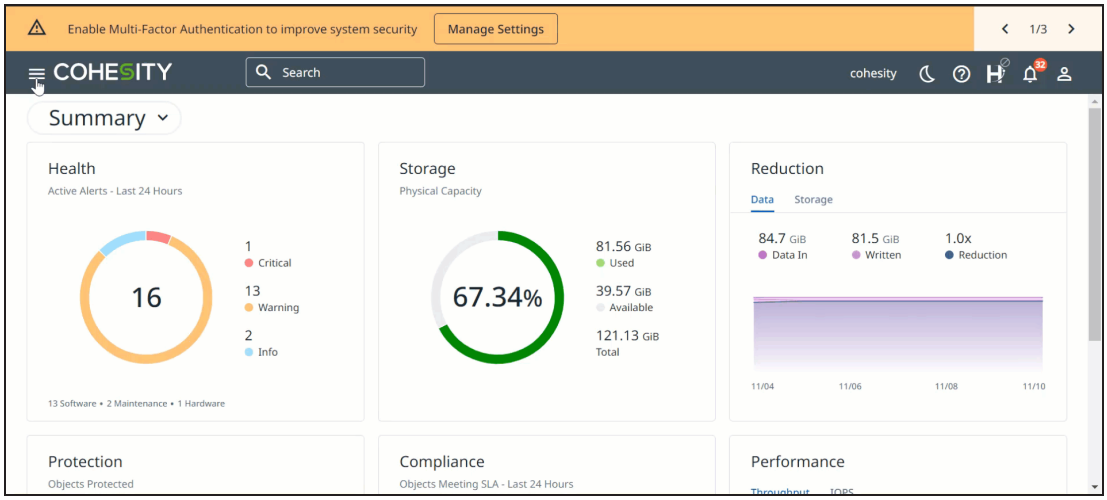

1. On the **Dashboard** page in the Cohesity cluster, navigate to **Data Protection** > **Sources**.
2. Select **Register** > **Virtual Machines**.
3. From Source Type, select **VMware: vCloud Director**.

4. Enter the following vCD details:
   - **Hostname or IP Address**: Specify the hostname or IP address of the vCloud Director that contains the vCenters to protect.
   - **Password**: Enter the password for the specified user.
   - **Username**: Specify the local administrator user.
5. Click **Connect**.
6. Enter the following vCenter details:
   - **Username**: You must specify a VMware vCenter user with adequate privileges on the VMware vCenter server. The role assigned to the specified VMware vCenter user should include privileges to turn on Change Block Tracking (CBT) and other privileges required by the Cohesity cluster. For detailed information about the necessary privileges, see Create a Role with Necessary vCenter Server Privileges.

     When specifying a Windows Active Directory username, use the following format: *<domain>\<username>* or *<username>@domain*. For all other cases, just enter: *<username>*.
   - **Password**: Enter the password for the vCenter.
   - **Detect VM migration across vCenter to Preserve backup chain**: Toggle this option on if you want your Cohesity cluster to detect the VMs that are migrated from the current VMware vCenter to vCD.
7. Click **Register**.

After source registration finishes successfully, the new source is displayed on the **Sources** page.

View the Registered vCD Source

To view a registered vCD Source:

1. Navigate to **Data Protection** > **Source**.

   The **Sources** page displays all the sources registered with that Cohesity cluster. The vCD ( ) icon indicates that the registered source is a vCD Source.

   Click the Actions menu (⋮) of that source to Edit, Refresh, or Unregister. From the Actions menu (⋮), you can also Deploy a Cohesity Agent and register Oracle, MS SQL, and Physical Servers.
2. Click on a source to view the hierarchy of the data centers, clusters, ESXi hosts, folders, VMs, vApps, and templates in that vCD source.
   - A selected check box ( ) indicates the object is selected to be backed up.
   - - Indicates a Linux VM.

- - Indicates a Windows VM.
- - Indicates a Windows or Linux template.
- - Indicates a vApp.
- - Indicates a VM
- - A blue shield indicates a protected VM.
- - Indicates an excluded object that is not protected. An ancestor of the object is auto-protected, but this object is excluded through inheritance (an ancestor is explicitly excluded).
- - Click to view the general information—type, server size, and version of the selected object.

3. Filters on the Objects tab help you display only the objects you want. The filtering options are:
   - **Protection Status**: Protected or Unprotected.
   - **VM Attributes**: VMware Tools Found, MS SQL Servers, vCenter Migrated, Cohesity Agent Installed or Cohesity Agent Not Installed.
   - **Host Type**: Linux or Windows.

4. Enter the name of an object in the **Search** box. You can even enter a partial name or an asterisk * wildcard. As you type, the objects that contain your search term appear.

   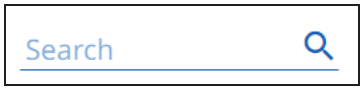

5. Display modes can show the objects in a hierarchy or a flat list. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection** > **Sources**, find the source in the list, and click the Refresh ( ) icon.

Update the vCD Source Configuration

You can update the vCD configuration you provided during the registration process with the latest configuration. To edit the vCD source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Actions** menu (⋮) next to the vCD source and select **Edit**.
3. In the **Edit Source** page, update the respective configurations.
4. Click **Save**.

Refresh the vCD Source

You can refresh the vCD source configuration to fetch the latest changes. To refresh the vCD source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Actions** menu (⋮) next to the vCD source and select **Refresh**.

Partial Refresh of the vCD Souce

You can perform a partial refresh of the vCD source using the API. A partial refresh is faster than a full source refresh. This enhancement allows you to selectively discover new VMs, vApps, Organizations, Org VDC, and vApp templates through the APIs into the vCD source, eliminating the need to wait for a full vCD refresh. You can refresh only the selected objects, such as specific VMs, vApps, Organizations, Org VDC, and vApp templates within the vCD source.

To do a partial refresh of the vCD source using the API:

Endpoint

```
PUT /v2/data-protect/sources/{sourceId}/objects
```

Parameters

Parameter	Type	Required	Description
sourceId	integer	Yes	The Entity ID of the data protection source

Finding the sourceId

You can find the sourceId from the UI by navigating to the Sources page and clicking your desired vCD source.

The URL will appear in the following format:

```
https://<Cluster IP>/protections/sources/details/{sourceId}/objects?environment=kVMware
```

The `{sourceId}` is the numeric ID you need for the API call. This value remains constant until the source is unregistered.

Headers

Header	Type	Required	Description
Content-Type	string	Yes	Must be `application/json`
Authorization	string	Yes	Bearer token for authentication

Request Body

The request body should include a JSON object with the following structure:

```
JSON
{
 "modificationRequests": [
 {
  "actionType": "Add",
  "vcdModifyParams": {
   "urn": "string"
 }
 }
 ]
 }
```

Request Body Schema

Field	Type	Required	Description
modificationRequests	array	Yes	Array of modification requests to process.
modificationRequests[].actionType	object	Yes	Action to perform. Valid values: `Add`.
modificationRequests []. vcdModifyParams	string	Yes	vCD-specific parameters for the modification.
modificationRequests []. vcdModifyParams.urn	string	Yes	The URN (Uniform Resource Name) of the VCD object to be discovered

vCD URN Format

VCD URNs follow the format:

```
`urn:vcloud:{type}:{uuid}`
```

Navigate to **Data Protection** > **Sources**.

The common VCD object types:

- `vapp` - Virtual Application
- `vm` - Virtual Machine
- `vapptemplate` - vApp Template
- `org` - Organization
- vdc` - Virtual Datacenter

To Perform a Partial Refresh:

1. **Authenticate**: Obtain Access Token using Authentication API:

   ```
   POST /irisservices/api/v1/public/accessTokens
   ```

2. Get Source ID of the vCD source as described above.
3. Discovery request API .

   ```
   PUT /v2/data-protect/sources/{sourceId}/objects
   ```

Example Requests

1. Discover a new VM.

   ```
   curl --location --request PUT
   'https://10.136.5.86/v2/data-protect/sources/1/objects' \
   --header 'Content-Type: application/json' \
   --header 'Authorization: Bearer YOUR_TOKEN_HERE' \
   --data '{
    "modificationRequests": [
    {
    "actionType": "Add",
    "vcdModifyParams": {
    "urn": "urn:vcloud:vm:96228d49-a079-4b0e-a19f-c5e78acd65b1"
    }
    }
    ]
    }'
   ```

2. Discover a new VCD org.

   ```
   JSON
   {
    curl --location --request PUT
    'https://10.136.5.86/v2/data-protect/sources/1/objects' \
    --header 'Content-Type: application/json' \
    --header 'Authorization: Bearer YOUR_TOKEN_HERE' \
    --data '{
    "modificationRequests": [
    {
   ```

```
"actionType": "Add",
"vcdModifyParams": {
urn": "urn:vcloud:org:96228d48-a079-4b0e-a19f-f5e78acd65b1"
}
}
]
}
```

Response Status Code and Body

1. Success Response (207 Multi-Status).

```
JSON
{
 "results": [
 {
 "succeeded": true,
 "vcdResult": {
 "urn": "urn:vcloud:org:96228d49-a079-4b0e-a19f-c5e78acd65b1"
 }
 }
 ]
 }
```

2. Error Responses.
   a. 403 Forbidden

```
JSON
{
 "errorCode": "KInvalidRequest",
 "message": "Invalid inputs. Unknown input urn type"
 }
```

   a. 500 Internal Server Error

```
JSON
{
 "errorCode": "KInternalError",
 "message": "VM not found in VCD"
 }
```

Response Status Codes

Status Code	Description
207	Multi-Status - Request processed, check individual results,

Status Code	Description
403	Forbidden - Invalid request or unknown URN type.
500	Internal Server Error - vCD object not found or server error.

> **Note:**
> - You can include multiple objects in the **modificationRequests** array to process several modification requests in a single API call. The batch operations are not supported for vCD partial refresh.
> - **URN Validation**: The API verifies that the specified URN exists in the connected vCD environment before processing the request.
> - **Hierarchical Dependencies**: When you add a vApp, OrgVDC, or Organization, all descendant entities are automatically included.
> - **vCD Hierarchy**: The system adheres to the vCD organizational hierarchy (Organization → VDC → vApp/vAppTemplate → VM) when processing modifications.
> - Ensure that your authorization token has the required permissions to modify data protection sources and access vCD objects.

Related Endpoints

Authentication API

```
POST /irisservices/api/v1/public/accessTokens` - Obtain access token for authentication
```

Authentication API Documentation: Obtain Access Token

Endpoint: `POST /irisservices/api/v1/public/accessTokens`

Description: Authenticate with the system and obtain a bearer token for API access.

Headers

| Header | Type | Required | Description |

| Content-Type | string | Yes | Must be `application/json` |

Request Body

Field	Type	Required	Description
username	string	Yes	Username for authentication.
password	string	Yes	Password for authentication.
domain	string	No	Authentication domain (e.g., "LOCAL") .

Example Request

```
curl --location --request POST
 'https://10.136.5.86/irisservices/api/v1/public/accessTokens' \
 --header 'Content-Type: application/json' \
 --data '{
 "username": "admin",
 "password": "admin",
 "domain": "LOCAL"
 }'
```

Success Response (201 Created)

```
JSON
{
 "accessToken": "eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9...",
 "tokenType": "Bearer",
 "privileges": [
 "AD_LDAP_MODIFY"
 "AD_LDAP_VIEW",
 "AGENT_UPGRADE_MODIFY",
 "PROTECTION_MODIFY",
 "PROTECTION_VIEW",
 "..."
 ]
 }
```

Response Fields

Field	Type	Description
accessToken	string	JWT token to use in Authorization header.
tokenType	string	Type of token (Bearer).

Field	Type	Description
privileges	array	List of permissions granted to the authenticated user.

Usage

Use the returned `accessToken` in the Authorization header for subsequent API calls:

Authorization: Bearer <accessToken>.

This API allows you to perform a partial refresh to discover VMware vCloud Director (VCD) objects within a data protection source.

You can selectively refresh and add VCD entities such as vApps, VMs, VCD Orgs, and other VCD resources for backup and protection.

Unregister vCD Source

If you plan to stop taking the backup of your vCD source, you can unregister the vCD source from the Cohesity cluster. To unregister the vCD source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Actions** menu (⋮) next to the vCD source and select **Unregister**.

Access Management: Privileges for Service Providers and Tenants

A Service Provider can perform the following:

- Register a vCD protection source
- Enable multi-tenancy in the Cohesity cluster
- Create tenants
- Assign a vCD Organization to the tenant
- Create and assign Storage Domains to tenants
- Create and assign Protection Policy to tenants

A tenant can perform the following:

- Log in to the Cohesity cluster
- Change password
- Refresh the registered vCD protection source
- Create protection job for vCD sources
- Monitor the protection runs
- Recover vCD sources

For details on the privileges of the service provider admin and the organization administrator, see Privileges for Service Provider Admin and Organization Admin.

Create a Protection Group for vCD objects: Protect Organizations, vApps, and VMs

To add or edit a Protection Group:

1. If you are creating a Protection Group, click **Protect > Virtual Machines**.

   The New Protection page appears.

   When protecting vApp and vApp template, you must enable auto-protect. If you are not enabling auto-protect, then Cohesity cluster will only back up the underlying objects of the vApp and vApp template. When protecting vApp VMs, Cohesity recommends using one protection group exclusively for all VMs within a vApp and not mixing vApp VMs with other VMs.

2. Navigate to **Data Protection** > **Protection** and do one of the following:

   - To create a Protection Group—Click **Protect** >**Virtual Machines**.

     The **New Protection** page appears.

     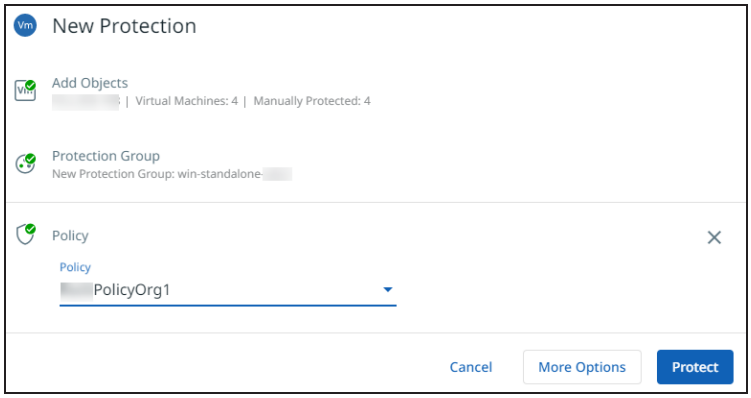

   - To edit a Protection Group—Click the Actions menu (⋮) next to the displayed list of Protection Groups and select **Edit**.

3. Click **Add Objects**, and from the **Registered Source** drop-down list, select the Virtual server whose VMs you want to protect. You can also Register a new source with the Cohesity cluster by clicking the **Register Source** option.

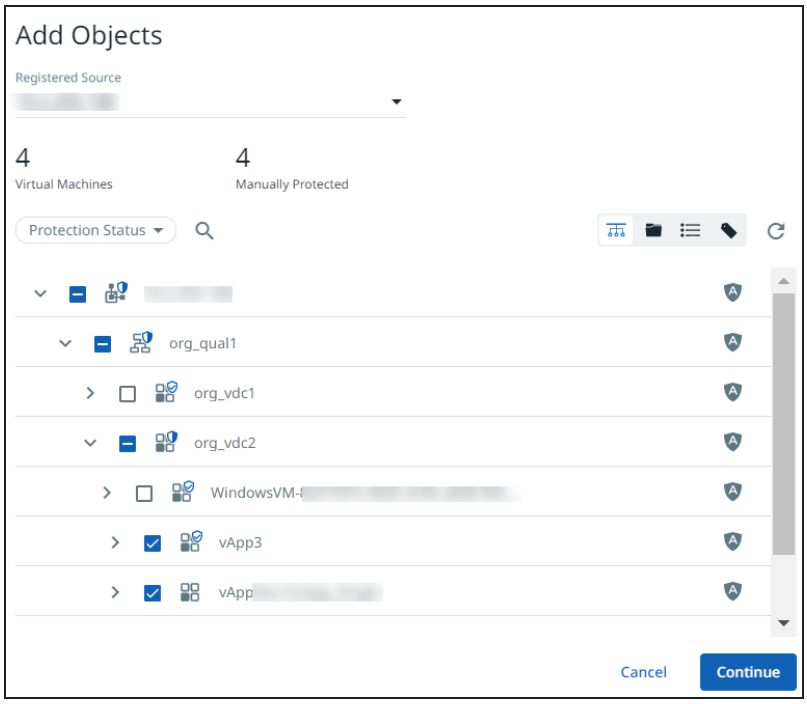

> **Note:**
>
> - VMs that are on ESXi hosts which are not part of a cluster in vCenter are not supported.
> - Cohesity can exclude the backup of placeholder VMs in a VMware VMs Protection Group. This is an Early Access feature and you must contact Cohesity Support to enable this feature.

4. Click **Add Objects**, and from the **Registered Source** drop-down list, select the vCD source you want to protect. The hierarchy of the organizations, virtual data centers, VMs, vApps, and vApp templates in that vCD source are displayed. For more details, see View the Registered vCD Source. You can also Register a new source with the Cohesity cluster by clicking the **Register Source** option.
5. Click **Continue** to go back to the New Protection page.
6. Select **Protection Group** and in the **Name** field, enter a name from the Protection Group.
7. Select **Policy** and from the **Policy** drop-down list, select an existing protection policy or create a new one by selecting **New Policy**. For more information, see Create or Edit a Standard Policy.
8. Select **Storage Domain** and select an existing storage domain or create a new one by selecting **Add Storage Domain**.

9. Select one of the following options:
    - **Protect**: Select this option if you want to complete the Protection Group creation without changing or configuring any of the additional settings.
    - **More Options**: Select this option and go to step 12 if you want to change or configure any of the additional settings.
10. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
11. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
12. **Pause Future Runs**: Enable this option to pause future protection runs for the Protection Group.
13. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
14. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.

- **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.

15. **Defer Incomplete Objects in Concurrent Runs**: (Applies only to VMware VMs) A Protection Group can protect both small (for example, 1 GB VM) and large VMs (for example, 1 TB VM). While a short VM usually completes its protection runs within a few minutes, a large VM might take longer and still be running when it's time for the next run. By default, when it's time to start the next run, the Cohesity cluster checks for any in-progress protection run in that Protection Group. If there are in-progress protection runs (typically a large VM), the Cohesity cluster cancels the next run for that Protection Group, causing the small VMs to also miss that run and the SLA. When the **Defer Incomplete Objects in Concurrent Runs** option is enabled, if there is an in-progress protection run, the Cohesity cluster postpones the subsequent run only for that VM. The postponed run resumes as soon as the current run is completed. The protection runs of the remaining VMs (typically small VMs) in that Protection Group are not affected and run on schedule.

    For example, Protection Group PG1 runs every 8 hours to protect VMs A, B, and C. When the first protection run starts, A and B complete running in a few minutes. VM C, a large VM, needs 12 hours to complete running. If **Defer Incomplete Objects in Concurrent Runs** is disabled for PG1, when it's time for the second run, the Cohesity cluster cancels the second run in all three VMs because C is still in progress. However, the third run is completed in all three VMs. Because C is a large VM, A and B miss every alternate run.

    If **Defer Incomplete Objects in Concurrent Runs** is enabled for PG1, when it's time for the second run, the Cohesity cluster starts the second run in A and B but postpones the run in C. C's second run resumes when its first run completes. As a result, A, B, and C do not miss any scheduled runs.

If PG1 also protects VM D, another large VM, which needs 20 hours to complete, when it's time for the second run, the Cohesity cluster postpones D's second run. When it's time for the third run, because D's first run is still in progress, the Cohesity cluster cancels its second run but postpones its third run. D's next run will resume automatically when its ongoing run completes.

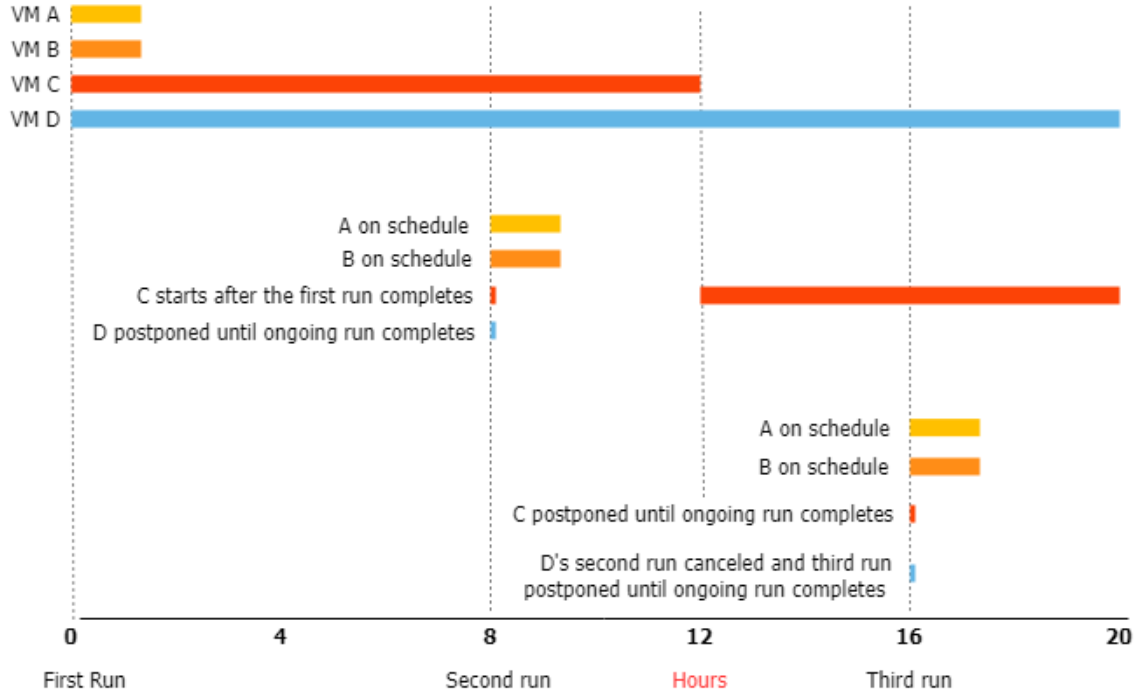

When the Cohesity cluster cancels the backup task of an object (such as D) in the Protection Group (such as PG1), the **Protection Group Details** page displays the following message at the corresponding protection run:

```
"Failed to back up some of the objects in the protection run as they are already backed in another protection run. ".
```

At the corresponding task (D's second backup task), the following message is displayed:

```
"cancelled the backup task of this object as it is already being backup up by another protection run."
```

Enabling the **Defer Incomplete Objects in Concurrent Runs** option helps you meet the defined SLAs, and improves the efficiency of the scheduled runs, thereby enhancing the backup performance.

> **Note: Defer Incomplete Objects in Concurrent Runs** can be enabled for runs with replication or archiving enabled.

16. **Leverage SAN Transport for Data Protection**: Cohesity supports SAN transport for Fibre Channel (FC) and iSCSI-based storage arrays and datastores that contain such storage arrays. Turn on this option to enable the SAN transport mode over FC or iSCSI for backing up VMware VMs to Cohesity cluster.

    > **Note:** Using the **Leverage SAN Transport for Data Protection** option with the **Leverage Storage Snapshots for Data Protection** option is not recommended.

    To use the SAN transport mode over FC, ensure the following pre-requirements:

    - Fibre channel is installed and zoned appropriately between the Cohesity cluster and the primary storage (Pure Flash Array and HPE Nimble Storage Array).
    - Volumes are mapped appropriately between the Cohesity cluster, ESXi hosts, and the primary storage.

    You can turn on the **Allow NBDSSL Transport Fallback** toggle if you do not want to enforce strict SAN transport mode for backing up VMware to the Cohesity cluster. With toggle enabled, the backup will not fail if the SAN transport is not available. If you disable this toggle, backup will fail if SAN transport mode is unavailable. NBDSSL encrypts all data passed over the TCP connection.

    > **Note:** Fallback to NBD is no longer supported if NBDSSL fails or is disabled.

    > **Note:** Ensure that the storage array connects to the primary environment (host) and the Cohesity node over the same protocol (FC on both ends or iSCSI on both ends). Cohesity does not support SAN transport for SAS (Serial Attached SCSI) based storage arrays.

17. **Exclude Disks**: (*Applies only to VMware VM*) By default, all the volumes on a selected server are protected. Toggle on to select disks to exclude for all VMs is this Protection Group. Provide the type of controller, bus number and unit number for each disk to exclude. Excluded disks are not backed up and are not recovered during VM recovery.

    > **Note:** If you exclude a disk that is part of a striped volume, the Cohesity cluster does not index the volume even if **Indexing** is toggled on for this Protection Group. You will not be able to search for or recover individual files in that volume.

18. **Exclude Physical RDM Volumes**: (*Applies only to VMware VM*) VMware vSphere Storage APIs cannot be used to snapshot physical compatibility RDM disks. Enable this option to allow the backup of VMs by excluding physical compatibility RDM disks. If you disable this option, the backup of VMs with physical compatibility RDM disks will fail. Excluded disks are not backed up and are not recovered during VM recovery. Cohesity recommends you to use agent-based backups to protect data on the physical RDM disks.

    > **Note:**
    > - VMware does not support taking the snapshot or backup of the VMs having RDM (physical and virtual) on Shared Bus.
    > - Backup of Virtual compatibility RDM disks is supported and does not depend on the **Exclude Physical RDM Volumes** toggle setting.

19. **App Consistent Backup**: Enable if you want the guest operating systems of all the VMs in the Protection Group to be quiesced before snapshots of these VMs are created. When this option is selected, the Cohesity cluster requests the VMware vSphere/AHV software to create a quiesced VM snapshot by invoking VMware Tools/Nutanix Guest Tools (installed on the guest operating systems of the VM). The VMware Tools/Nutanix Guest Tools requests that applications on the guest OS quiesce their state so that application-consistent snapshots can be created. Quiescing of VMs before capturing snapshots ensures the integrity of the data saved in the snapshots.

    If the Cohesity cluster fails to capture an app-consistent snapshot, you can use the **Take a Crash Consistent backup if unable to perform an App Consistent backup** option to capture a crash-consistent snapshot. The option is available only if the **App-Consistent backups option** is enabled.

    For example, the Cohesity cluster may fail to perform an app-consistent backup for the following reasons: when VMware Tools/Nutanix Guest Tools is not installed on the VM, the VM (VMware) is powered off, or if the VM cannot be quiesced. If the **Take a Crash Consistent backup if unable to perform an App Consistent backup** is enabled in that Cohesity cluster, a crash-consistent snapshot is captured.

    > **Note:** App-consistent backups apply to VMs only. For physical servers, Windows is app-consistent by default and Linux is crash-consistent. For more information, see Creating Application-Consistent Snapshots.

20. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.

21. **Cloud Migration:** Enable this option to support the cloud migration of VMs between hypervisors (such as vCenter) servers and cloud services for failover and failback. Cohesity agents must be installed on the Windows VMs prior to backing up the snapshots on the on-premises Cohesity cluster. For more information, see Disaster Recovery of VMs Using AWS Cohesity Cloud Edition.

    After enabling cloud migration, you can download the Cohesity Agent directly from the Cohesity Dashboard.
22. **Cancel Runs at Quiet Time Start**: Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
23. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.
24. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
25. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

    > **Note:**
    > The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).
26. **Description:** Optional. Click the plus sign and enter a description for the Protection Group.
27. Click **Protect**.

The Protection page is displayed and includes the new Protection Group.

> **Note:**
>
> - If a server is being backed up and files will also be recovered to that same server, the file recovery task starts after the backup is completed. The converse is also true; if files are being recovered to a server and that same server will be backed up, the backup starts after the file recovery is completed.
> - For a File-Based backup, the "Protected" and "Protected Size" fields on the source page for physical sources are updated after the source refresh.

Recover Organizations, vApps, and VMs to a New or Original Location

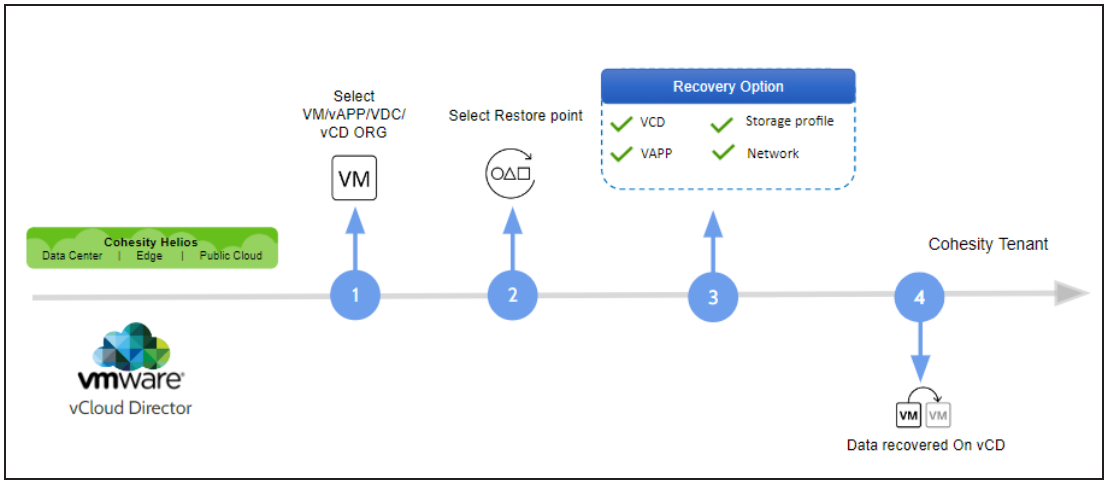

You can recover any object in vCD, such as an organization, data center, VM, vApp, and so on. You can recover the objects to the original (same location as the source) or a new location. After selecting the restore point, you can modify the storage domain and network. While recovering, you can either auto-deploy a Cohesity Agent, use an existing agent, or use VMware tools.

You recover the VM files (such as the VMDK files) to their original datastores and create new instances of the VMs in the original location in the original Source.

Before you can recover VMs to their original location, snapshots of those volumes must exist on the Cohesity cluster. To create Snapshots, create and run a Protection Group.

The following GIF illustrates the instant recovery of a Linux VM to a vApp in a new location:

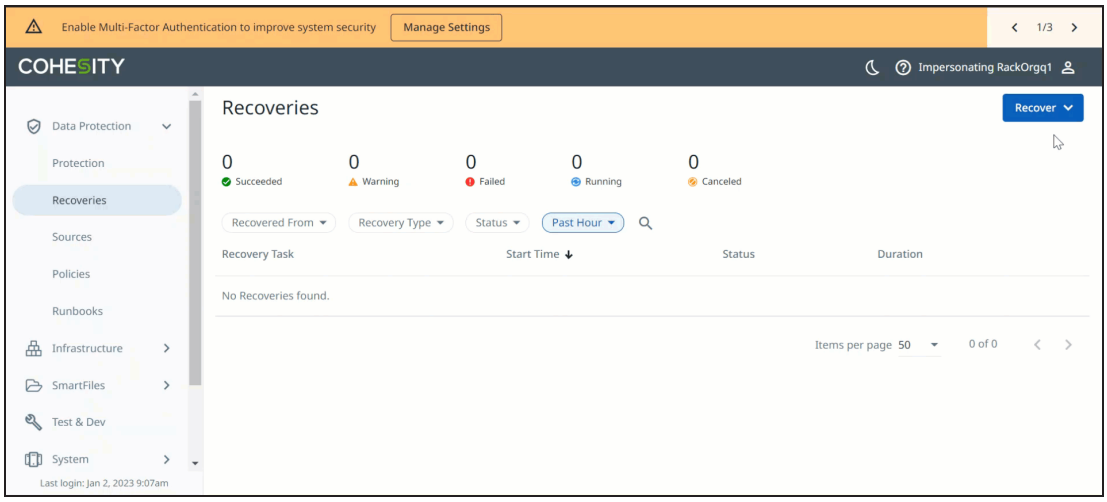

The following GIF illustrates the copy recovery of a vAPP to the original location:

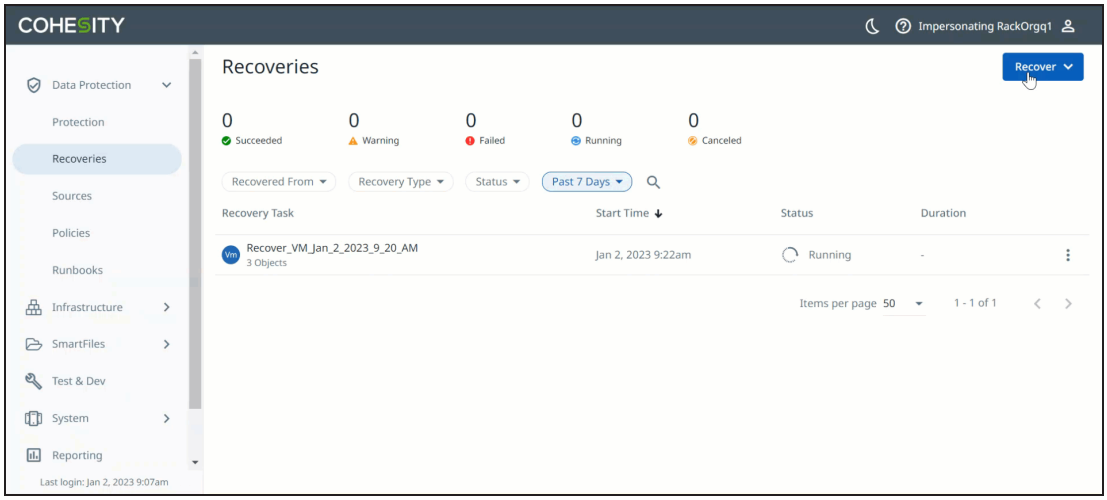

To recover VMs to its original location:

1. Navigate to **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > VMs**.

   The Virtual Machines page appears.
3. Search for a Protected Object (VM) or Protection Group that contains the snapshots to recover from by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.

4. From the list of search results, select one or more VMs or Protection Groups that contain the snapshots to recover.

   > **Note:**
   > - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
   > - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

5. By default, the latest snapshot of the VMs is selected for recovery.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   To recover from a different snapshot, perform the following steps depending on the object selected:

   - From the **Selected** list, click the edit icon.
   - From the **Protection Run** drop-down, select the snapshot from which you want to recover the Organization. By default, the snapshot is recovered from the local Cohesity cluster.
   - Click **Select Recovery Point**.

6. Click **Next: Recover Options**.
7. On the **New Recovery** page, select **Original Location** to recover the VM to the original location or **New Location** to recover the organization or its entities to an alternate datastore and create new instances in the new location of a registered Source.
8. If you selected New Location in step 7, enter:
   - **Registered Source**: Select a registered vCloud Director source to which you plan to recover the VMs or vApp.
   - **Organization VDC**: Select the target organization VDC.

- **vApp**: Optional. Select the target vApp. If not specified, VMs are recovered as standalone VMs inside the target organization VDC.
- **Storage Profile**: Displays all the Storage Profiles assigned to that Organization. Select a target Storage Profile. The Cohesity cluster chooses any compliant datastore associated with the Storage Profile to store the VM files. If you are recovering multiple VMs, the Cohesity cluster chooses one or more datastores associated with the Storage Profile.
- **Catalog**: (Applies only for vApp template recovery) Select the catalog for the vApp template recovery.
- **Placement Policy**: Select a placement policy to control where the recovered VM is placed. The selected placement policy determines the eligible hosts used during recovery.

9. If you selected, Original Location in Step 7 to recover the VMs, from the Existing VM Handling options list, choose how you want to handle the original VM when recovering to the original location:
    - **None**: Recovers the VM as a new VM and retains the original VM. In the Rename field available in the Recovery Options, you can specify the name of the recovered VM by adding prefix and/or suffix strings to the name of the original VM. By default, the recovered VM is named copy-<original_VM_name>.
    - **Overwrite Existing VM**: Recovers the VM by deleting the original VM. The recovered VM will have the original VM name. The original VM is deleted before the recovery. Therefore a recovery failure will also lead to the loss of the original VM.
    - **Attempt Differential Recovery**: Attempts to recover the VM by overwriting only the difference between the original VM and the snapshot selected for recovery. Any newly added data in the original VM is deleted.
        - This option is available only if you have selected **Overwrite Existing VM** and the **Recovery Type** option as **Copy Recovery**.
        - This option is disabled if you have selected **Leverage SAN Transport** under Recovery Options.
        - You can choose this option if there are no hardware configuration changes involved in the original VM.
        - If you want to reclaim free space for thin-provisioned disks, then Cohesity recommends not to attempt differential recovery and only perform **Overwrite Existing VM** recovery.
        - Differential recovery reduces the amount of data transfer in a recovery process. In the pulse log, you can view the amount of data transfer saved by selecting differential recovery.

- If the original VM is not present or if the attempt to differential recovery fails, then Cohesity will perform **Overwrite Existing VM** recovery.
- In the original VM, if there are any newly added disks or any disks that are excluded during backup, then the recovered VM will not have these newly added disks and the disks excluded during backup. If you want to retain the newly added disks and the excluded disks in the original VM, contact Cohesity Support. However, Cohesity recommends you to choose the **Keep Existing VM** option if you want to retain the newly added disks and the excluded disks in the original VM while recovering.
- All the snapshots present on the original VM are consolidated and removed as part of differential recovery.

- **Keep Existing VM**: Recover the VM as a new VM with the original VM name. The original VM is retained and powered off. The original VM will be renamed to obsolete_<time_stamp><Orginal_VM_Name>.

10. Select the **Recovery Type**:

> **Note:**
> - You should have the `VirtualMachine.Config.Rename` permission on your user account to use this recovery option.
> - For VMware VMs, Virtual RDM disks are recovered as flat disks.

- **Instant Recovery**: The VMs will be instantly available after recovery in the target location. The data will be moved to the target storage later. The VM can be used in production during Instant Recovery when data is copied to the VM from the Cohesity cluster. In such scenarios the VM performance may be low as the Storage vMotion is still in progress.

  You can instantly recover vCloud Director VMs to their original location if all the vCloud Director VMs of the Protection Group belong to the same 'tag based placement rule' Storage profile.

  If you choose to keep the VM in a powered-off state when initiating instant recovery, wait for the Storage vMotion to complete before powering the VM on. Cohesity recommends that you choose Copy Recovery for powered-off VMs for better performance.

  > **Note:** For instant recovery to Standalone ESXi, the VM(s) will be available immediately in the target environment and will have to be moved to the target storage manually.

- **Copy Recovery**: The VMs will be available in the target location only after all the data is copied to the target storage from the source location. Once the restore is complete, the VM will be available to use and the performance of the VM will be normal.

11. (Applies only for vApp template recovery) From the **Recover As** options, select how you want to recover the vApp template:
    - **vApp Template**: Select this option if you want to recover the vApp template as a vApp template itself.
    - **vApp**: Select this option if you want to recover the vApp template as a vApp.
12. Optional. Edit the default settings of the following **Recovery Options**:

Recovery Options	Description
Network	By default, the VMs that are to be recovered do not have a virtual Network Interface Card (vNIC) attached. Enable the **Attach** option to attach a virtual Network Interface Card (vNIC) to each VM that is to be recovered. With the **Attach** option enabled, each VM will connect to its original network and the following option is displayed: • **Start Connected**– Enable this option to connect to the original network when the VMs reboots. If this option is disabled, the VMs are not connected to any network on reboot. When recovering VMs to their original location, the original IP mode is restored to the respective NICs and the original IP address on each NIC is restored only if the original IP mode is set to “Static-Manual.” When recovering VMs to a new location, the original IP address and IP mode are not recovered to the NIC cards. By default, the IP mode is set to "none" and the IP address remains empty. Contact Cohesity support to recover the original IP addresses to a new location. While recovering VMs to the original location, there is no explicit control to preserve the MAC address. If the **Attach** option is enabled and the VM is configured to have a manual MAC address assigned at the time of backup, the same address will be recovered. If the VM is configured to have an automatic MAC address at the time of backup, the same setting will be applied to the recovered VM. For information on preserving MAC address when recovering VMs to a new location, see Recover VMs to a New Location.

Recovery Options	Description
Rename	(*Applies for VM recovery*) Add Prefix and/or Suffix strings to the names of the new VMs created by this task. For example you could specify Prefix `Test_ and Suffix _QA` to the VM name WinOS and when the VM is recovered, the new VM that is created is called `Test_WinOS_QA`.
Rename vApp/vApp Template	(*Applies for vApp or vApp template recovery*) Add prefix and/or suffix strings to the names of the new vApp or vApp template created by this task. For example, you could specify prefix `Test_ and Suffix _QA` to the vApp template name WinOS and when the vApp template is recovered, the new vApp template that is created is called `Test_WinOS_QA`.
Rename VMs within vApp	(*Applies vApp or vApp template recovery*) Add prefix and/or suffix strings to the names of the new VMs within vApp created by this task. For example, you could specify prefix `Test_ and Suffix _QA` to the VM name WinOS and when the VM within the vApp is recovered, the new VM that is created is called `Test_WinOS_QA`.
Power State	Disable **Power On** if you want the recovered VMs to remain powered off after they are created. If you recover VMs from a Storage Domain that has Cloud Tier enabled, Cohesity recommends leaving the VMs powered off. By default, the VMs are powered on after they are created.
Continue on Error	Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cohesity cluster will still attempt to create the other VMs. If this option is disabled, then the recovery task is discontinued if VM creation fails. However, if the VM creation succeeds, but VM storage migration fails, the task continues. By default this option is disabled.
Preserve VMs UUID	Select **Preserve VMs UUID** if you want to preserve the UUID for the recovered VMs. This is applicable only to Microsoft Hyper-V VMs.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
Task Name	Change the default name of the recovery task.

13. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Files or Folders in a vCD Source to a New or Original Location

To recover files and folders to the original server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select...	Then...
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. > **Note:** When entering the path of a Windows VM, be sure to reformat as shown in the example: Instead of *C:\Windows\System32\winevt* use *C:/Windows/System32/winevt/*. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. 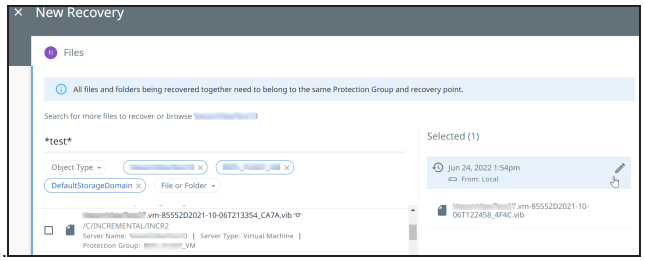 The Edit recovery point page appears. 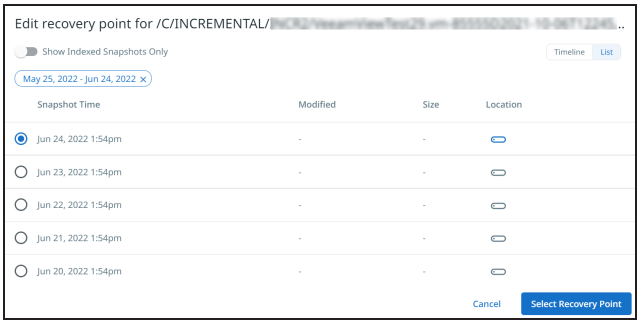 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option.

If you select...	Then...
	3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select...	Then...
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select...	Then...
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. From the **Recover To** drop-down:

- Select **Original Server** to recover the files or folders to the original server.
- Select **New Server** to recover the files or folders to a new location.

5. If you selected **New Server**, in the previous step, enter:
   - From the **Source** drop-down, select a registered server to which you want to recover the files and folders. You can also register a new source by clicking Register Source from the Source drop-down.
   - From the **Target** drop-down, select a VM to which the files and files are to be recovered.

     For recovering files and folders of a VMware VM, ensure that VMware Tools are installed and running on the target VM. The target list displayed includes only VMs with VMware Tools installed and running.
6. Once you select the files and folder to recover, select one of the following options:
   - **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
     Example:
     - The original name is "a:b" then the archived name is "a_COLON_b"
     - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

     If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
     Example:
     - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
     - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"
7. Select one of the following **Restore Method** options and provide the required information:
   - **Auto Deploy Cohesity Agent**: Select this option to automatically deploy the Cohesity Agent on the target VM. This allows you to use domain credentials to access the target VM. If you are using domain credentials, then you should provide the username in the format <username@domain.com> and the user

should be either a domain administrator or a member of the local administrators group on the server.

- **Use Existing Cohesity Agent**: Select this option to use the existing Cohesity Agent deployed on the target VM.

  > **Note:** To recover files and folders of VMware VMs using Cohesity Agent, VMware tools must be installed and running on the target VM.

- **Use VMware Tools**: Select this option to use the VMware Tools for recovery. You have to enter the credentials required to access the target VM. If you are using domain credentials to access a target VM, then you should provide the username in the format <username@domain.com>. Ensure to install the VMware Tools on the target VM.

8. Optional. Edit the default settings of the following **Recovery options**:

   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

     - C:\restore\1.txt
     - C:\restore\cohesity<taskid>_0\1.txt
     - C:\restore\bar
     - C:\restore\cohesity<taskid>_1\bar

     If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

   - **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only, files, then files will receive the new timestamps.
   - **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.

- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.
- **Save Success Files**: Enable this option to download the success logs for a recovery run.

9. Click **Recover**.

Protect vCD resources using the Cohesity vCD Plugin

Cohesity Plugin for VMware vCloud Director enables cloud providers to offer data protection as a service in multi-tenant environments. The vCD Plugin provides a seamless experience by facilitating backup and recovery operations in the same vCD console. It delivers self-service securely through role-based access control. The vCD Plugin's integrated data protection service streamlines operations and provides a rich tenant experience.

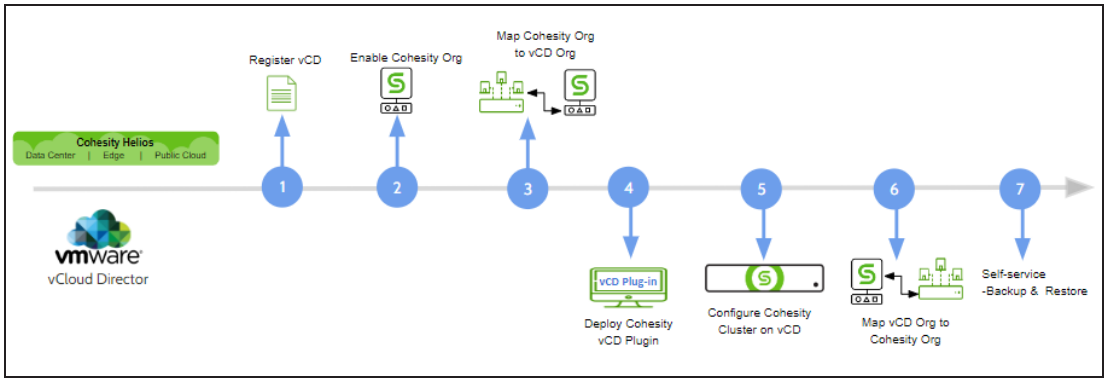

Data Protection using the Cohesity vCD plugin involves the following steps:

1. Ensure that you meet the requirements.
2. Manage access to the tenants and providers.
3. Setup the Cohesity vCD Plugin
   a. Register vCD as a source with the Cohesity cluster
   b. Enable Multi-tenancy on Cohesity cluster
   c. Create Organizations and Map to vCD Source
   d. Install the Cohesity vCD plugin
   e. Configure the Cluster and Map Tenants
4. Backup vCD Objects

   For details on how to protect, a single vCD object or multiple objects as a tenant, see Backup VMs, vApps, or vApp Templates.
5. Recover vCD Objects

- Restore VMs
- Restore vApps
- Restore Files or Folders

Prerequisites for Cohesity vCD Plugin

The prerequisites are as follows:

- Signed and trusted certificate must be installed on the Cohesity cluster.
- Cohesity cluster must be accessible through port 443.

Supported Software

Software	Version
vCloud Director	See the Compatibility Matrix.
Cohesity Data Platform	See the Compatibility Matrix.
Web browsers	Google Chrome, Firefox, Safari

Compatibility Matrix

Plugin Version	vCD Version	Cohesity Cluster Version	Notes
2.5.x	10.5.x, 10.4.x	7.0.x, 7.1.1, 7.1.2	vCD 10.5 and 10.4.2 are supported in 7.0.1 and above versions.
	10.3.x	6.6.x, 6.8.x, 7.0.x, 7.1.1, 7.1.2	
2.4.x	10.4.2	7.0.1	
	10.4, 10.4.1	7.0	vCD 10.4 and 10.4.1 are supported in 7.0 and above versions.
	10.1.x, 10.2.x, 10.3.x	6.6.x, 6.8.x, 7.0	

Plugin Version	vCD Version	Cohesity Cluster Version	Notes
2.3.6	10.0.x, 10.1.x, 10.2.x, 10.3.x	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x, 6.8.x	• vCD 10.3 is supported in 6.6.0d_cpd and above versions. • vCD 10.3.1 and 10.3.2 are supported in 6.6.0d-p6. You must contact Cohesity Support to enable the backup and recovery support for vCD 10.3.1 and 10.3.2. • vCD 10.3.3 is supported in 6.6.0d_u5 and above versions. • Cohesity cluster version 6.8 is supported with an additional patch. Contact Cohesity Support for the patch binaries.
2.3.5	10.0.x, 10.1.x, 10.2.x, 10.3.x	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x, 6.8.x	• vCD 10.3 is supported in 6.6.0d_cpd and above versions. • vCD 10.3.1 and 10.3.2 are supported in 6.6.0d-p6. You must contact Cohesity Support to enable the backup and recovery support for vCD 10.3.1 and 10.3.2. • vCD 10.3.3 is supported in 6.6.0d_u5 and above versions. • Cohesity cluster version 6.8 is supported with an additional patch. Contact Cohesity Support for the patch binaries.
2.3.4	10.0.x, 10.1.x, 10.2.x, 10.3.x	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x, 6.8.x	• vCD 10.3 is supported in 6.6.0d_cpd and above versions. • vCD 10.3.1 and 10.3.2 are supported in 6.6.0d-p6. You must contact Cohesity Support to enable the backup and recovery support for vCD 10.3.1 and 10.3.2. • vCD 10.3.3 is supported in 6.6.0d_u5 and above versions. • Cohesity cluster version 6.8 is supported with an additional patch. Contact Cohesity Support for the patch binaries.

Plugin Version	vCD Version	Cohesity Cluster Version	Notes
2.3.3	10.0.x, 10.1.x, 10.2.x, 10.3.x	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x	• vCD 10.3 is supported in 6.6.0d_cpd and above versions. • vCD 10.3.1 and 10.3.2 are supported in 6.6.0d-p6. You must contact Cohesity Support to enable the backup and recovery support for vCD 10.3.1 and 10.3.2. • vCD 10.3.3 is supported in 6.6.0d_u5 and above versions.
2.3.1	10.0.x, 10.1.x, 10.2.x, 10.3.0, 10.3.1, 10.3.2	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x	• vCD 10.3.0 support is only with Cohesity cluster 6.6.0d (CPD patch). • vCD 10.3.1 and 10.3.2 are supported in 6.6.0d-p6. You must contact Cohesity Support to enable the backup and recovery support for vCD 10.3.1 and 10.3.2.
2.3.0	10.0.x, 10.1.x, 10.2.x, 10.3.0, 10.3.1, 10.3.2	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x	• vCD 10.3.0 support is only with Cohesity cluster 6.6.0d (CPD patch). • vCD 10.3.1 and 10.3.2 are supported in 6.6.0d-p6. You must contact Cohesity Support to enable the backup and recovery support for vCD 10.3.1 and 10.3.2.
2.2.0	10.0, 10.1, 10.2, 10.2.2	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x	
2.1.2	10.0, 10.1, 10.2, 10.2.2	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x	
2.0.3	10.2	6.5.1b (CPD), 6.5.1d (CPD), 6.6.x	
2.0.3	9.5, 9.7, 10.0, 10.1	6.3.1, 6.4.1, 6.5.1, 6.6.x	

Access Management: Privileges for Service Providers and Tenants

A Service Provider can perform the following:

- Register vCD as a source with the Cohesity cluster
- Enable Multi-tenancy on Cohesity cluster
- Create Organizations and Map to vCD Source
- Install the vCD plugin
- Configure the Cluster and Map Tenants
- Edit or remove the Cohesity cluster configuration in the vCD Web Console.

A tenant can perform the following:

- Discover VMs
- Backup VMs, vApps, or vApp Templates
- Auto-protect vApps and vApp Templates
- Exclude VMs from an auto-protect job
- Unprotect a protected VM
- Edit Protection Group
- Delete Protection Group (including or excluding snapshots)
- On-demand Protection Runs
- Restore VMs
- Restore vApps
- Restore Files or Folders
- Monitor vCD Backup and Restore Tasks
- Access Cohesity Cluster from the Cohesity vCD Plugin

For more details, see Privileges for Service Provider Admin and Organization Admin.

Setup the Cohesity vCD Plugin

Setting up the Cohesity vCD Plugin involves the following steps:

1. Register vCD as a source with the Cohesity cluster
2. Enable Multitenancy on Cohesity cluster
3. Create Organizations and Map to vCD Source
4. Install the Cohesity vCD plugin
5. Configure the Cluster and Map Tenants

Register and Manage a vCD Source

Register vCD as a source

Once registered, its objects can be protected or backed up on the Cohesity cluster.

To register a source to back up:

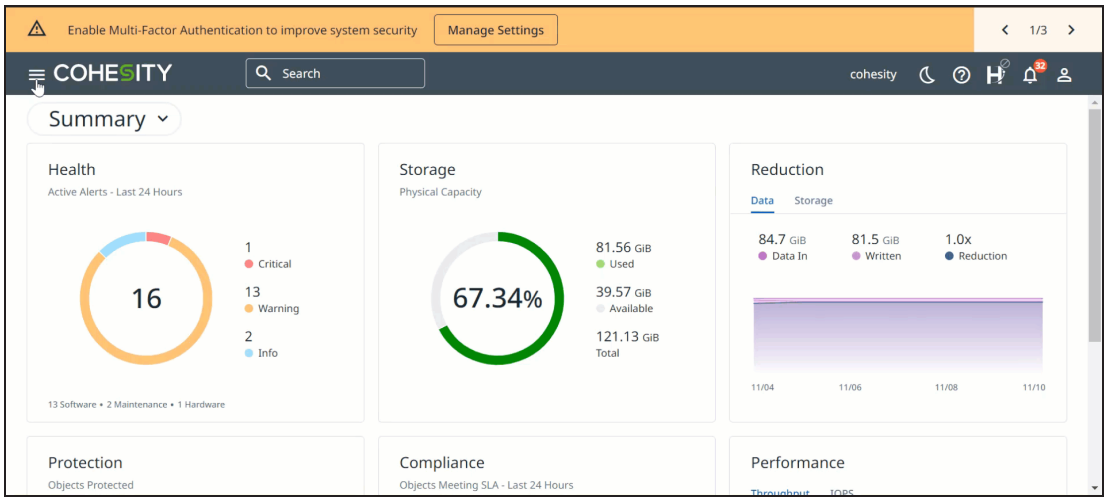

1. On the **Dashboard** page in the Cohesity cluster, navigate to **Data Protection** > **Sources**.
2. Select **Register** > **Virtual Machines**.
3. From Source Type, select **VMware: vCloud Director**.
4. Enter the following vCD details:
   - **Hostname or IP Address**: Specify the hostname or IP address of the vCloud Director that contains the vCenters to protect.
   - **Password**: Enter the password for the specified user.
   - **Username**: Specify the local administrator user.
5. Click **Connect**.
6. Enter the following vCenter details:
   - **Username**: You must specify a VMware vCenter user with adequate privileges on the VMware vCenter server. The role assigned to the specified VMware vCenter user should include privileges to turn on Change Block Tracking (CBT) and other privileges required by the Cohesity cluster. For detailed information about the necessary privileges, see Create a Role with Necessary vCenter Server Privileges.

     When specifying a Windows Active Directory username, use the following format: *<domain>\<username>* or *<username>@domain*. For all other cases, just enter: *<username>*.
   - **Password**: Enter the password for the vCenter.
   - **Detect VM migration across vCenter to Preserve backup chain**: Toggle this option on if you want your Cohesity cluster to detect the VMs that are migrated from the current VMware vCenter to vCD.
7. Click **Register**.

After source registration finishes successfully, the new source is displayed on the **Sources** page.

View the Registered vCD Source

To view a registered vCD Source:

1. Navigate to **Data Protection** > **Source**.

   The **Sources** page displays all the sources registered with that Cohesity cluster. The vCD ( ) icon indicates that the registered source is a vCD Source.

   Click the Actions menu (⋮) of that source to Edit, Refresh, or Unregister. From the Actions menu (⋮), you can also Deploy a Cohesity Agent and register Oracle, MS SQL, and Physical Servers.
2. Click on a source to view the hierarchy of the data centers, clusters, ESXi hosts, folders, VMs, vApps, and templates in that vCD source.
   - A selected check box ( ) indicates the object is selected to be backed up.
   - - Indicates a Linux VM.
   - - Indicates a Windows VM.
   - - Indicates a Windows or Linux template.
   - - Indicates a vApp.
   - - Indicates a VM
   - - A blue shield indicates a protected VM.
   - - Indicates an excluded object that is not protected. An ancestor of the object is auto-protected, but this object is excluded through inheritance (an ancestor is explicitly excluded).
   - - Click to view the general information—type, server size, and version of the selected object.
3. Filters on the Objects tab help you display only the objects you want. The filtering options are:
   - **Protection Status**: Protected or Unprotected.
   - **VM Attributes**: VMware Tools Found, MS SQL Servers, vCenter Migrated, Cohesity Agent Installed or Cohesity Agent Not Installed.

- **Host Type**: Linux or Windows.

4. Enter the name of an object in the **Search** box. You can even enter a partial name or an asterisk * wildcard. As you type, the objects that contain your search term appear.

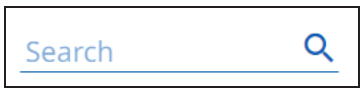

5. Display modes can show the objects in a hierarchy or a flat list. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection** > **Sources**, find the source in the list, and click the Refresh ( ) icon.

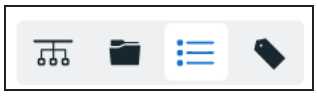

Update the vCD Source Configuration

You can update the vCD configuration you provided during the registration process with the latest configuration. To edit the vCD source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Actions** menu (⋮) next to the vCD source and select **Edit**.
3. In the **Edit Source** page, update the respective configurations.
4. Click **Save**.

Refresh the vCD Source

You can refresh the vCD source configuration to fetch the latest changes. To refresh the vCD source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Actions** menu (⋮) next to the vCD source and select **Refresh**.

Partial Refresh of the vCD Souce

You can perform a partial refresh of the vCD source using the API. A partial refresh is faster than a full source refresh. This enhancement allows you to selectively discover new VMs, vApps, Organizations, Org VDC, and vApp templates through the APIs into the vCD source, eliminating the need to wait for a full vCD refresh. You can refresh only the selected objects, such as specific VMs, vApps, Organizations, Org VDC, and vApp templates within the vCD source.

To do a partial refresh of the vCD source using the API:

Endpoint

```
PUT /v2/data-protect/sources/{sourceId}/objects
```

Parameters

Parameter	Type	Required	Description
sourceId	integer	Yes	The Entity ID of the data protection source

Finding the sourceId

You can find the sourceId from the UI by navigating to the Sources page and clicking your desired vCD source.

The URL will appear in the following format:

```
https://<Cluster IP>/protections/sources/details/{sourceId}/objects?environment=kVMware
```

The `{sourceId}` is the numeric ID you need for the API call. This value remains constant until the source is unregistered.

Headers

Header	Type	Required	Description
Content-Type	string	Yes	Must be `application/json`
Authorization	string	Yes	Bearer token for authentication

Request Body

The request body should include a JSON object with the following structure:

```
JSON
{
 "modificationRequests": [
 {
  "actionType": "Add",
  "vcdModifyParams": {
   "urn": "string"
```

```
}
}
]
}
```

Request Body Schema

Field	Type	Required	Description
modificationRequests	array	Yes	Array of modification requests to process.
modificationRequests[].actionType	object	Yes	Action to perform. Valid values: `Add`.
modificationRequests []. vcdModifyParams	string	Yes	vCD-specific parameters for the modification.
modificationRequests []. vcdModifyParams.urn	string	Yes	The URN (Uniform Resource Name) of the VCD object to be discovered

vCD URN Format

VCD URNs follow the format:

```
`urn:vcloud:{type}:{uuid}`
```

Navigate to **Data Protection** > **Sources**.

The common VCD object types:

- `vapp` - Virtual Application
- `vm` - Virtual Machine
- `vapptemplate` - vApp Template
- `org` - Organization
- vdc` - Virtual Datacenter

To Perform a Partial Refresh:

1. **Authenticate**: Obtain Access Token using Authentication API:

   ```
   POST /irisservices/api/v1/public/accessTokens
   ```

2. Get Source ID of the vCD source as described above.
3. Discovery request API .

   ```
   PUT /v2/data-protect/sources/{sourceId}/objects
   ```

Example Requests

1. Discover a new VM.

```
curl --location --request PUT
'https://10.136.5.86/v2/data-protect/sources/1/objects' \
--header 'Content-Type: application/json' \
--header 'Authorization: Bearer YOUR_TOKEN_HERE' \
--data '{
 "modificationRequests": [
 {
 "actionType": "Add",
 "vcdModifyParams": {
 "urn": "urn:vcloud:vm:96228d49-a079-4b0e-a19f-c5e78acd65b1"
 }
 }
 ]
 }'
```

2. Discover a new VCD org.

```
JSON
{
 curl --location --request PUT
 'https://10.136.5.86/v2/data-protect/sources/1/objects' \
 --header 'Content-Type: application/json' \
 --header 'Authorization: Bearer YOUR_TOKEN_HERE' \
 --data '{
 "modificationRequests": [
 {
 "actionType": "Add",
 "vcdModifyParams": {
 urn": "urn:vcloud:org:96228d48-a079-4b0e-a19f-f5e78acd65b1"
 }
 }
 ]
 }
```

Response Status Code and Body

1. Success Response (207 Multi-Status).

```
JSON
{
 "results": [
 {
```

```
"succeeded": true,
"vcdResult": {
"urn": "urn:vcloud:org:96228d49-a079-4b0e-a19f-c5e78acd65b1"
}
}
]
}
```

2. Error Responses.
   a. 403 Forbidden

```
JSON
{
 "errorCode": "KInvalidRequest",
 "message": "Invalid inputs. Unknown input urn type"
 }
```

   a. 500 Internal Server Error

```
JSON
{
 "errorCode": "KInternalError",
 "message": "VM not found in VCD"
 }
```

Response Status Codes

Status Code	Description
207	Multi-Status - Request processed, check individual results,
403	Forbidden - Invalid request or unknown URN type.
500	Internal Server Error - vCD object not found or server error.

> **Note:**
> - You can include multiple objects in the **modificationRequests** array to process several modification requests in a single API call. The batch operations are not supported for vCD partial refresh.
> - **URN Validation**: The API verifies that the specified URN exists in the connected vCD environment before processing the request.
> - **Hierarchical Dependencies**: When you add a vApp, OrgVDC, or Organization, all descendant entities are automatically included.
> - **vCD Hierarchy**: The system adheres to the vCD organizational hierarchy (Organization → VDC → vApp/vAppTemplate → VM) when processing modifications.
> - Ensure that your authorization token has the required permissions to modify data protection sources and access vCD objects.

Related Endpoints

Authentication API

```
POST /irisservices/api/v1/public/accessTokens` - Obtain access token for authentication
```

Authentication API Documentation: Obtain Access Token

Endpoint: `POST /irisservices/api/v1/public/accessTokens`

Description: Authenticate with the system and obtain a bearer token for API access.

Headers

| Header | Type | Required | Description |

| Content-Type | string | Yes | Must be `application/json` |

Request Body

Field	Type	Required	Description
username	string	Yes	Username for authentication.
password	string	Yes	Password for authentication.
domain	string	No	Authentication domain (e.g., "LOCAL") .

Example Request

```
curl --location --request POST
 'https://10.136.5.86/irisservices/api/v1/public/accessTokens' \
 --header 'Content-Type: application/json' \
 --data '{
 "username": "admin",
 "password": "admin",
 "domain": "LOCAL"
 }'
```

Success Response (201 Created)

```
JSON
{
 "accessToken": "eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9...",
 "tokenType": "Bearer",
 "privileges": [
 "AD_LDAP_MODIFY"
 "AD_LDAP_VIEW",
 "AGENT_UPGRADE_MODIFY",
 "PROTECTION_MODIFY",
 "PROTECTION_VIEW",
 "..."
 ]
 }
```

Response Fields

Field	Type	Description
accessToken	string	JWT token to use in Authorization header.
tokenType	string	Type of token (Bearer).
privileges	array	List of permissions granted to the authenticated user.

Usage

Use the returned `accessToken` in the Authorization header for subsequent API calls:

Authorization: Bearer <accessToken>.

This API allows you to perform a partial refresh to discover VMware vCloud Director (VCD) objects within a data protection source.

You can selectively refresh and add VCD entities such as vApps, VMs, VCD Orgs, and other VCD resources for backup and protection.

Unregister vCD Source

If you plan to stop taking the backup of your vCD source, you can unregister the vCD source from the Cohesity cluster. To unregister the vCD source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Actions** menu (⋮) next to the vCD source and select **Unregister**.

Configure the Organizations on Cohesity Cluster

Enable Multitenancy on Cohesity Cluster

To create a tenant or organization on a Cohesity cluster, you must first enable organizations on your cluster and then create a new organization.

> **Note:** The term tenant also refers to an organization, and these terms have been used interchangeably throughout these sections.

To enable the Multitenancy in a Cohesity cluster:

1. In the Cohesity cluster, navigate to **Settings** > **Summary** and click **Configure**.
2. Turn on the **Enable Organizations** toggle.

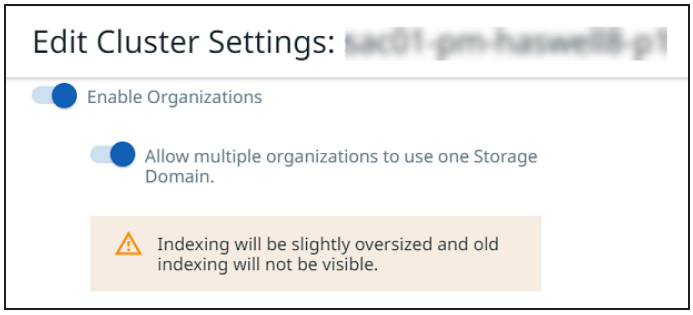

3. *(Optional)* Turn on the **Allow multiple Organizations to use one Storage Domain** toggle.

   When you select this option, the storage consumption on the cluster reduces by providing Cohesity deduplication across all the Organizations using the storage domain. Organizations will be isolated from other Organizations whether this option is on or off.

4. Click **Save**.

Create Organizations and Map to vCD Source

Before you Begin

When creating a new organization or tenant, vCD must be registered as a source.

To create a new organization in the Cohesity cluster:

1. On your Cohesity cluster, navigate to **Settings** > **Organizations**.
   a. Click **Add Organization**.
   b. Enter the Organization Name, ID, and Description.
   c. Click **Add**.

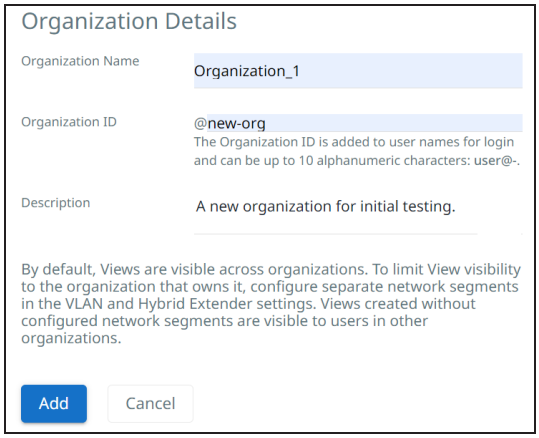

2. Select Storage Domain for the organization from the drop-down list. You can either select an existing domain or add a new Storage Domain.

   For more details on creating Storage Domains, see Create or Edit Storage Domains.
3. *(Optional)* Select **Views** from the drop-down list.
4. In the **Edit Organizations** page, click **Assign Users/Groups**.
5. In the **Assign Users/Groups** dialog, select an existing user and assign as an Administrator or create a new user. To create and assign a new user:
   a. Click **Add User**.
   b. Enter the user details and click **Add**.
   c. Select the newly created user and click **Assign**.

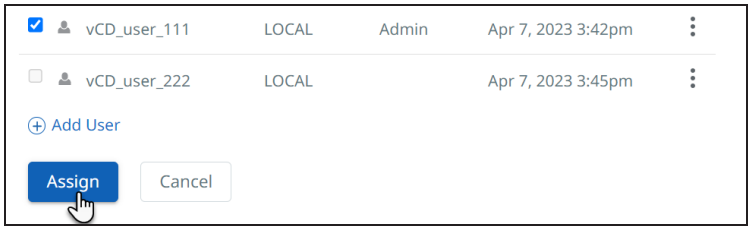

   For more details on adding users or groups, see Users and Groups for Organizations.
6. Under **Sources**, select **Assign Objects** to assign a source to this organization.
   a. Click on a vCD source to view the hierarchy of the data centers, clusters, ESXi hosts, folders, VMs, vApps, and templates in that vCD source.

b. Select the objects that are to be assigned to the newly created Cohesity tenant and select **Assign**. For more details, see Register and Manage a vCD Source.

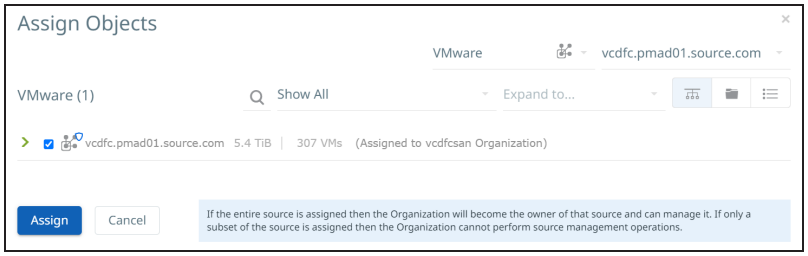

7. Click **Assign Policies and Protection Groups** to assign policies to the organization and select **Assign**. For more details, see Data Protection for Organizations.
8. Click **Save**.

   An organization or tenant is successfully created.

Install the Cohesity vCD Plugin

To install the vCD plugin, follow these steps:

1. Log into Github.
2. Under the **Release** tab, find your vCD plugin version.
3. Under **Assets**, click on **plugin.zip** to download the file.
4. Login to vCD as the provider.
5. Select **More** > **Customize Portal**.
6. Click **Upload**.
7. Click **Select plugin** file, select the **plugin.zip** file, and click **Open**. The plugin is uploaded and details of the plugin are displayed.

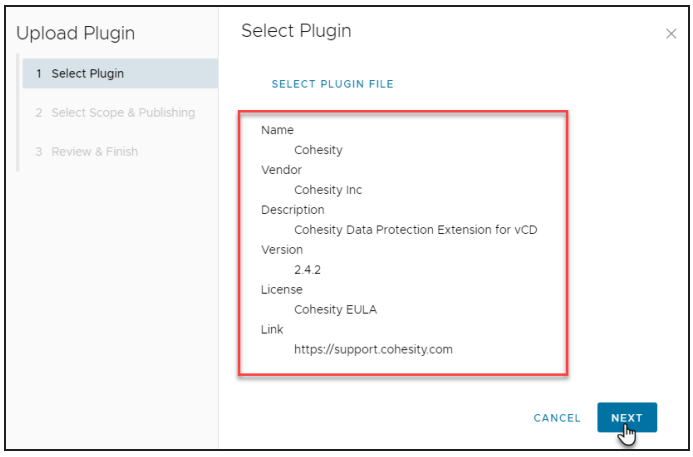

8. Click **Next**.
9. Select the scope for this plug-in:

- **Service Providers**: The plug-in function becomes available for all in the VMware Cloud Director Service Provider Admin Portal.
- **Tenants**: The plug-in function becomes available in the VMware Cloud Director Service Provider Admin Portal of the organizations that you select.

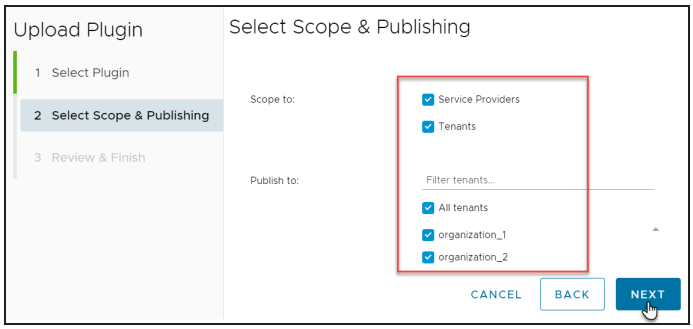

10. *(Optional)* Select the organizations to which you want to publish this plug-in. You can choose between all tenants or specific tenants.
11. Click **Next**.
12. Review the information and click **Finish**.

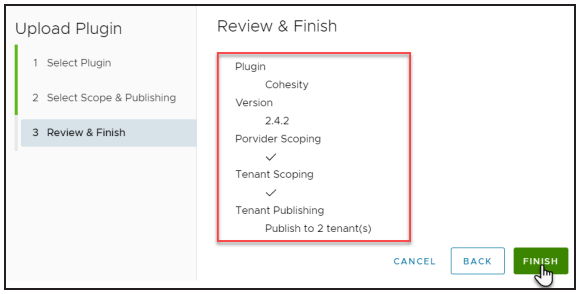

Verify the Cohesity extension is successfully installed

Log in to vCD, and select **More**.

The **Data Protection** option is displayed in the menu.

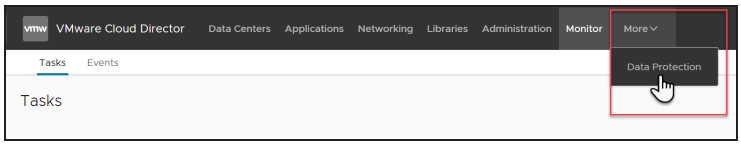

Upgrade the Plugin

> **Note:** Before you upgrade the plugin, delete the existing plugin.

The procedure for upgrading the vCD plugin is similar to installing the plugin. To upgrade the plugin, simply repeat the steps in the Install the Cohesity vCD Plugin section.

> **Note:** After upgrading the vCD plugin, you must manually refresh the vCD source before you start protecting your vCD objects.

Verify the Plugin Version

1. Log into vCD as a provider.
2. Select **More** > **Customize Portal**.

Among the list of plugins, the Cohesity vCD plugin version is displayed as shown in the below image.

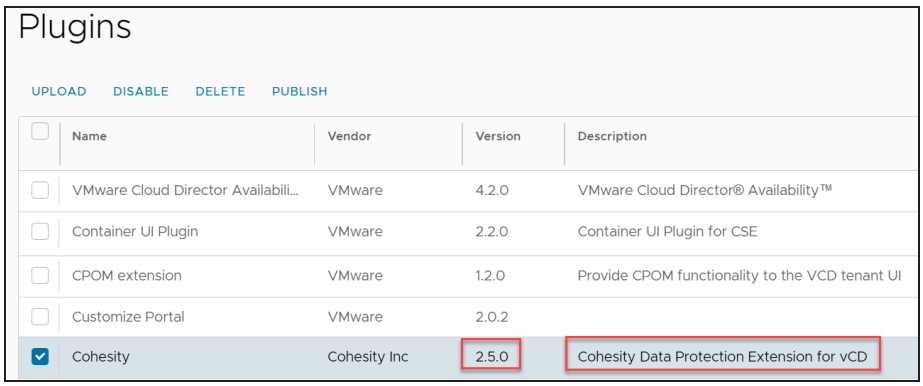

Configure the Cluster and Map Tenants

Before you Begin

The Cohesity vCD plugin makes cross-origin API requests to the Cohesity cluster. To enable support for Cross Origin Resource Sharing (CORS) headers on the Cohesity cluster, contact your Cohesity account team.

Configure the Cohesity Cluster in vCD Console and Map Tentants

After the Cohesity vCD Plugin is successfully deployed in vCD, you must configure the Cohesity cluster in the vCD Console and map the tenants or organizations created in the Cohesity cluster with a vCD organization.

> **Note:** At least one tenant must be mapped before you start protecting your vCD objects.

To map the organizations in a Cohesity cluster with a vCD organization, follow these steps:

1. Log into the vCD console as a provider with admin credentials.
2. Select **More** > **Data Protection**. The **Configuration** page displays.

3. Click **Add**. The **Add Cohesity Cluster** page displays.

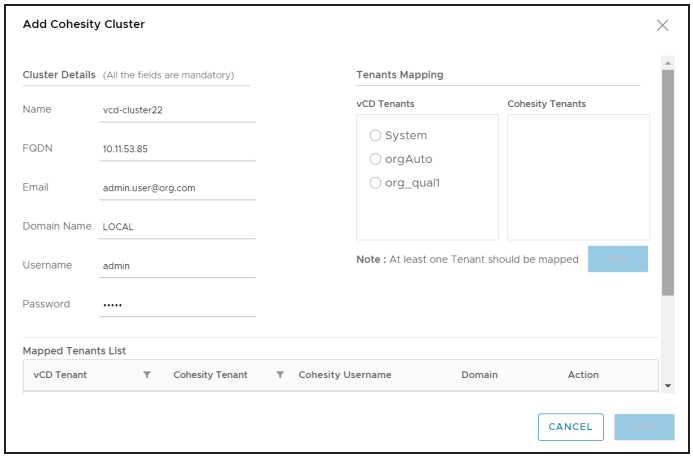

4. Enter the Cohesity cluster details such as IP address or hostname of the cluster along with the Service Provider credentials. These credentials are used to fetch the Cohesity tenants.
5. Once the credentials are validated, the Cohesity tenants are listed in the **Tenants Mapping** section.
6. Select the **vCD Tenant** and the **Cohesity Tenant** and click **Map**.

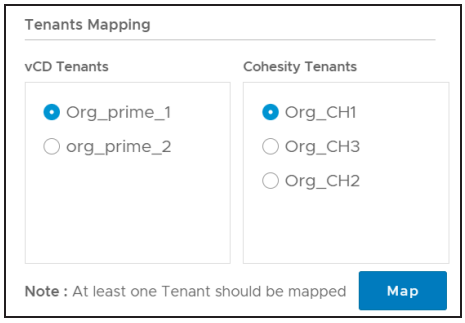

7. In the **Tenant User Credentials** prompt, you can either **Create New Local User** or select an **Existing User**. The tenant user credentials are used to execute all workflows for the tenant.

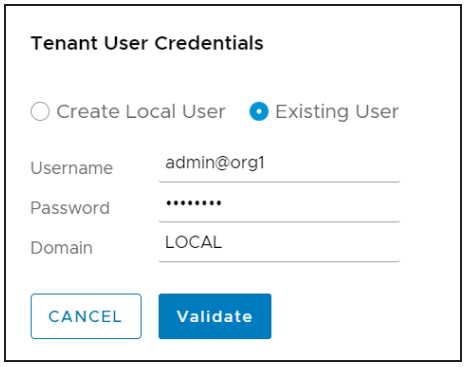

> **Note:**
> - If you create a new user, the Service Provider user must have privileges listed in the Access Management: Privileges for Service Providers and Tenants section.
> - If you are an existing vCD plugin user and upgrading your Cohesity cluster to 6.6.0c or later version, then you need to change the passwords of all mapped Cohesity Tenant Users to access Cohesity workflows.

8. Click **Save**.

   The Cohesity tenant is mapped successfully to the vCD tenant. The Cohesity cluster has been added and the vCD tenant user can now access Cohesity protection workflows through the tenant portal.

Backup VMs, vApps, or vApp Templates

Before you Begin

Discover VMs

The first step before we execute the vCD workflows is to make sure that VMs are discovered by Cohesity. You can check the status based on the **Protection Status** field.

1. Log into the vCD Console as a tenant user.
2. Navigate to **Data Protection** > **vApps** or **VMs**.
3. Check the status of all the VMs in the **Protection Status** field.
4. If the status is **Undiscovered** for any VM, click **Discover**. This initiates a refresh operation on the vCD source.

> **Note:** After upgrading the vCD plugin, you must manually refresh the vCD source before you start protecting your vCD objects.

Backup vCD Resources

There are two workflows for protecting and backing up your VMs, vApps, and vApp Templates as a tenant:

- Protect Individual VMs, vApps, and vApp Templates
- Protect Multiple Objects in a Single vCD Protection Group

Protect Individual VMs, vApps, and vApp Templates

To protect a specific vCD object, you must first create a Protection Group following the process below.

To protect individual VMs, vApps, and vApp Templates:

1. Log into the vCD Console as a tenant user.
2. Navigate to **Data Protection** > **vApps** or **VMs**.
3. Use the **Org vDC** and **Object Type** filter on the top to list to filter the objects based on the selected vDC or choose from the listed VMs, vApps, and vApp Templates.

   Verify the status of the object in the **Protection Status** column. If any of the objects that you want to back up are in an **Undiscovered** state, then you must first discover that object. See Backup VMs, vApps, or vApp Templates.

   Once the status of the object changes to **Unprotected**, click on the ellipsis (⋮) on the row for the unprotected object and select **Protect VM** (or **vApp** or **vApp Template**). If you choose **AutoProtect Object** (for a vApp or vApp Template), any VM that belongs to that object or any new VM created in that object will be protected automatically.
4. Click **Create New Group** in the **Protect Object** window.
   If you already have one, you can select an existing Protection Group instead.
5. Configure the Protection Group details and click **Save**.

   > **Note:** For more information on Leverage SAN Transport, see SAN Transport.
6. The Protection Status is updated as **Protected**.

   > **Note:** You can also log in to the Cohesity cluster at this point to verify that the Protection Group was successfully created.
7. Click the ellipsis (⋮) and either **Backup VM** (or **vApp** or **vApp Template**) **Now** or **Unprotect VM** (or **vApp** or **vApp Template**).
   If you enabled **Allow Archival** or **Replication Target Selection**when you configured vCD above, specify the **Archive to** and **Replicate to** targets. After that, click **Backup Now**.

Protect Multiple Objects in a Single vCD Protection Group

Another way of protecting vCD resources is by creating a Protection Group in vCD. This method allows you to select multiple resources to be protected in a single vCD Protection Group.

To protect multiple vCD resources in the same vCD Protection Group:

1. Log into the vCD Console as a tenant user.
2. Navigate to **Data Protection** > **Protection Group**> **New**.
3. In the **New Protection** dialog, enter the relevant information according to your requirement and click **Next** to select the objects to be protected.
   You can use options like setting **Email Recipients** to receive alerts for Protection Group **Success** and **Failure** runs.

   > **Note:** For more information on Leverage SAN Transport, see SAN Transport.

4. In the **Objects** page, click **Add Objects** to protect the vApp, vApp Template, or VM. There is also an option to specify **Auto Protection** for a vApp or vApp template. If you choose **AutoProtect Object** (for a **vApp** or **vApp Template**), any VM that belongs to that object or any new VM created in that object will be protected automatically.
5. Select the objects to be protected and click **Save Selection**.
6. Review the objects that are protected and click **Finish** to create the vCD Protection Group.
7. Run the Group by clicking the ellipsis (⋮) and selecting **Run Now**.
   You can also edit the Protection Group to add or remove protection objects or change the configured parameters.

   If you enabled **Allow Archival** or **Replication Target Selection** when you configured vCD above, specify the **Archive to** and **Replicate to** targets. After that, click **Backup Now**.

SAN Transport

Leverage SAN Transport for Data Protection: Cohesity supports SAN transport for Fibre Channel (FC) and iSCSI-based storage arrays and datastores that contain such storage arrays. Turn on this option to enable the SAN transport mode over FC or iSCSI for backing up VMware VMs to the Cohesity cluster.

> **Note:** Using the Leverage SAN Transport for Data Protection option with the Leverage Storage Snapshots for Data Protection option is not recommended.

To use the SAN transport mode over FC, ensure the following pre-requirement:

- Fibre channel is installed and zoned appropriately between the Cohesity cluster and the primary storage (Pure Flash Array and HPE Nimble Storage Array).

To use the SAN transport mode over FC, ensure the following pre-requirement:

- Volumes are mapped appropriately between the Cohesity cluster, ESXi hosts, and the primary storage.

  You can turn on the Allow NBDSSL Transport Fallback toggle if you do not want to enforce strict SAN transport mode for backing up VMware to the Cohesity cluster. With toggle enabled, backup will not fail if SAN transport is not available. If you disable this toggle, the backup will fail if the SAN transport mode is unavailable.

> **Note:**
> - Fallback to NBD is no longer supported if NBDSSL fails or is disabled.
> - Ensure that the storage array connects to the primary environment (host) and the Cohesity node over the same protocol (FC on both ends or iSCSI on both ends). Cohesity does not support SAN transport for SAS (Serial Attached SCSI) based storage arrays.
> - VMware SAN transport mode is not supported with Cohesity virtual nodes. To use SAN transport mode for backing up VMware VMs, zoning and granting datastores access to the Cohesity nodes has to be done manually. VMware SAN transport mode is not supported with Cohesity virtual nodes.

Recover VMs from a VCD Source

To restore a VM:

1. Log into the vCD Console.
2. Navigate to **Data Protection** > **Recovery**.
3. Click **VM**.
4. Search for a VM to restore.
5. Select the VM and click **Continue**.
6. In the **Restore VM** page, select the snapshot and provide a prefix and suffix.

> **Note:**
> - Prefix or suffix is required to ensure that the VM name is unique. By default, the VM is restored back to the original location. However, you can choose to recover to a new location by toggling on the option.
> - Restore to an alternate location is only supported for the Cohesity cluster running version 6.4.1c, 6.5.1 and later.
> - Using the Leverage SAN Transport for Data Protection option with the Leverage Storage Snapshots for Data Protection option is not recommended.
> - The placement and sizing policies for VMs are restored during recovery to both the original and alternate locations.

To use the SAN transport mode over FC, ensure the following pre-requirements:

- Fibre channel is installed and zoned appropriately between the Cohesity cluster and the primary storage (Pure Flash Array and HPE Nimble Storage Array).
- Volumes are mapped appropriately between the Cohesity cluster, ESXi hosts, and the primary storage.
- You can turn on the Allow NBDSSL Transport Fallback toggle if you do not want to enforce strict SAN transport mode for backing up VMware to Cohesity cluster. With toggle enabled, backup will not fail if SAN transport is not available. If you disable this toggle, backup will fail if SAN transport mode is unavailable. NBDSSL encrypts all data passed over the TCP connection.

> **Note:**
> - Fallback to NBD is no longer supported if NBDSSL fails or is disabled.
> - Ensure that the storage array connects to the primary environment (host) and the Cohesity node over the same protocol (FC on both ends or iSCSI on both ends). Cohesity does not support SAN transport for SAS (Serial Attached SCSI) based storage arrays.

7. Click **Restore**.
8. To check the status, navigate to **Recovery** to check for the recovery task as described in Monitor vCD Backup and Restore Tasks.

Recover vApps from a vCD Source

To restore a vApp:

1. Log into the vCD Console.
2. Navigate to **Data Protection** > **Recovery**.
3. Click **vApp**.
4. You can search for a specific vApp name or use wildcard characters such as * to list all the VMs in the vApp.
5. Select a specific vApp, and click **Continue**.
6. Configure the details:
    1. Under **vApp Details**, enter a prefix and suffix that will help you identify the vApp being restored.
    2. Under **VM Details**, enter a similar prefix.
7. Under **Restore Properties**, select the snapshot that you want to recover.

    > **Note:** Using the Leverage SAN Transport for Data Protection option with the Leverage Storage Snapshots for Data Protection option is not recommended.

    To use the SAN transport mode over FC, ensure the following pre-requirements:

    - Fibre channel is installed and zoned appropriately between the Cohesity cluster and the primary storage (Pure Flash Array and HPE Nimble Storage Array).
    - Volumes are mapped appropriately between the Cohesity cluster, ESXi hosts, and the primary storage.
    - You can turn on the Allow NBDSSL Transport Fallback toggle if you do not want to enforce strict SAN transport mode for backing up VMware to the Cohesity cluster. With toggle enabled, backup will not fail if SAN transport is not available. If you disable this toggle, the backup will fail if the SAN transport mode is unavailable. NBDSSL encrypts all data passed over the TCP connection.

    > **Note:**
    > - Fallback to NBD is no longer supported if NBDSSL fails or is disabled.
    > - Ensure that the storage array connects to the primary environment (host) and the Cohesity node over the same protocol (FC on both ends or iSCSI on both ends). Cohesity does not support SAN transport for SAS (Serial Attached SCSI) based storage arrays.

8. Click **Restore**.
9. To check the status, navigate to **Recovery** to check for the recovery task as described in Monitor vCD Backup and Restore Tasks.

Recover Files or Folders from a vCD Source

To recover a file or folder:

1. Log into the vCD Console.
2. Navigate to **Data Protection** > **Recovery**.
3. Click **Files and Folders**.
4. Click **Files and Folders** to search for a specific indexed file. You can search for a partial name or use a wildcard expression such as * to search.

   The search output is displayed in the table.
5. Select the file from the table and click **Add to Cart**.
6. From the list of snapshots, select the recovery point which will be applied to all files or folders recovered in the task.

   To restore multiple files in the same task, follow these additional steps:

   a. Click **Add More**.

   b. Select the files from the filtered list and click **Add to Cart**.

   > **Note:** Cohesity allows you to recover multiple files for recovery only when the selected files belong to the same Server, Protection Group, Cluster, and Snapshot.
7. Enter the user credentials for the VM. You can also choose to recover to a new location, in which case the recovery path should be mentioned.
8. Click **Restore**.

   To check the status, navigate to **Recovery** and click **Refresh** to check for the recovery task as described in Monitor vCD Backup and Restore Tasks.

If you have chosen **Browse or Specify Path**, you can search for non-indexed files. Enter the VM name to search, select the VM, click **Continue**. From the tree structure, browse and check the box of the file(s) that you want to recover, select the snapshot version from the drop-down menu on the top left corner, and click **Continue**.Fill in the configuration details and click **Restore**.

Download files or folders

You can also download files locally. Search for the file, click **Continue**, select the snapshot and click **Download Now.** To download the file, navigate to **Recovery**, select the job, right click and select **Show Details** and click **Download Files**.

Monitor vCD Backup and Restore Tasks

You can check the status of the backup and recovery tasks under the **Recovery** tab. The following can be monitored:

- Scheduled or on-demand backup
- VM or vApp recovery
- File or Folder recovery
- File or Folder download

Monitor Backup

To monitor backup and restore tasks, follow these steps:

1. Log into the vCD Console as a tenant user.
2. Navigate to **Data Protection** > **Protection Group**. All the protection groups that are present in the cluster are listed.
3. Select the protection group whose runs you need to monitor.
4. Click on the **Protection Runs** tab.

   All the backup runs (scheduled or on-demand) which were triggered in the last 7 days are listed.
5. You can use the **Date** filter to filter the protection runs.

Monitor Recoveries

You can monitor the status of all recovery tasks which were triggered through the vCD plugin.

To monitor the status, follow these steps:

1. Log into the vCD Console as a tenant user.
2. Click **Recoveries**. All the recovery requests made in the last 30 minutes.
3. Click on the task name to get the status of the recovery task.

   You can use the **Filter By** option to filter recovery tasks based on their type. Use the **Duration** filter to fetch requests based on time.

   > **Note:** You can only monitor the status of recovery tasks which were triggered in the last 7 days through the plugin.

Access Cohesity Cluster from the Cohesity vCD Plugin

vCD tenants can launch Cohesity UI from the vCD extension and access advanced workflows that are currently not available in the vCD extension.

> **Note:** This feature is disabled by default. Contact your Cohesity account team to enable this feature.

Enable Access to the Cohesity Cluster

To enable access to the Cohesity cluster:

1. Log into the vCD Console as a provider with admin credentials.
2. Navigate to **More** > **Data Protection.** The **Configuration** page is displayed.
3. Select **Settings** from the navigation pane on the left.
4. Click **Edit**. The **Edit Settings** dialog is displayed.

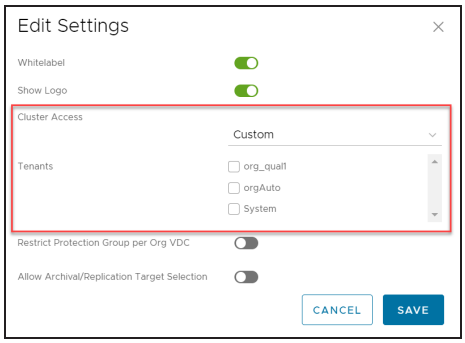

5. The **Cluster Access** field is disabled by default. Turn the **Cluster Access** toggle ON to enable Cohesity cluster access. You can select **All Tenants** to enable this feature for all your vCD tenants or select **Custom** to enable it for selective tenants.
6. Click **Save**.

Access the Cohesity cluster

To access the Cohesity cluster:

1. Login to vCD as a tenant user.
2. Navigate to **More** > **Data Protection**. The **Dashboard** page displays.
3. From the navigation pane on the left, expand the **Advanced** option to see the Cohesity clusters to which the tenants have access to.
4. Click on any cluster to launch the Cohesity cluster.

OpenStack

OpenStack is an open-source cloud computing platform that provides Infrastructure-as-a-Service (IaaS) for compute, storage, networking, and identity at scale. It is widely used by service providers, telecom operators, and large enterprises for private and managed clouds.

Protection of OpenStack VMs using Cohesity

OpenStack-native snapshot facilities (Nova, Cinder, Glance) are effective for short-term operational rollback, but they are not a complete enterprise data-protection solution. Cohesity protects OpenStack VMs through an integration with Storware: Storware acts as

the OpenStack-aware orchestrator and data mover, while Cohesity DataProtect provides the secure, scalable, policy-driven backup repository. In this solution, a Cohesity SmartFiles NFS 4.1 View is used as the Storware staging space and backup destination.

> **Note:** OpenStack protection is delivered through a third-party integration with Storware. Storware-specific configuration is performed in Storware and documented in the Storware documentation; this guide focuses on the Cohesity-side actions and references Storware where appropriate.

OpenStack-native snapshots lack the enterprise data-protection characteristics that customers require. The integration addresses the following gaps:

Capability	Enterprise Need	Gap with OpenStack-Native Snapshots
Cyber resilience	Immutable, isolated, encrypted backup copies	Native snapshots do not adequately protect against ransomware and malicious data destruction.
Retention management	Long-term retention with compliance support	Snapshots are short-term and often stored on the same underlying storage as production.
TCO	Deduplication and compression to optimize storage and cost	Native snapshots add storage overhead and fragmented management.
Disaster recovery	Geographically distributed protection	Limited DR orchestration and cross-site recovery.
Long-term archival	Archival tiers and tape for compliance	No integrated support for archival workflows, tape export, or tiering.
Visibility and reporting	Centralized reporting, monitoring, chargeback	Native snapshots lack centralized governance and analytics.
Enterprise-wide protection	Unified protection across many workloads	Native snapshots are platform-specific, creating operational silos.

Storware uses a disaggregated architecture consisting of a Storware server (UI and API) and one or more Storware nodes that connect to the OpenStack source, move data, and run backup and restore operations. Storware does not provide its own storage; the nodes mount a Cohesity SmartFiles NFS 4.1 View and use it as the filesystem backup destination.

During backup, Storware nodes read VM data from OpenStack and write it to the Cohesity NFS View. During recovery, the nodes read the selected recovery point from the Cohesity View and restore the VM (or files) back to OpenStack. For NFS/filesystem backups,

Storware uses the same NFS share for both staging and backup, avoiding a separate staging copy.

Setup Protection with OpenStack and Cohesity

The steps involved in setting up protection with OpenStack and Cohesity are as follows:

- Plan and prepare for OpenStack VMs protection through an integration with Storware and Cohesity. For more information, see Plan and Prepare for OpenStack VMs Protection with Storware and Cohesity.
- Configure NFS 4.1 view in Cohesity. For more information, see Plan and Prepare for OpenStack VMs Protection .
- Backup OpenStack VMs with Storware and Cohesity. For more information, see Backup OpenStack VMs.
- Recover OpenStack with Storware and Cohesity. For more information, see Recover OpenStack VMs.

Plan and Prepare for OpenStack VMs Protection

This topic helps Cohesity users plan and prepare a joint Storware along with Cohesity DataProtect solution for OpenStack VM protection. It focuses on Cohesity-related preparation and links to Storware documentation for Storware-specific configuration.

Prerequisites

- Cohesity DataProtect 7.3 or later.
- Storware server and nodes deployed in the OpenStack environment to be protected; a Storware node is required in each Availability Zone that needs protection.
- An OpenStack user with admin-level access to all projects to be protected.
- Network connectivity from the Cohesity cluster to the Storware nodes on NFS ports 111 and 2049.

OpenStack protection through the Storware integration is supported in Cohesity DataProtect 7.3 and later. For more information, see Supported Software section.

The core components of OpenStack are:

Component	Role
Keystone	Identity service; authentication and authorization for users and services.

Component	Role
Nova	Compute service; manages virtual machines across hypervisors.
Cinder	Block storage service; persistent volumes for VMs.
Glance	Image repository service; VM image templates.
Neutron	Networking service; connectivity, IP addresses, and security groups.
Swift	Object storage service for unstructured data.
Ceilometer / Aodh	Telemetry and alerting.

Planning Inputs to Collect

Category	Information to Capture	Owner
OpenStack	Keystone URL, region, domains/projects, endpoint type, availability zones, target VMs, storage backend, network segments	OpenStack administrator
Storware	Storware server FQDN/IP, node FQDN/IP per required AZ, node placement, licensing, Internet/fallback API reachability	Storware administrator
Cohesity	Cluster/VIP/FQDN, SmartFiles View name, storage domain, NFS export path, allowed subnets, protection/security settings	Cohesity administrator.
Operational	Backup window, retention goals, restore targets, RPO/RTO expectations, network bandwidth, largest VM size, expected concurrency	Joint project team

Considerations

- Storware advises scheduling no more than 10 tasks per Storware node.
- Keypairs the Storware user cannot access are not protected or recovered; restore them manually under the same name.
- A Storware node is required in each Availability Zone that needs protection.
- Protection policies are independent chains, and a VM/instance can belong to only one policy.
- Storware retains one backup copy indefinitely as long as the source exists.
- For NFS (filesystem) backups, a separate staging directory is not required; staging and destination share the same NFS share.

- If the primary backup location task fails, secondary targets are not attempted.
- NFS recoveries are not instant because merge operations are required to hydrate images.
- Snapshots are garbage-collected once a day (configurable in Global Settings).

For the complete and current list of considerations, supported features, and limitations, see Storware OpenStack documentation.

Best Practices

- Back up the Storware internal database. Without it, backups can be lost if the Storware server node must be redeployed.
- When OpenStack uses an NFS storage backend, enable QCOW2 file support in the Cinder configuration so backup snapshots create QCOW2 files instead of RAW.
- Place backup storage on the same network or close to the Storware nodes for high throughput.
- Size the solution using the backup window, network bandwidth, storage type, and largest VM size; scale Storware nodes up or down as needed.

Configuration of NFS 4.1 View in Cohesity

1. Prepare the SmartFiles View
    1. In Cohesity DataProtect, create a SmartFiles View for use as the Storware backup target.
    2. Use Category: Backup Target; Type: General Backup; Read/Write Protocol: NFS 4.1; Performance Policy: Backup Target High.
    3. Use a clear naming convention such as OpenStack-<region>-Storware or OpenStack-<tenant>-BackupTarget.
    4. Record the NFS export path in the form <Cohesity_Cluster_Hostname_or_VIP>:/<NFS_View_Name>.

    For more information see, SmartFiles (File and Object Services).
2. Restrict NFS Access
    1. Create either a global or view-level subnet allowed list for the SmartFiles View.
    2. Allow only the Storware server/node subnets that require read/write access.
    3. Avoid broad 0.0.0.0/0 access except for temporary lab validation.
3. Validate Network Requirements

Flow	Ports/Protocol	Notes
Cohesity DataProtect NFS View <-> Storware Node	2049/111 TCP/UDP	Required for NFS backup destination access. Validate firewall and routing before Storware configuration.
Storware Node -> OpenStack services	OpenStack service endpoint ports	Refer to Storware OpenStack documentation for Keystone, Nova, Glance, Cinder, and Neutron requirements.
Storware Node -> Ceph monitors, if applicable	3300/tcp, 6789/tcp	Needed when using Ceph RBD-related capabilities described by Storware.

You can verify the readiness as follows:

- Cohesity SmartFiles NFS 4.1 View created and export path recorded.
- Subnet allowed list configured for Storware nodes only.
- Storware nodes can mount the Cohesity NFS export with read/write access.
- Backup window, retention, and expected concurrency agreed with stakeholders.
- Storware OpenStack strategy selected (disk attachment with generic incremental is the preferred strategy).
- OpenStack access model and authentication domain approach agreed with the OpenStack administrator.
- Restore test plan includes both full VM recovery and file-level or mounted-backup validation where applicable.

For more information on software preparation, refer to the following:

- Storware OpenStack provider setup, authentication domains, supported strategies, and limitations
- Storware deployment methods
- Storware server/node component requirements

Backup OpenStack VMs

Storware initiates and orchestrates OpenStack VM backups and writes backup data to the Cohesity SmartFiles NFS 4.1 View mounted on the Storware nodes.

Backup Architecture

- Storware initiates and orchestrates OpenStack VM backup activities.

- Cohesity SmartFiles NFS 4.1 View is mounted on Storware nodes and used as the filesystem backup destination.
- For NFS/filesystem-based backups, the Storware can use the same NFS share for staging and backup storage, avoiding a separate staging directory.

Cohesity Backup Target Setup

Complete these Cohesity-specific steps before creating or running Storware backup policies:

1. Create the SmartFiles View as Backup Target > General Backup, with NFS 4.1 and Backup Target High performance policy.
2. Apply a subnet allowed list that allows the Storware node subnet(s) to access the view.
3. Share the NFS export path with the Storware administrator.
4. Confirm the Storware node OS user/group ownership requirement is met on the mount point before backup execution.

Reference Mount Command

Use the following as the reference, after replacing the placeholders with values specific to your environment.

```
mount -t nfs -o rw,re-
latime,ver-
s=4.1,-
rsize-
=1048576,w-
wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys
<Cohesity_Cluster_Hostname_or_VIP>:/<NFS_View_Name> <View_Mount_Path>
```

For persistence across reboot, add the equivalent mount line to /etc/fstab after validating the mount manually.

Configure Backup in Storware

Perform the following in the Storware UI. Refer to the Storware documentation for detailed steps:

1. Upload the Storware license and run the configuration wizard; verify the Storware nodes are visible.
2. Choose OpenStack as the hypervisor and register the source with the Keystone API URL, region, endpoint type, and authentication domain settings.
3. Select Backup Destination type File System and enter the mount point where the Cohesity SmartFiles View is mounted.

4. Create a backup policy: select projects/instances, configure retention and schedule (each schedule is an independent snapshot chain), and target the configured nodes.
5. Optionally test a small instance backup before broad rollout, and adjust the maximum tasks per node if needed.

Cohesity Validation Before Running Backups

Validation	Expected Result
NFS mount from Storware node	Mount succeeds and is read/write.
File creation test on mounted path	Test file can be created and removed by the vprotect user/group.
Cohesity View access policy	Only intended Storware subnets are allowed.
Performance baseline	Initial small VM backup completes without NFS errors before broad policy rollout.
Capacity monitoring	SmartFiles capacity and growth are monitored during initial full backups.

Operational Guidance

- Start with a small VM backup to validate end-to-end connectivity, permissions, and backup target behavior.
- Scale concurrency carefully. Storware advises scheduling no more than 10 tasks per Storware node.
- If node count or load is too high, stagger full backups across multiple schedules and policies to reduce peak throughput demand.
- For NFS-based backups, reads from the source are limited by the write speed of the backup destination, because staging and destination share the same NFS share.
- Monitor backup target utilization, throughput, and errors from both Cohesity and Storware during the first production backup cycles.

Verification list for backup completion

- Initial test backup completed successfully.
- Backup data appears on the Cohesity SmartFiles View.
- Backup policy retention and cadence match the agreed protection plan.
- At least one restore validation is scheduled after the first successful backup.
- Storware internal database backup is included in the operational plan (backups can be lost if the Storware server node must be redeployed without its database).

For more information on Storware backup, refer to the following:

- Configure OpenStack provider, strategy, authentication domains, inventory sync, and backup behavior
- Configure filesystem backup destination in Storware

Recover OpenStack VMs

Storware performs the OpenStack-aware recovery workflow and recreates the VM in OpenStack based on the selected recovery point and restore options. Cohesity provides the backup repository through the SmartFiles NFS View that contains the Storware backup data.

Prerequisites

- Verify the Cohesity cluster and SmartFiles View are online and reachable from the Storware node.
- Confirm NFS mount access using the same mount point used during backup operations.
- Confirm the View still permits read/write access from the Storware node subnet.
- Validate sufficient capacity and performance headroom for recovery operations.
- For isolated or alternate-environment tests, confirm routing and DNS/IP reachability to the Cohesity cluster/VIP.

To recover OpenStack:

1. Confirm the Cohesity SmartFiles View used as the backup destination is accessible from the Storware node.
2. In Storware, select the VM recovery point and recovery mode (recovery to virtualization provider, VM instance mount, or file-level recovery) per the Storware documentation.
3. During recovery, monitor Cohesity-side NFS availability, throughput, and any View access errors.
4. After Storware reports completion, validate from OpenStack that the VM was recreated and is reachable according to the restore objective.
5. Document the recovered VM name, target project, target flavor, target network, recovery point used, and any deviations from the original VM configuration.

Recovery Options in the Storeware UI:

Option	Description
Restore to original location	Overwrites the existing VM.

Option	Description
Restore to new location	Creates a new VM with a different name.
Target project	OpenStack project for the restored VM.
Target flavor	VM size; can differ from the original.
Target network	Network to attach the restored VM to.
VM instance mount / file-level recovery	Mount a backup to a Storware node to browse, download, or restore individual files.

For more information, see Storware documentation.

You can validate the recovery completion using the following:

Validation Area	What to Confirm
Cohesity repository	SmartFiles View is mounted and readable/writable from the Storware node.
Restore job	Storware restore job completes without repository or NFS errors.
OpenStack object	Recovered VM is present in the intended project.
Compute/network	Target flavor and network are as planned.
Access	Required keypair/access key handling is completed per Storware/OpenStack requirements.
Application validation	Application owner confirms workload-level readiness where applicable.
Post-restore protection	Recovered VM is added to the correct protection policy if ongoing protection is required.

References

- Storware OpenStack
- Storware deployment
- Storware component requirements
- Storware filesystem backup destination

Chapter 13: Databases

Cohesity supports the backup and recovery of the following databases:

- Amazon RDS and Amazon Aurora 2786
- Cassandra 2795
- Couchbase Using Connector Agent 2843
- MongoDB 2857
- MongoDB via Ops Manager 2878
- Pinecone 2896
- Microsoft SQL Server 2906
- Oracle Database 3022
- CockroachDB 3146
- Epic Database Protection using Cohesity Physical Agents 3164
- IBM DB2 3168
- MySQL EE 3195
- Percona MySQL on Linux 3213
- SAP HANA 3232
- SAP Sybase IQ 3264
- SAP MaxDB on Linux 3286
- SAP MaxDB on Windows 3309
- SAP Sybase ASE on Linux 3331
- SAP Sybase ASE on Microsoft Windows 3352
- SAP on Oracle for Linux 3372
- SAP on Oracle for Microsoft Windows 3399
- PostgreSQL 3419

Amazon RDS and Amazon Aurora

You can use a Protection Group to back up Amazon Relational Database Service (Amazon RDS) database instances and Amazon Aurora. Supported database engines include PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. Cohesity protects the data using RDS snapshots (no data is stored on the cluster).

Prerequisite

For replicating encrypted RDS snapshots from one location to another:

- The KMS key Aliases used to encrypt the RDS snapshots in the source location must be available in the target location.
- The IAM user/role used to register the target source must have access to the KMS key at the target location.

Considerations

- The backup and recovery of Amazon RDS and Amazon Aurora at instance level (snapshot-based protection) is supported on physical clusters, virtual edition clusters, and Cohesity Cloud Editions.
- The backup and recovery of Amazon RDS and Amazon Aurora at database level (ingest-based protection) is supported only on Cohesity Cloud Edition on AWS.
- Cohesity uses the Amazon RDS API to take snapshots according to the scheduled protection policy.
- Cohesity supports automated and manual RDS snapshots.
- You can restore the backups to the original location or a new location.
- Cohesity does not support recovery to different accounts with different regions but recovery to different accounts with the same region, the same account with the same region, and the same account with different regions are supported.
- Point in time recovery is supported only if you enable automated backups on the RDS instance. You can select a time and restore the database instance to the time selected. Automated backups occur once daily at the time selected. You can restore to any point in time within your backup retention period.
- Amazon Aurora Serverless backups are not supported.
- Recovery of an RDS instance encrypted with the default CMK (Customer Managed Key) to a different account is not supported as AWS does not support sharing a snapshot that is encrypted using the default KMS key of the AWS account.
- You cannot recover an Amazon RDS database to the same or alternate location if the name of the database you are restoring and at the target location are the same. AWS

does not allow creating more than one RDS database instance with the same name.

Workaround: Specify a unique name for the RDS database instance and then perform a recovery.

- The following table lists the service endpoints required for the backup and recovery of AWS RDS databases:

AWS Resource Type	Service Endpoints
IAM	iam.amazonaws.com
KMS	kms.<region>.amazonaws.com
RDS	rds.<region>.amazonaws.com

- If you are protecting Amazon RDS and Aurora instances within a private VPC, then you must:
  - Register the AWS cloud source using an IAM role. For more information, see Register AWS Source.
  - Ensure the following endpoints are open.
    -

AWS Resource Type	Service Endpoints
STS	sts.<region>.amazonaws.com
RDS	rds.<region>.amazonaws.com
KMS	kms.<region>.amazonaws.com

  - Enable AWS Security Token Service (STS) in the default region where the Cohesity cluster is deployed.

To learn more, see the following topics:

- Manage Protection Policy
- Add or Edit a Protection Group for Amazon RDS or Amazon Aurora

Add or Edit a Protection Group for Amazon RDS or Amazon Aurora

Cohesity supports backing up Amazon Relational Database Service (RDS) databases and Amazon Aurora.

For information about required permissions, see Minimum Permissions for Individual Workloads and Workflows.

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit a Hypervisor Source
- Create or Edit a Standard Policy
- Create or Edit Storage Domains

To create a Protection Group for Amazon RDS:

1. Navigate to **Data Protection > Protection** and do one of the following:
   - To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
   - To edit a Protection Group—Click the action icon ( ) next to the Protection Group and select **Edit**.

   If creating a Protection Group for Amazon RDS, select **Databases > Amazon RDS**.

   If creating a Protection Group for Amazon Aurora, select **Databases > Amazon Aurora**.
2. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
3. **Source:** Select the database you want to protect. You can also register a new source with the Cohesity cluster so you can protect it.
1. Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.
2. Use the filtering options to tailor the list of servers. You can also filter the objects based on the tags by clicking the **Tags** icon and selecting the required tag from the **Tags** drop-down list.
3. Click the arrow adjacent to the Object (**Expand All**) to expand the entire Object hierarchy. Click the arrow again (**Collapse All**) to collapse the entire Object hierarchy.
4. The graphical representations in the table are described below.

   When creating a new Protection Group, a shield next to the Object name indicates the Object is already been protected by another Protection Group.

   Click the icon to see details about the object.

If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.

5. Click **Save Selection**.
6. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.

   > **Note:** For Amazon RDS and Amazon Aurora snapshot protection, the actual snapshot remains on AWS and metada about the snapshot is stored in the cohesity cluster. The data lock policy applies to the metadata stored within Cohesity. Changes made from the AWS end for the snapshots cannot be controlled by Cohesity

   The selected Storage Domain will contain the new View where the script results are stored.
7. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

   The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
8. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
9. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
10. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

For more information, see Alerts.

11. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
12. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

    > **Note:**
    > The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

13. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
14. **Description:** Optional. Click the plus sign and enter a description for the Protection Group.
15. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Recover Amazon RDS or Amazon Aurora

Cohesity provides the ability to recover the Amazon RDS database or Amazon Aurora to its original location or an alternate location. For more information, see:

- Recover Amazon RDS or Amazon Aurora to the Original location
- Recover Amazon RDS or Amazon Aurora to an Alternate Location

Recover Amazon RDS or Amazon Aurora to the Original location

To recover the Amazon RDS or Amazon Aurora to the original location:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Databases > Amazon RDS**.

   The Amazon RDS page appears.
3. Search for the database you want to recover by entering the characters of the database name or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select the database you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date for the recovery point.
     b. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
     c. Select the snapshot to be recovered by clicking the corresponding blue dot on the timeline bar.

        If multiple snapshots are available, then you can choose any point in timeline (Point-in-Time snapshot) between the two snapshots either by adjusting the timeline bar or by selecting the time from the timer field. If you have selected an invalid time, then the nearest valid snapshot will be selected for recovery.
     d. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the List mode:
     a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page.
     b. By default, the page displays the snapshots available in the last one month.

c. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

d. Select the snapshot to be recovered.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the Amazon RDS to the original location.
8. Provide the following **Network Settings** information:
    1. **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered RDS.
    2. **Subnet**: From the drop-down list, select a subnet in the Amazon VPC to store the recovered Amazon RDS.
    3. **Availability Zone**: From the drop-down list, select an availability zone in AWS to recover the RDS.
9. Configure the following **Additional Settings**:

    The **Database Engine** field auto-populates the database engine name used for this instance.

    1. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.
    2. **DB Port**: Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port.
    3. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
    4. **DB Parameter Group**: Select the database parameter group to associate with the DB instance.
    5. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
    6. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.

7. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
8. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
9. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.

10. Optional. Change the default name of the recovery task in the **Task Name** field.
11. Click **Start Recovery**.

The recovery process is initiated. You can monitor the recovery task from the Recoveries page.

Recover Amazon RDS or Amazon Aurora to an Alternate Location

To recover the Amazon RDS or Amazon Aurora to an alternate location:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Databases > Amazon RDS**.

   The Amazon RDS page appears.
3. Search for the database you want to recover by entering the characters of the database name or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select the database you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date for the recovery point.
     b. The snapshots available on the selected date are displayed in blue dots on the timeline bar.

c. Select the snapshot to be recovered by clicking the corresponding blue dot on the timeline bar.

   If multiple snapshots are available, then you can choose any point in timeline (Point-in-Time snapshot) between the two snapshots either by adjusting the timeline bar or by selecting the time from the timer field. If you have selected an invalid time, then the nearest valid snapshot will be selected for recovery.

d. Click **Select Recovery Point** and then go to step 6.

- To select a snapshot from the List mode:

  a. Select the **List** tab from the top-right corner of the Edit recovery point for <VM_name> page.

  b. By default, the page displays the snapshots available in the last one month.

  c. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

  d. Select the snapshot to be recovered.

  e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** to recover the Amazon RDS to an alternate Amazon RDS instance
8. From the **Source** drop-down menu, select a registered RDS instance to which you want to restore the Amazon RDS . You can also register a new source by clicking **Register Source** from the **Source** drop-down menu.
9. From the **Region** drop-down menu, select a region to restore the Amazon RDS.
10. Provide the following **Network Settings** information:
    1. **VPC**: Select an Amazon Virtual Private Cloud (VPC) to store the recovered RDS.
    2. **Subnet**: From the drop-down list, select a subnet in the Amazon VPC to store the recovered Amazon RDS.
    3. **Availability Zone**: From the drop-down list, select an availability zone in AWS to recover the RDS.
11. Configure the following **Additional Settings**:

    The **Database Engine** field auto-populates the database engine name used for this instance.

    1. **Database Instance Identifier**: Specify the unique key that identifies the database instance that will be recovered.

2. **DB Port**: Specify the TCP/IP port that the DB instance will use for application connections. The connection string of any application connecting to the DB instance must specify the port number of the DB instance. Both the security group applied to the instance and your company's firewalls must allow connections to this port.
3. **DB Option Group**: Select an option group that contains the option you want to attach to the DB instance that will be recovered. If there are not any option groups compatible with the selected engine, a default group will be created at launch.
4. **DB Parameter Group**: Select the database parameter group to associate with the DB instance.
5. **Multi A-Z Deployment**: Enable this option if you want the DB instance to be recovered to have a standby instance deployed in another availability zone. This option is disabled by default.
6. **IAM DB Authentication**: Enable this option if you want to manage your database user credentials through AWS IAM users and roles. This option is disabled by default.
7. **Public Accessibility**: Enable this option if you want the DB instance to also have a public IP address in addition to the private IP address. This option is disabled by default.
8. **Copy Tags To Snapshots**: Enable this option for copying tags to snapshots. This option is disabled by default.
9. **Auto Minor Version Upgrade**: Enable this option if you want the DB instance to automatically upgrade when a new minor database engine version is available. This option is disabled by default.

12. Optional. Change the default name of the recovery task in the **Task Name** field.
13. Click **Start Recovery**.

The recovery process is initiated. You can monitor the recovery task from the Recoveries page.

Cassandra

27 July 2026

The Cohesity Cassandra DataProtect solution natively integrates with Apache Cassandra and Datastax Enterprise to provide consistent backup and recovery for clustered Cassandra deployments. See the following topics for details:

Based on the policy settings, the Cohesity NoSQL and Hadoop Service initiates a backup of the Cassandra Database. The initial protection run/first backup is a full-backup and everything after that is an incremental backup only. The service automatically determines the number of backup channels, based on the compute resources on Cohesity and the

topology of the Cassandra cluster. The necessary Cohesity Views are automatically generated where the backups are going to be written. Communication between Cohesity and the Cassandra Databases is managed by the NoSQL and Hadoop Service as described in the section below.

Core components

The following lists the core components of the Cohesity Cassandra solution:

- Physical or virtual Cassandra instance.
- Cohesity Data Cloud.
- Cohesity Compute Nodes.
- Cohesity NoSQL and Hadoop Service running on Cohesity Compute Nodes.
- Optional. Archive or tiering target at your data center or cloud provider of choice.

As illustrated in the figure below, you can integrate Cohesity with an existing Cassandra infrastructure by installing the Cohesity NoSQL and Hadoop Service on Compute Nodes. Once the service is installed and running, you can create protection policies, and Protection Groups for the Cassandra database. Once the database is protected, it can be archived, sent to tape, and or replicated to another instance running in the cloud or on-premises as shown in the image below.

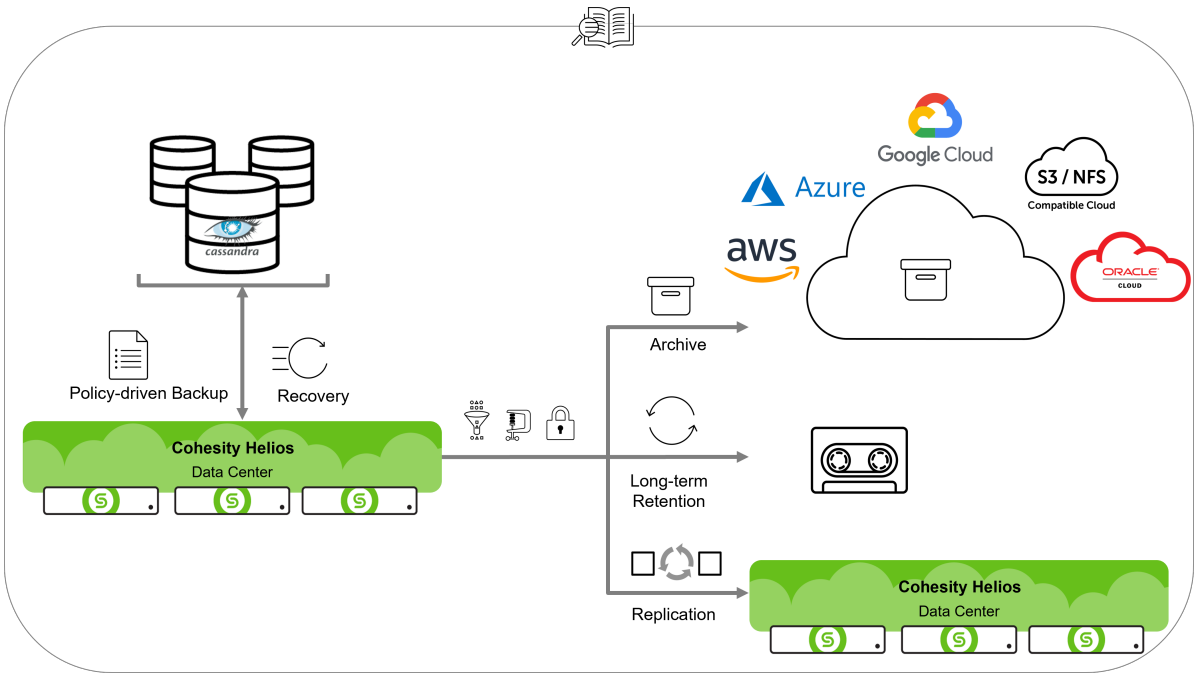

Cassandra Features

Following are the key Cassandra-specific features of the solution:

- Incremental-forever backups - after the first full backup, all subsequent backups are incremental only forever
- Granular backup and recovery - backup and recover tables or keyspaces
- Schema & data backed up and restored
- Multi-DC aware - option to back up data from a local data center
- Optimized restore - much faster recovery if Cassandra cluster topology is unchanged
- Backup & Recover DSE Solr indexes
- Handle topology changes seamlessly (node addition, removal, IP changes, etc.)
- Support Kerberos and LDAP authenticated clusters
- Support SSL-enabled clusters
- Change schema on restore

Plan and Prepare for Cassandra Protection

27 July 2026

The following topics provide guidance that helps you get ready to implement the Cassandra DataProtect Solution:

Supported Versions

For more information, see NoSQL and Hadoop Platforms.

> **Note:**
>
> Users whose existing Cassandra setups match any of the following criteria must edit their Cassandra source and go through the registration process.
>
> - Cassandra cluster set up with different public and private IP addresses(rpc_address/broadcast_rpc_address and different listen_address)
> - Cassandra setup with different listen_address and broadcast_rpc_address (listen_address = address1, rpc_address = 0.0.0.0, broadcast_rpc_address = address2)

Prerequisites

Make sure the following prerequisites are met before backing up Cassandra databases:

- The NoSQL and Hadoop Service is in Running status. See NoSQL and Hadoop Service for details.

- Refer to the NoSQL and Hadoop Service - Post Installation Procedures to get a general idea about the following operations:
  - Manage SSL enabled NoSQL primary clusters
  - Work with Kerberos Security enabled primaries for data protection
- Linux User on Cassandra Primary
- SSL Enabled Cassandra Clusters
  - Steps Required to Support SSL Enabled Cassandra Clusters
  - Obtain SSL Certificates in Cassandra Nodes and Import them on the Cohesity cluster that are Running the NoSQL and Hadoop Services
  - Import NoSQL and Hadoop Services Certificates
- Setting the Number of Open Files Limit for Cassandra
- Enabling JMX

Linux User on Cassandra Primary

For adding a Cassandra data source, either an SSH username and password or a private key is needed. SSH keys can be used to establish a secure connection. For more details on the minimum permissions needed for the user, see Minimum Permissions.

> **Note:** You need to use the same Linux user & password or SSH key on all Cassandra nodes that belong to the same source/cluster.

SSL Enabled Cassandra Clusters

Steps Required to Support SSL Enabled Cassandra Clusters

The `cassandra.yaml` file is the main configuration file for Cassandra. Check the following SSL options in the cassandra yaml file:

- `client_encryption_options` (Client-to-node encryption)
  - `enabled`: By default, the enabled option is set to `False`. If the enabled option is set to True, then you must Setting up truststore:.
  - `require_client_auth`: By default, the require_client_auth is set to `False`. If the `require_client_auth` is set to True, then you must Setting up keystore and import them in all Cassandra nodes, and restart Cassandra services for the changes to take effect
- `server_encryption_options` (Node-to-node encryption)
  - `internode_encryption`: If this option is set to `all`, dc or rack, then you must copy Cassandra certificates into NoSQL and Hadoop Services TrustStore using `imanis_cli.sh -c ssl`.

    - `require_client_auth`: By default, the `require_client_auth` is set to False. If set to true, then you must copy NoSQL and Hadoop Services SSL certificates into every Cassandra nodes TrustStore and restart Cassandra node for the changes to take effect.
- Additionally, if encryption is enabled for JMX communication, the truststore and keystore parameters are specified in cassandra-env.sh using the following JVM options:
    - `javax.net.ssl.keyStore`
    - `javax.net.ssl.keyStorePassword`
    - `javax.net.ssl.trustStore`
    - `javax.net.ssl.trustStorePassword`

    Use the values from these options to create SSL certificates and import them into the NoSQL and Hadoop Service truststore, following the same process as client_encryption_options or server_encryption_options.

Obtain SSL Certificates in Cassandra Nodes and Import them on the Cohesity cluster that are Running the NoSQL and Hadoop Services

This section describes the procedure of creating SSL certificates in Cassandra nodes and importing them in NoSQL and Hadoop Services cluster. However, if you already have SSL certificate (*.cer) files, you can start from Step #4.

1. Check the location of Truststore and password in Cassandra Primary. You must execute the following command from the directory which contains the cassandra.yaml file.

```
[root@nosql13 conf]# grep "truststore:" cassandra.yaml
truststore:  <trustore location1>/cassandra.truststore
truststore: <trustore location2>/cassandra.truststore
[root@nosql13 conf]#
[root@ nosql13 certs]# grep password <configuration file location>/cassandra.yaml
# - PasswordAuthenticator relies on username/password pairs to authenticate users. It keeps usernames and hashed passwords in system_auth.credentials table.


# The passwords used in these options must match the passwords used when generating
keystore_password: nosql2016
truststore_password: nosql2016
```

```
keystore_password: nosql2016


# Set trustore and truststore_password if require_client_auth is true
truststore_password: nosql2016
```

2. Go to truststore location from the grep command (e.g.: /home/cassandra/certs/) and create SSL certificates for both the client and server encryption options using the password from the grep command (e.g.: nosql2016). The "/root/jdk1.7.0_45" is the Java home directory in Cassandra nodes.

```
[root@nosql13 certs]# /root/jdk1.7.0_45/jre/bin/keytool -export -alias
nosql13 -file nosql13new.cer -keystore cassandra.keystore


Enter keystore password:


Certificate stored in file <nosql13new.cer>
```

3. Repeat Step #2 in all the Cassandra nodes:

```
[root@nosql14 certs]# /root/jdk1.7.0_45/jre/bin/keytool -export -alias
nosql14 -file nosql14new.cer -keystore cassandra.keystore


Enter keystore password:


Certificate stored in file <nosql14new.cer>


[root@nosql14 certs]#
```

   Where /root/jdk1.7.0_45 is the Java home directory in Cassandra nodes.

4. Import all the SSL certificates. For more details, refer Setting up truststore:.

Import NoSQL and Hadoop Services Certificates

1. For more information on creating the SSL certificates, refer Setting up keystore.
2. Import NoSQL and Hadoop Services certificates on each of the Cassandra nodes. Ensure that this step is repeated for all the remaining Cassandra nodes.

```
[root@nosql13 certs]# keytool -import -v -trustcacerts -alias nosql1 -
file nosql1.cer -keystore .truststore
Enter keystore password:
Owner: CN=nosql1, OU=, O=, L=, ST=, C=
```

```
Issuer: CN=nosql1, OU=, O=, L=, ST=, C=

Serial number: 203a2356

Valid from: Tue Jul 11 19:14:33 IST 2017 until: Mon Oct 09 19:14:33 IST
2017

Certificate fingerprints:

MD5: 17:4B:40:00:DA:A0:38:82:3A:88:37:4D:BE:8F:92:C7

SHA1: A7:C1:5E:B8:A3:7F:49:C3:C8:66:76:51:7B:E3:28:D4:F5:D7:62:C9

SHA256:
FC:C3:99:69:9F:EC:40:EC:8B:E3:DD:FD:68:3E:30:01:66:F7:E8:DA:97:6F:71:1
3:19:1B:35:6D:38:4B:30:77

Signature algorithm name: SHA256withRSA

Version: 3

Extensions:

#1: ObjectId: 2.5.29.14 Criticality=false

SubjectKeyIdentifier [

KeyIdentifier [

0000: BA 16 CB 9A CE DD 71 8B 98 4F 5F 75 39 CB F3 65 ......q..O_u9..e

0010: BF 8A B7 FF ....

]

]

Trust this certificate? [no]: yes

Certificate was added to keystore

[Storing .truststore]

[root@nosql13 certs]# keytool -import -v -trustcacerts -alias nosql2 -
file nosql2.cer -keystore .truststore

Enter keystore password:

Owner: CN=nosql2, OU=, O=, L=, ST=, C=

Issuer: CN=nosql2, OU=, O=, L=, ST=, C=

Serial number: 456a0a54

Valid from: Tue Jul 11 19:14:31 IST 2017 until: Mon Oct 09 19:14:31 IST
2017

Certificate fingerprints:

MD5: 0E:36:CD:50:86:EE:47:BB:17:EA:E3:B9:60:9F:A7:C2
```

```
SHA1: 37:FD:A0:BF:B1:E3:15:1F:99:6B:AC:7D:91:66:43:65:85:0C:41:AC
SHA256:
8E:1F:11:71:1C:EC:BA:52:F1:42:CE:53:3C:63:C4:47:09:71:23:13:21:B0:E0:4
B:AC:CE:73:44:D5:B8:77:7D
Signature algorithm name: SHA256withRSA
Version: 3
Extensions:
#1: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 41 ED 3E 7D 10 56 62 01 A7 A1 21 5F F6 D9 7C D0 A.>..Vb...!_....
0010: D0 2E 03 34 ...4
]
]
Trust this certificate? [no]: yes
Certificate was added to keystore
[Storing .truststore]
```

> **Note:** Ensure that permissions of Cassandra's truststore file are given to the Cassandra user.

Setting the Number of Open Files Limit for Cassandra

The default setting for the number of open files in Cassandra primary is 1024. Please set the number of open files to 100,000 by running the following command on the Cassandra primary.

```
ulimit -n 100000
```

> **Note:** Without this step you may face an error on the Primary cluster while executing Cassandra restore.

Enabling JMX

Enable remote JMX connection on all the Cassandra nodes before you add Cassandra as a source. Cassandra DataProtect Solution connects to the primary using the remote JMX connection so a local JMX cannot be used in this case.

> **Note:** You need to use the same JMX credentials on all Cassandra nodes that belong to the same source/cluster.

To enable JMX:

1. Enable the Remote JMX connection: For Cassandra 2.0.14 (2.0 branch) and Cassandra 2.1.4 (2.1 branch) onwards.
2. Access the cassandra-env.sh from the Cassandra configuration directory on all the Cassandra nodes and set ["$LOCAL_JMX" = "yes"] entry to "no".
3. For the default DSE installation, the preceding file is available in /etc/dse/cassandra.
4. For a successful data source discovery in DSE 6.7.x, the following parameter must be set in /etc/dse/cassandra/cassandra-env.sh.

   "JVM_OPTS"="JVM_OPTS -Djava.rmi.server.hostname=<ip>"
5. Restart Cassandra services.

> **Note:** If JMX communication is SSL-encrypted, see Steps Required to Support SSL Enabled Cassandra Clusters for information on importing the corresponding SSL certificates.

Considerations

- For users who have their Cassandra cluster set up with **public address = rpc_address/broadcast_rpc_address** and **private address = listen_address**, for DSE below 6.x (5.x and below), a bulk restore will not work if the private IPs for the Cassandra nodes are not accessible by the Cohesity cluster (SCP restore will work as expected).
- Users must refresh a data source when configuration settings of the data source are modified. For example, if changes are made to cassandra.yaml, cassandra-env.sh, or dse.yaml file of a data source, then the data source must be re-discovered and re-verified.
- Backup is not supported for the following system keyspaces:
  - system
  - system_auth
  - system_distributed
  - system_schema
  - system_traces
  - dse_leases

- dse_perf
- dse_security
- dse_system
- solr_admin
- dse_insights_local
- dse_insights
- dse_system_local
- system_virtual_schema
- system_views
- system_backups

> **Note:** The system_auth and system_schema keyspaces are not protected as part of the regular auto-protection job. Although these keyspaces may appear as protected in the **Protected Objects** and **All Objects** tabs on the **Sources** page when auto-protect is enabled, protection of system keyspaces is not supported as part of regular protection jobs. To back up these keyspaces, create a separate protection job by enabling the **Backup System Keyspaces** option on the **New Protection** page.

- Restore of system keyspaces such as OpsCenter is only supported for the same cluster from which it was backed up.
- Backup and recovery of system keyspaces and of keyspaces using SimpleStrategy are not supported in hybrid Cassandra environments, because these keyspaces require access to all data centers for consistent backup and restore operations.
- Backup and recovery of tables containing semicolon (';') in the column name is not supported.
- Data recovery from Cassandra 2.x to 3.x is not supported if the backed-up tables contain dropped columns.
- For Cassandra 2.x or DSE 4.x, backup recovery is not supported for tables where all columns are part of the Primary Key.
- Appending suffix is not supported while restoring graph objects. This is because individual graph keyspaces cannot be renamed. This issue occurs during the jobtag recovery process when the global suffix is applied to all keyspaces thus wrongly renaming the '_system' keyspace for a graph. For example, we need to back up two types of keyspaces for a graph. If the name of the graph is ‘exampleGraph’, then we should back up ‘exampleGraph’ and ‘exampleGraph_system’. However, during jobtag

recovery, if global suffix is applied to all keyspaces it wrongly renames the keyspaces 'exampleGraph_suffix' and 'exampleGraph_system_suffix' for a graph.

- Support for non-default store_type, algorithm, and protocol in cassandra.yaml is not available in this release.
- **For DataStax Enterprise 6.x, the following considerations are applicable:
    a. To restore DSE 5.x backups to DSE 6.x cluster, the minimum version must be DSE 6.0.6.
    b. If a protection run gets killed or the primary Cassandra node goes down during a restore, the temporary directories on the primary nodes are not deleted. In such cases, the user has to manually delete the temporary directories on the primary Cassandra nodes.

       The directories are created in each of the storage directories on the primary nodes and are named by the restore protection run's UUID.
- In DSE 6.x, restoring a Cassandra table where the primary key contains UDT would fail if the corresponding keyspace object had been renamed.
- Apache Cassandra 4.0 does not support the experimental features.
- Data restore from Cassandra 3.x to 4.0 is not supported.
- Review and understand the following considerations of **Restrict access to specific Objects**:
    - For Log backups, **Restricted access to specific Objects** must be provided at the source level. Log backups are necessary for Cassandra PITR.
    - For recovery, to select the specific data centers, **Restricted access to specific Objects** must be provided at the source level. If access is not provided or if the user does not select the data center for recovery the data is bulk loaded to all data centers.
    - **Restricted access to specific Objects** on replicated clusters is not supported.

> **Note: Restrict access to specific Objects** for NoSQL and Hadoop Service is an Early Access feature. Contact your Cohesity account team to enable the feature.

- Supports IPv6 on DataStax Enterprise 6.x only.
- Cohesity does not support registering a Cassandra cluster using a mix of IPv4 and IPv6 addresses.
- For IPv6, Kerberos authentication is not supported.
- During the recovery, roles are not deleted or updated.

- Backup and restore of role passwords are not supported.
- During the recovery, If there is already a role with the same name on the target, even if it has different properties, the permissions will still be applied to that role.

Firewall Ports

The following table lists the required firewall ports that are to be open to enable communication between Cohesity and Cassandra DataProtect Solution.

Port	Source	Target	Direction (From Node)	Network Protocol	Usage Notes	Type of Traffic
22	Cohesity cluster	Cassandra nodes	Outgoing	TCP	Cassandra backup Cassandra optimized restore/DSE 6.x restore	Backup and restore
7000	Cohesity cluster	Cassandra nodes	Outgoing	TCP	Cassandra bulk restore	Restore
7001	Cohesity cluster	Cassandra nodes	Outgoing	TCP	Cassandra bulk restore incase of ssl enabled cluster	Restore
7199	Cohesity cluster	Cassandra nodes	Outgoing	TCP	Cassandra repository registration Cassandra backup (snapshot operation) Fetching Cassandra configurations (list of nodes, token map etc)	Management
8983	Cohesity cluster	Cassandra nodes	Outgoing	TCP	Cassandra solr index management	Management

Port	Source	Target	Direction (From Node)	Network Protocol	Usage Notes	Type of Traffic
9042	Cohesity cluster	Cassandra nodes	Outgoing	TCP	Cassandra backup Cassandra restore Cassandra source registration	Management
9160	Cohesity cluster	Cassandra nodes	Outgoing	TCP	Cassandra bulk restore (DSE 4.x / Cassandra 2.x)	Management

Minimum Permissions

Authentication Enabled Cassandra Cluster

In case of authentication enabled Cassandra cluster, the specified cassandra user must be a super user.

Discovery:

ssh user should have read/execute access to Cassandra config directory.

Backup:

ssh user should have read/execute access to Cassandra data directories.

Restore:

For optimized recovery:

- The ssh user must have read/write/execute access to Cassandra data directories.
- The ssh user must have privileges to change the file ownership to Cassandra user.

Additionally, for DSE 6.x and restore with staging dir option selection, ssh user should have permission to execute sstableloader cmd on the Cassandra nodes. This means that sstableloader cmd should be in the PATH of ssh user, and the ssh user should have write access to the following directories:

- commitlog directory
- cdc_raw directory (if this dir is missing, need to create an empty dir)
- data directory
- saved_caches directory

Write-access to these directories are DSE requirements since the sstableloader tool checks for the same.

Additionally, if password encryption is enabled through client_encryption_options and server_encryption_options in cassandra.yaml, the SSH user must have read access to the system_key file (defined by the config_encryption_key_name setting in dse.yaml), as it contains the encryption key.

Register a Cassandra Source

27 July 2026

To configure Cassandra for data protection on Cohesity, you need to register it as a source.

Before you register Cassandra as a source, ensure that you have reviewed and completed the instructions provided in the Plan and Prepare for Cassandra Protection topic.

To register Cassandra as a source:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** on the upper-right of the page and then select **Databases > Cassandra**. The **Register Cassandra** dialog is displayed.
3. **Primary Host**: Specify any one seed node of the Cassandra cluster.
4. **customName**: Specify a generic source name of the Cassandra cluster.

> **Note:** You can register two sources with the same IP address on the Cohesity cluster by adding a different custom name.

> **Note:** This is an Early Access feature. Contact your Cohesity account team

> to enable the feature.

5. **Cassandra Configuration Directory**: Specify the directory path for the Cassandra configuration YAML file.
6. **Password**: Select **Password** if you want it as the authentication type.
    1. **Username**: Specify the username.
    2. **Password**: Specify the password.
7. **Private Keys**: Select **Private Keys** if you want it as the authentication type.
    1. **Login**: Specify the username.
    2. **SSH Private Key**: Specify the private key.
    3. **Passphrase (Optional)**: Specify the password.
8. **Cassandra Distribution**: Select the required Cassandra distribution:
    - **DSE**: Select **DSE** to use DataStax Enterprise as the Cassandra distribution.
        a. **DSE Configuration Directory**: Specify the directory path for DSE-specific configuration.
        b. **DSE Tiered Storage**: Select this check box if the cluster has DSE tiered storage.
        c. **DSE Authenticator**: Select this check box if the cluster has a DSE Authenticator.
        d. **DSE Solr**: Select this check box if the cluster has DSE Solr.
            i. **Solr Nodes**: Enter all the Solr node IP addresses
            ii. **Solr Ports**: Enter the Solr node Port.
    - **Apache Cassandra**: Select **Apache Cassandra** to use Apache Cassandra as the Cassandra distribution.
9. Click **Register**.
10. Enter the details for the additional authentication and configurations as needed.

> **Note:** For users who have their Cassandra cluster set up with different public and private IP addresses (different listen_address and rpc_address/broadcast_rpc_address), when registering/updating a Cassandra source, the full list of seeds may not be displayed on the verification screen. Instead, the public address of the node supplied in the registration screen is displayed in the seeds list. After verification is complete and the cluster is fully registered, the settings page for the Cassandra cluster displays the full correct seed list.

Field	Description
JMX Authentication	JMX Credentials for this cluster. This should be the same for all the nodes.
Cassandra Authentication	Cassandra Credentials for this cluster.
Kerberos Principal	Specify the Kerberos principal to access the Cassandra primary.
Data Centers (optional)	Data centers for this cluster.
Commit Log Archive Location (optional)	Specify the commit log archive location. You must specify the commit log archive location when registering the Cassandra cluster as a source for point-in-time recovery. You can edit or remove the commit log archive location if no backup snapshot and local log Protection Group exists for the source. Removing this value disables PITR for the source.
Cassandra Configuration Details	Displays the information from the Cassandra configuration directory. If DSE Cassandra is selected, DSE-specific configuration is also displayed.

11. Click **Register**.

If the registration is successful, the registered cluster appears in the list of the **Sources** page.

Note:

- Registering a faulty source triggers an error. Users must Update Cassandra Configuration to correct the information and save it to resolve the errors.
- When registering hybrid Cassandra environments where some data centers are not reachable from the Cohesity cluster, you may encounter a Deadline Exceeded error during registration. In such cases, the `imanis_grpc_connection_timeout_msecs Magneto` gflag may need to be adjusted on the Cohesity cluster. Please contact Cohesity Support for assistance with configuring the appropriate timeout settings.

Update Cassandra Configuration

You can update the Cassandra details that you provided during the registration process with the latest Cassandra configuration.

To edit the Cassandra Source Configuration:

1. Navigate to **Data Protection > Sources**.
2. Browse to the required Cassandra source you want to edit, click ⋮ and select **Edit**. The **Edit Cassandra Registration** dialog is displayed.
3. Edit the authentication details and other configurations as required.

   Note: Editing or removing the commit log archive location is not allowed while local log Protection Groups and backup snapshots exist for the Cassandra source. You must delete associated local log Protection Groups and backup snapshots before you edit or remove this setting.

   1. If the entire Cassandra configuration has changed and new nodes are to be added, click ✎ adjacent to **Host Details**.

      Note: Performing step 4 leads to rediscovering the Cassandra source. The source ID, protection runs and Protection Groups that are already created are retained.

   2. Edit the required details and click **Update**.
   3. Enter the Cassandra authentication details.
4. To update the details, click **Update**.

Refresh the Cassandra Source

You can manually refresh the Cassandra source configuration and fetch the latest changes on the source.

To refresh the Cassandra Source Configuration:

1. Navigate to **Data Protection > Sources**.
2. Click ⋮ next to the Cassandra source and select **Refresh**.

When you refresh the Cassandra source, Cohesity auto-updates the Cassandra version and seed nodes. During the refresh, seed nodes are retrieved using the nodetool getseeds API.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Consideration

The output of the nodetool getseeds API should be consistent across all nodes in the registered data centers of the cluster.

Unregister the Cassandra Source

If you want to stop taking backups of the Cassandra source, you can unregister the source from Cohesity.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the Cassandra source.

To unregister the Cassandra Source:

1. Navigate to **Data Protection > Sources**.
2. Click ⋮ next to the Cassandra source and select **Unregister**.

Backup Cassandra

27 July 2026

Cohesity provides a simple sequence of tasks that constitute the procedure for backing up your Cassandra source. Once Cassandra is registered as a source on Cohesity, you can start configuring your backups.

> **Note:** If the regular backup run for the Cassandra Primary is canceled, Cohesity will automatically delete the snapshots of the keyspaces that were created on Cassandra during the backup process.

Protect Cassandra Objects

Cassandra Objects are protected using a Protection Group and a Protection Policy. A Protection Group enables you to select the objects that you want to back up. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured, retained, archived, and replicated.

A Protection Policy is a reusable list of settings that define the type of backup, backup schedule, and backup snapshot retention. Also, for disaster recovery, you can add replication for the Protection Group using the Protection Policy. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

To protect a Cassandra object:

1. **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases > Cassandra**. The **New Protection** page is displayed.
3. **Source**: Select the required Data Source from the drop-down menu.

   > **Note:** For commit log backup, select the Cassandra source that has the **Commit Log Archive Location (optional)** specified.

4. **Data Centers**: Select the required **Data Center** from the drop-down menu.
5. **Log Backup Selection**: Toggle on the **Log Backup** setting to schedule a backup of the log archives of the Cassandra source.

   > **Note:**
   > - Once **Log Backup** is enabled you cannot select any objects.
   > - You can only create one Log Backup Protection Group for a Cassandra source.

6. Select the required objects to protect.

   > **Note:** This option is available only if the **Log Backup Selection** toggle setting is off.

7. To auto-protect an object, click adjacent to it. For more information, see Enable or Disable Auto Protect.
8. Click **Save Selection**.
9. Under **Protection Group**, enter a Protection Group name. Existing Protection Groups cannot be used while creating a new Cassandra protection.
10. Under **Policy**, select an existing protection policy or create a New Policy if the available policies do not meet your schedule and retention needs.

    > **Note:** For commit log backup, only policies with Log Backup are made available to select.

11. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
12. After selecting a Storage Domain, if you need to change any of the additional settings on the Cassandra Protection page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
13. Click **Protect**.

Your new Protection Group is now active and running.

Additional Settings

Field	Description
Start Time	The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection group by selecting a different time zone.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Field	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the cluster. By default, the maximum scaling is set to 16.
Bandwidth Throttling	Maximum network bandwidth that each backup stream can use for exchanging data with the cluster.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • Success: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Field	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run.
Description	Optional. Enter a brief description about the Protection Group.
Pause Future Runs	Toggle on this setting to stop new runs of this Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current protection run continues to execute and only future protection runs of this protection group are paused. **Note:** Once the **Pause Future Runs** option is enabled, protection runs are not scheduled.

Manage the Cassandra Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to an Existing Protection Group

You can add new objects to the Protection Group created for Cassandra protection. The newly added objects are backed up from the next protection run.

To add objects to the Protection Group:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cassandra.
3. Click ⋮ next to the Protection Group and select **Edit**.

4. In the Edit Protection page, Click ✎ next to **Objects**.
5. Make the required changes. For more information, see Protect Cassandra Objects.
6. Click **Protect**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the objects that are added to the Cassandra source. You can auto protect either the entire cluster or a Keyspace.

> **Note:** When auto-protect is enabled on the entire Cassandra source, the system_auth and system_schema keyspaces may appear as protected in the **Protected Objects** and **All Objects** tabs on the **Sources** page. However, protection of system keyspaces is not supported as part of regular protection jobs. To protect system keyspaces, create a separate protection job by enabling the **Backup System Keyspaces** option on the **New Protection** page.

When the Auto Protect option is enabled:

- All the objects that are added to the Cassandra source in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the objects that are removed from the Cassandra source in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backups are preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly added objects in the Cassandra source to the Protection Group.
- Remove existing objects that are removed from the Cassandra source. The existing backups are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cassandra. For more information, see Protect Cassandra Objects.
3. Click ⋮ next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** area, click **Edit**.
5. In the pop-up window , toggle the to enable or disable Auto Protect.
6. Click **Protect**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cassandra. For more information, see Protect Cassandra Objects.
3. Click ⋮ next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Protect**.

Start an On-demand Protection Run

You can manually start a run for the Protection Group that you have created to backup Cassandra objects. It provides you options to:

- Back up all the objects in the Protection Group.
- Select an incremental or full backup.
- Add replication and archive.

To start a protection run:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cassandra.
3. Click ⋮ next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. The **Backup** drop-down list is defaulted to **Backup all Objects in the Protection Group**.
    2. In the **Backup Type** drop-down list, select the backup type.
    3. If required add details for Replication and Archival.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group. However, if a protection run is in progress, it continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cassandra.
3. Click ⋮ next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered.

To resume future protection run:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click ⋮ next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the snapshots (local, archived, and replicated) that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete the Protection Group

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived snapshots generated by the Protection Group and reclaim the space utilized by the Protection Group.

> **Note:** The storage space reclaimed by deleting the Protection Group takes some time to reflect.

To delete the Protection Group:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Cassandra.
3. Click ⋮ next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:

- Delete Protection Group Only
- Delete Protection Group and Snapshots

Enable/Disable Megafile support

Megafile feature improves Cassandra backup/restore task performance by copying mega files (greater 32 GB in size) by splitting them into subfiles and copying the subfiles in parallel. It requires a modified version of SCP binary on the primary nodes to achieve this. By enabling this feature, you allow the NoSQL and Hadoop Service to copy this binary to primary cluster nodes on demand and execute them as needed.

> **Note:** The directory to copy the binary should be accessible by the user to allow the backup/restore functionality to work. You may also need to explicitly allow this executable in puppet/chef tools (if required).

Run imanis_cli tool from any Cohesity node where the NoSQL and Hadoop Service is running.

```
$ imanis_cli.sh -c megafile --help
INFO : Checking reachability of pod imanis-nodemanager-0
-----------------------------------------------------------------------
imanis_cli.sh -c megafile --connector <cname> [--workdir <dirpath>] --enable
true|false
--connector value can be cassandra-backup or cassandra-restore
--workdir value will be the directory path on primary cluster where
talena-scp will be copied
NOTE: --workdir is not needed in case of '--enable false'
--workdir is optional in case of cassandra-restore
-----------------------------------------------------------------------
```

Enable Megafile support

Run the following commands to enable megafile support:

```
$ imanis_cli.sh -c megafile  --connector cassandra-backup  --workdir
/data/binaries/talenascp --enable true
INFO : Checking reachability of pod imanis-nodemanager-0
NOTE:
Megafile feature improves the performance of Cassandra backup/restore tasks
by
```

```
copying mega files ( > 32GB in size) by splitting them into subfiles and copying the
subfiles in parallel. It requires a modified version of scp binary on the primary
nodes to achieve this.
By enabling this feature, you are allowing the NoSQL app to copy this
binary to primary cluster nodes on demand and execute them as needed. Please
note that the directory to copy the binary should be accessible by the user with
which you are allowing the backup/restore functionality to work. You may also need to
allowlist this executable in puppet/chef tools, if needed
Enable megafile for cassandra-backup connector :  (yes/no) yes ? yes
INFO : Updating value for cassandra.backup.megafile.enable
INFO : Updating value for cassandra.backup.megafile.workdir
INFO : core-site.xml updated successfully.
```

Disable Megafile support

Run the following commands to disable megafile support:

```
$ imanis_cli.sh -c megafile  --connector cassandra-backup   --enable false
INFO : Checking reachability of pod imanis-nodemanager-0
INFO : Updating value for cassandra.backup.megafile.enable
INFO : core-site.xml updated successfully.
```

Recover Cassandra

27 July 2026

You can recover the protected Cassandra Objects using Cohesity. You can recover to the same location or an alternate location.

Considerations

- Recovering at the protection group level is supported for Cassandra.
- Recovering with a suffix at the protection group level applies the suffix to the leaf-level objects in the selected protection group.
- Recovering multiple protection groups in a single recovery task is not supported.

- Recovering items in the same protection group but from different recovery points is not supported in the same recovery task. You need to create different recovery tasks for different recovery points.
- Recovering objects from more than one protection group in a single recovery task is not supported.

Recover Cassandra Objects

To recover Cassandra objects you can search through the backups using the protection group, keyspaces or table name for the objects you plan to recover, and then start the recovery.

To recover Cassandra objects:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Databases > Cassandra**. The **Databases | Cassandra** page appears.
3. Enter characters of the object name or Protection Group name or source name, or Use Wildcard * in Search for Recovery. A list of protected items appears.

   You can optionally filter the search results by Source, Protection Group, or Storage Domain.
4. By default, the latest snapshot is selected for recovery.

   Perform the following steps to recover from a different snapshot:

   1. From the selected list, hover the mouse pointer over the highlighted snapshot, and then click .

      The **Edit recovery point** page appears.
   2. From the built-in calendar, select a date range.
   3. Select the snapshot point for recovery.
   4. Click **Select Recovery Point**.

   Perform the following steps to recover to a point in time:

   1. From the selected list, hover the mouse pointer over the highlighted snapshot, and then click .

      The **Edit recovery point** page appears.
   2. Select **Timeline** view.
   3. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar.
   4. For snapshot recovery, select the snapshot to be recovered by clicking the corresponding blue dot. You can also select any point from the green solid line

on the slider if you want to restore to a specific point in time.

5. Click **Select Recovery Point**.

5. Click **Next: Recover Options**. The recovery details are displayed.
6. Under **Recover As** select **Replace Original**, **Modify Original**, or **New Object**.
   - **Replace Original** - Select the **Data Center** to replace the object(s) back to the source from where it was backed up.
   - **Modify Original** - Choose this option and select the Data Centers to perform the object(s) restore on live Cassandra tables. This option allows you to skip truncate and retain the current Cassandra data while recovering objects.

     > **Warning:** The live Cassandra table restore operation may result in undesirable outcomes, ensure that you review and understand the recommendations and limitations before performing live table restore. For more information, see Considerations for Live Cassandra Table Restore.

   - **New Object**:
     a. **Target** - Specify the target for the recovered object.
     b. **Overwrite existing objects with the same name** - Toggle this setting on if you want to overwrite existing objects with the same name. A warning is displayed, type **YES** to proceed.
7. **Log Restore Directory**: Specify the path of the directory on the Cassandra cluster to which you want to recover the commit logs. You must specify the same path for all Cassandra nodes.

   > **Note:**
   > - The SSH user must have read and write permissions to the archive directory.
   > - The user running the Cassandra process must have read permissions to the archive directory.

8. **Staging Directory List**: Specify the staging directory path to which you want to recover the files. You can specify one or more directories separated by comma.

   > **Note:** Applicable only when restoring DSE 6.x.

9. **Restart Cassandra Cluster**: Toggle the **Restart Using Cohesity** setting on if you want Cohesity to restart the Cassandra cluster.
    a. **Command**: In the **Command** field, specify the command or script that can be used to restart the Cassandra cluster.
        - If the command is provided, Cohesity runs the command as the SSH user specified during the registration process. You must create a script or command so that the Cassandra services can be either stopped or started with the appropriate credentials and permissions.
        - If the command is not provided:
            - For DSE, Cohesity uses nodetool stopdaemon API using remote JMX connection to Cassandra nodes to stop the Cassandra services.

              Cohesity runs the `service dse start` command as the SSH user to start the Cassandra services.
            - For Apache, Cohesity uses nodetool stopdaemon API using remote JMX connection to Cassandra nodes to stop the Cassandra services.

              Cohesity runs the `<Config_Dir>/../bin/cassandra -R` command as the SSH user to start the Cassandra services.
    b. **Restart Immediately**: Select this option to restart the Cassandra cluster immediately after recovery is complete.
    c. **Restart On**: Select this option if you want to schedule the restart of the Cassandra cluster on a specific date and time after recovery is complete.
        i. **Date**: From the built-in calendar specify a date.
        ii. **Time**: From the time picker, select the time.

> **Note:** If you do not want Cohesity to restart the Cassandra cluster, perform the procedure mentioned in Configure Log Replay.

10. Under **Recovery Options**, select the required options for the recovery and complete the following:

Option	Description
Roles and Permissions	Restores the roles and permissions associated with the objects. The option is enabled by default.

Option	Description
Rename	Rename objects. Choose **Bulk Rename** to append a suffix to all the selected objects, or choose **Rename Individual Objects**. This option is displayed only if you select **New Object** in the **Recover As** option.
Object Properties	The following property overrides are supported: • Keyspace property override: This option is only available when the user chooses keyspace objects to be restored. The only available property that can be overridden is **replication**. Replication property example : Simple strategy - {'class': 'SimpleStrategy', 'replication_factor': '2'} Network strategy - {'class': 'NetworkTopologyStrategy', 'dc1': '3'} • Table property override: This option is only available when the user chooses table objects to be restored. The only available properties that can be overridden are **compression** and **compaction**. Compression example - {'chunk_length_in_kb': '64', 'class': 'org.apache.cassandra.io.compress.LZ4Compressor'} Compaction example - {'class': 'org.apache.cassandra.db.compaction.SizeTieredCompactionStrategy', 'max_threshold': '32', 'min_threshold': '4'}
Concurrency	Maximum number of recovery streams that will be created to exchange data with the cluster.
Bandwidth Throttling	Maximum network bandwidth that each recovery stream can use for exchanging data with the cluster.
Task Name	Recovery task name.

11. Click **Start Recovery** to start the recovery operation.

Considerations for Live Cassandra Table Restore

Review and understand the following recommendations and limitations before performing live table restore:

- Live Cassandra table restore can only be performed if you restore the objects on the same Cassandra cluster.
- Ensure you run Imanis Apps on all data centers where the Cassandra cluster will be deployed.
- During live Cassandra table restore operation, renaming or changing object properties is prohibited.

- Point-in-time recovery (PITR) recovery is not supported.
- Sequentially perform the restore operations. Only start live table restore on other data centers once the expected restoration of one data center is successfully finished.
- Live-table restore only applies to objects with a "**NetworkTopologyStrategy**" as the replication strategy.
- Live Cassandra table restore is only possible if the objects being restored exist in the target Cassandra database.
- Live table restore can only be performed if the schema of the objects to be restored is the same as that of the backed-up objects.
- To ensure efficient data restoration, Cohesity recommends taking backup snapshots of different data centers, including the data center from which regular restores will be performed and the data centers from which live-table restores will be performed. These snapshots should be taken almost simultaneously and should contain similar objects.
- If different snapshots have different objects, then live-table restore is only possible for those common objects present in all the snapshots.
- Cohesity recommends replicating snapshots using a replication policy between clusters for both restore types (regular and live-table).
- Live Cassandra table restore only applies if table properties (**compaction** & **compression**) match their backup properties.
- Avoid using any **mutation operation** (**insert**/**update**/**delete**) on any objects (backed up by the same snapshot) after starting a regular restore on one data center and before completing the live-table restore on any other data center.
- When performing a live-table restore, it is crucial to carefully choose the source cluster or data center from where the data will be restored. This will help ensure a successful restoration process.

Cassandra Instant Recovery

27 July 2026

Cassandra Instant Recovery is designed to significantly reduce the Recovery Time Objective (RTO) for Apache Cassandra and DataStax Enterprise (DSE) clusters protected with Cohesity. By retaining native Cassandra snapshots on the original source nodes and coordinating recovery using those snapshots, Cohesity enables a near-immediate recovery experience.

Actual RTO can vary based on factors such as such as source cluster size, snapshot integrity, cluster health, network conditions, and whether any portion of the recovery must fall back to Cohesity-stored backups. For more information, see Continue on Error.

How Instant Recovery Works

During backup, Cohesity invokes the native Cassandra snapshot APIs and retains those snapshots on the original Cassandra source nodes instead of deleting them at the end of the

backup cycle. As a result, two synchronized copies of each snapshot are maintained:

- One copy on the Cohesity cluster (primary backup copy).
- One copy on the original Cassandra source nodes (native source snapshot).

Because the native source snapshots remain available, Instant Recovery can directly reuse them without fully rehydrating data from the Cohesity cluster. Data rehydration refers to the traditional process of copying all backed-up data back from the Cohesity cluster to the source systems. By avoiding this step, Instant Recovery achieves substantially faster restore times.

Prerequisites

Before performing Cassandra Instant Recovery, ensure that all the following prerequisites are met:

- **Configure Protection Policy with Native Source Snapshot Retention**
  Instant Recovery requires keeping native snapshots on the Cassandra source cluster. This must be explicitly defined in the protection policy. To configure, see Configure Protection Policy with Native Source Snapshot Retention.
- **Enable Preserve Native Snapshot in Protection Group**
  Instant Recovery requires that snapshots be retained on the Cassandra source cluster. This setting must be explicitly enabled in the protection group. To configure, see Enable Preserve Native Snapshot in the Protection Group.

Considerations

- **Verification of Source Snapshots**

  Cohesity verifies that valid native Cassandra snapshots exist on the source nodes before starting Instant Recovery. If required snapshots are missing or compromised, Instant Recovery cannot be completed solely from source snapshots. In such scenarios, the recovery behavior depends on how the backup job is configured:

  - If the backup job is run without enabling Continue on Error, the job fails and notifies the user of the failure details, including affected nodes and snapshot files that are missing or compromised.
  - If the backup job is run with Continue on Error enabled, the recovery can proceed by falling back to Cohesity-stored backup copies. This fallback may result in increased recovery time compared to recovery from native Cassandra snapshots.
- **Source and Target Cassandra Clusters Must Match**

  Instant Recovery requires restoring data back to the same Cassandra cluster from which the backup was originally taken. Recovery to a different cluster is not supported.
- **No Topology Changes Allowed**

The Cassandra cluster's topology such as number of nodes, their IP addresses, token ranges, and data distribution must remain unchanged. In addition, the replication factor for each keyspace must match exactly between the source (at the time of backup) and the target environment during recovery. Cohesity validates both the cluster topology and keyspace replication factors at recovery time and prevents Instant Recovery if any mismatch is detected.

Limitations

Cassandra Instant Recovery does not support Point-in-Time Recovery (PITR) or commit-log-based recovery.

Configure Protection Policy with Native Source Snapshot Retention

You can create a new protection policy or modify an existing one to enable native source snapshots.

To create or modify a policy:

1. Navigate to **Data Protection** > **Policies**.
2. Click **Create Policy** located at the top-right of the page or select an existing policy to edit.
3. Click **More Options** and go to **Backup Options**.
4. Select **Native Source Snapshot**.

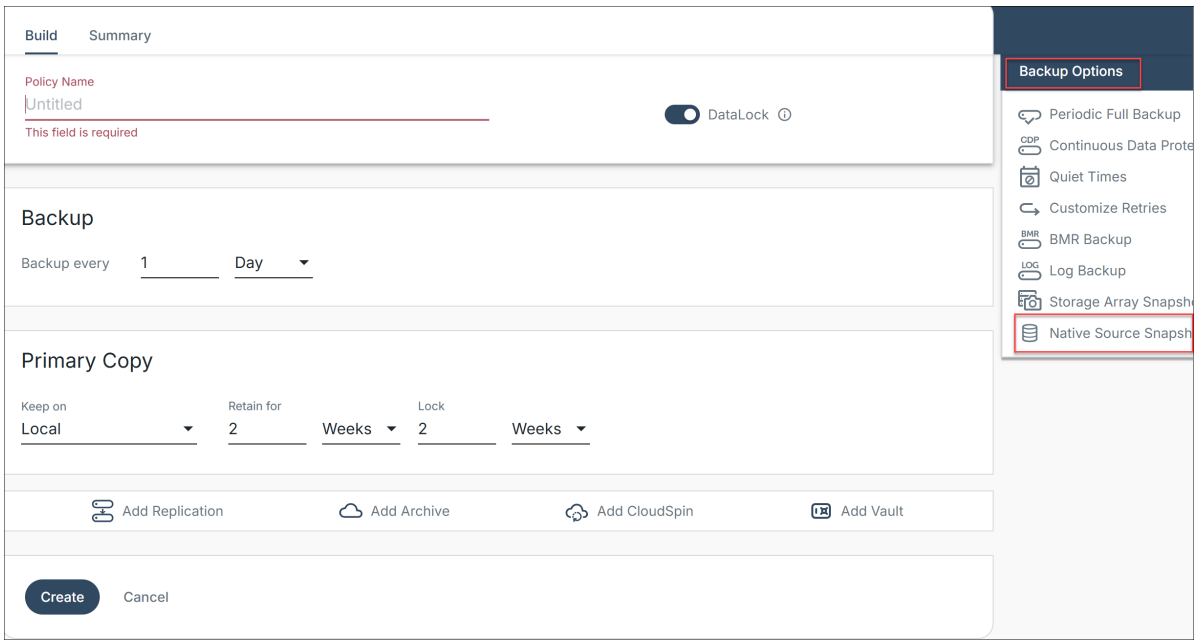

5. Provide the following information:

1. **Policy Name** - Specify a protection policy name.
2. **Data Lock** - DataLock cannot be enabled for native source snapshots because they reside on the user's infrastructure. DataLock applies only to snapshots stored on the Cohesity cluster.
3. **Backup Every** - Specify the minutes, hours, weeks, months, or years to set the frequency for how often backups are captured by the Protection Group. This same schedule applies to both the Cohesity primary copy and the native source snapshot copy.
4. **Native Source Snapshot**
    1. **Retain for** - Specify how many days, weeks, months, or years backups are retained on the native source before they are deleted.

    > **Note:** The retention period for native source snapshots must be shorter than the retention period of Cohesity-stored backup snapshots. If a Cohesity-stored backup snapshot has expired and recovery relies only on a native source snapshot, the recovery may fail or result in corrupted data if any files in the native snapshot are missing or corrupted—even if only a small number of files is affected. Native source snapshots are stored and managed in the customer environment. Cohesity does not control or guarantee their long-term integrity. To ensure a reliable and authoritative recovery point, Cohesity requires that Cohesity-stored backup snapshots be retained longer than native source snapshots.

5. **Primary Copy**
    1. **Keep on** - Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster.
    2. **Retain for** - Specify how many days, weeks, months, or years backups are retained on the Cohesity cluster before they are deleted.
6. Click **Save**.

Enable Preserve Native Snapshot in the Protection Group

To enable native snapshots for a protection group:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **Cassandra** or edit an existing Cassandra protection group.
3. Select any Cassandra source.

4. Under **Protection Group**, enter a Protection Group name.
5. Under **Policy**, select any Cassandra Instant Recovery policy or create a New Policy with Native Source Snapshot option if the available policies do not meet your schedule and retention needs.
6. Under **Additional Settings**, toggle ON the **Preserve Native Snapshot** setting and click **Protect**.

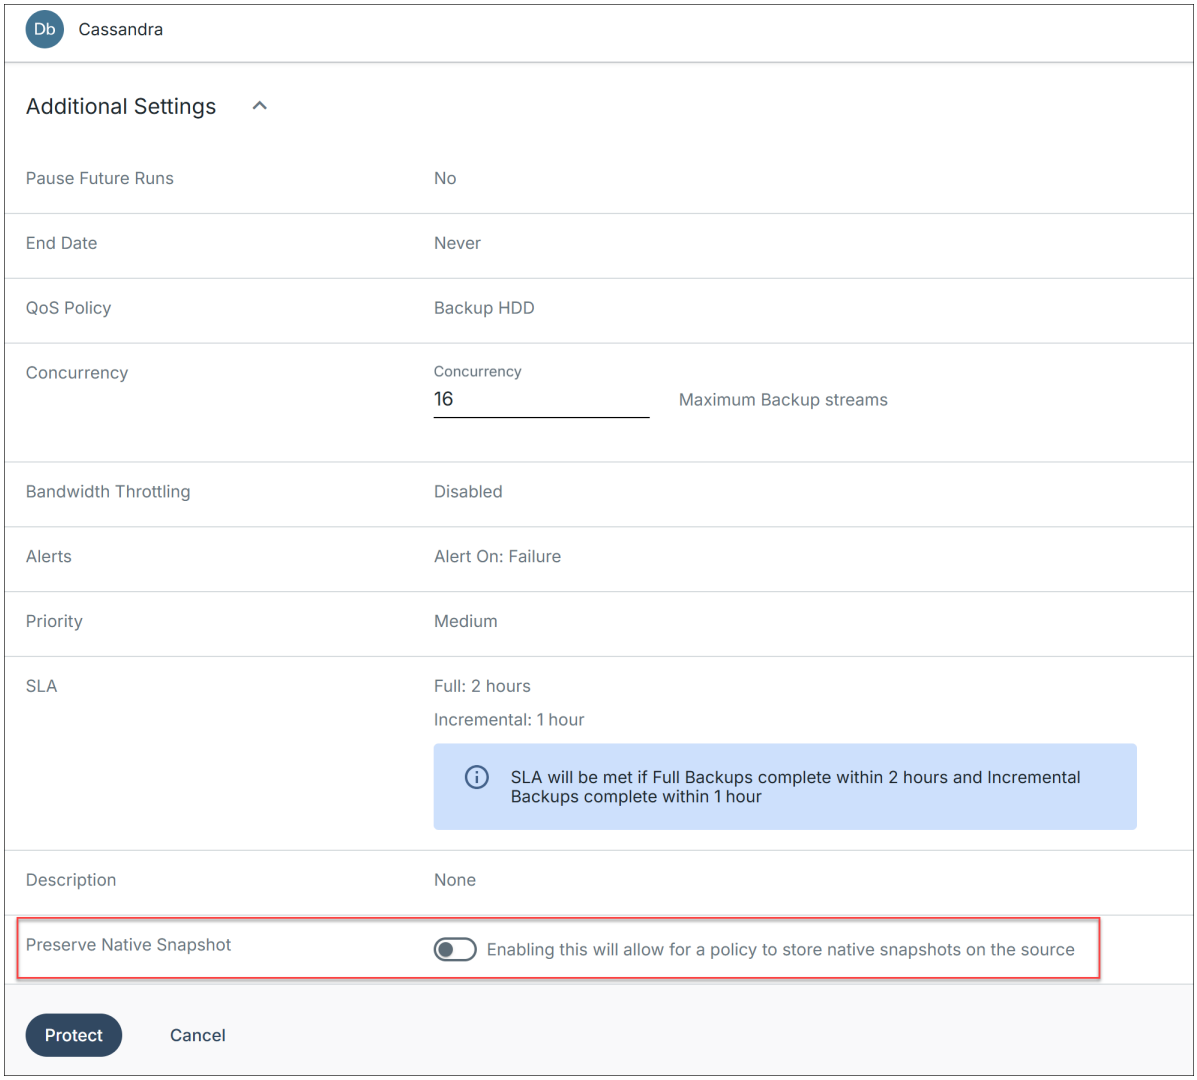

To perform Instant Recovery:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Databases** > **Cassandra**. The **Databases** | **Cassandra** page appears.
3. Enter characters of the object name or Protection Group name or source name, or Use Wildcard * in Search for Recovery. A list of protected items appears.

You can optionally filter the search results by Source, Protection Group, or Storage Domain.

4. From the selected list, hover the mouse pointer over the highlighted snapshot, and then click the edit icon.

   The **Edit Recovery Point** page appears with the following fields:

   1. **Built-in calendar** - Select a date range to filter the available recovery points.
   2. **Snapshot Time** - Shows the exact timestamp when the snapshot was taken. Select the desired recovery point by clicking the radio button next to the timestamp.
   3. **Backup Type** - Indicates the nature of the backup.
   4. **Location** - Displays icons indicating where the snapshot is stored.
      1. **Local** - Snapshot is stored in the Cohesity cluster. To perform the standard recovery, select **Local**.

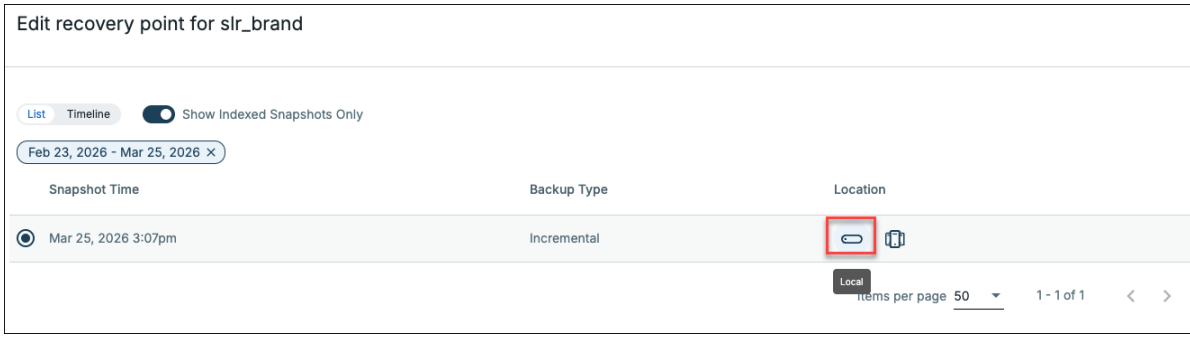

      2. **Source** - Snapshot also exists on the native Cassandra source node. To perform the Instant Recovery, select **Source**.

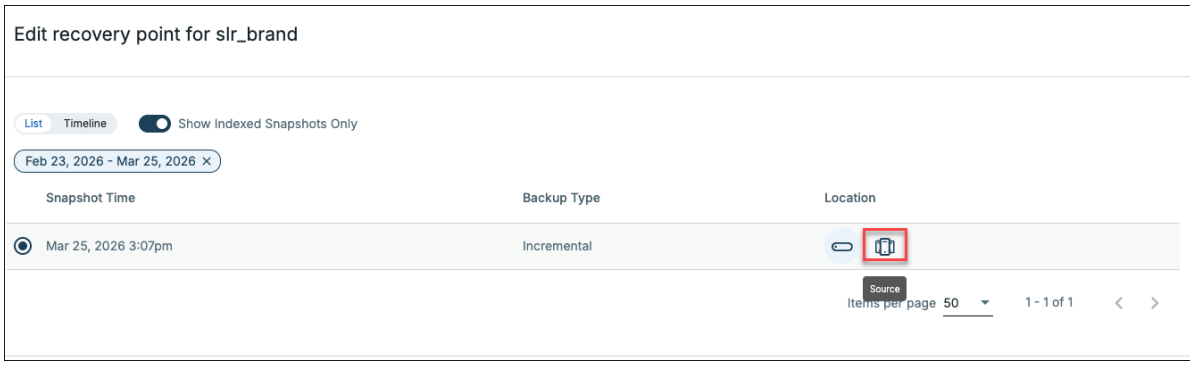

5. Click **Select Recovery Point**.
6. Click **Next: Recover Options**. The recovery details are displayed.
7. Under **Recover As** select **Replace Original** or **New Object**.

1. **Replace Original** - Select the **Data Center** to replace the object(s) back to the source from where it was backed up.
2. **New Object**
    1. **Overwrite existing objects with the same name** - Toggle ON if you want to overwrite existing objects with the same name. A warning is displayed, type YES to proceed.

8. Under **Recovery Options** select the required options for the recovery and complete the following:
    1. **Roles and Permissions** - Restores the roles and permissions associated with the objects.
    2. **Concurrency** - Maximum number of recovery streams that will be created to exchange data with the cluster.
    3. **Bandwidth Throttling** - Maximum network bandwidth that each recovery stream can use for exchanging data with the cluster.
    4. **Task Name** - Recovery task name.
    5. **Rename** - Rename objects. Choose **Bulk Rename** to append a suffix to all the selected objects, or choose **Rename Individual Objects**.
    This option is displayed only if you select **New Object** in the **Recover As** option.
    6. **Continue on Error** – This option is selected as **No** by default. Toggle ON if you want to continue recovery even if one of the objects encounters an error.
9. Click **Recover** to start the Instant Recovery.

Continue on Error

In a standard recovery, the Continue on Error option allows the recovery process to proceed even if some objects fail to recover. Instead of stopping the entire job, the system completes recovery for all valid objects and reports the failures separately.

Instant Recovery relies on native snapshots present on Cassandra source nodes. In some cases, one or more snapshot files may be compromised, for example due to a checksum mismatch.

If Continue on Error is NOT enabled:

- The Instant Recovery fails completely when compromised snapshot files are detected.
- The system reports:
    - Which snapshot files are compromised
    - Why the Instant Recovery failed

This ensures transparency and allows you to decide the next step.

When to Use Continue on Error for Instant Recovery

Use Continue on Error when:

- You want to maximize recovery speed, even if some snapshot files are unusable.
- You are willing to recover healthy data instantly and allow the system to fall back to backup copies for problematic files.

How it Works with Instant Recovery

When Continue on Error is enabled during Instant Recovery:

- All intact native snapshot files are recovered instantly from the Cassandra source nodes.
- Any compromised snapshot files are automatically recovered from the primary backup copy stored in the Cohesity cluster.
- The recovery completes successfully without stopping the entire job.

> **Note:** In this scenario, not all files can be recovered using Instant Recovery. For example, if there are 1,000 files and 2 of them have issues, the remaining 998 files are recovered using Instant Recovery. The 2 problematic files will be restored from the Cohesity backup copy. Because a portion of the data must be recovered from a backup copy rather than the snapshot, we cannot commit to a specific RTO. However, the overall recovery process will still be significantly faster.

Cassandra Point-in-time Recovery (PITR)

Point in Time Recovery (PITR) enables you to have aggressive Recovery Point Objectives (RPO) for Cassandra source. PITR provides more granular recovery to scenarios of accidental data error or corruption. For PITR, commit logs on the Cassandra cluster are backed up which helps you to select a recovery point on a time scale.

Prerequisites

- On the Cassandra cluster:
  - You must enable commit log backup for successful point-in-time recovery.
  - You must ensure that the location to which the commit logs are restored must have sufficient disk space for new logs.
  - Cassandra service must be restarted on all nodes after enabling Commit Log Archival for backup to copy the commit logs.
  - You must run one full backup after enabling point in time on the Cassandra cluster.
  - The SSH user must have read and write access to the backup location.

  - The backup location must be the same for all nodes on the Cassandra cluster where PITR is enabled.
- For PITR, the Cassandra cluster must be registered on one Cohesity cluster only.

Configure Commit Log Archival

For PITR, the archival of commit logs must be enabled on the Cassandra cluster.

You need to configure the commit log archival on the Cassandra cluster in the `commitlog_archiving.properties` configuration file. The commit log is put to the archive location at node startup and when a commit log is written onto disk, or at a specific point-in-time.

The command `archive_command` expects only a single command with arguments. The parameters must be entered verbatim. `STDOUT` and `STDIN` or multiple commands cannot be executed. To execute multiple commands you need to create a script with multiple commands and then call the script from the configuration file.

Enable Commit Log Archival

You must specify the parameters in the `archive_command` to enable commit log archival.

The following is a sample command:

`archive_command=/bin/ln %path /backup/%name` where

- `%path` is the fully qualified path of the segment to archive.
- `/backup` is the location where you want to archive the commit log.
- `%name` is the name of the commit log.

> **Note:**
> - The SSH user must have read and write access to the `/backup` location.
> - The backup location must be the same for all nodes on the Cassandra cluster where PITR is enabled.

Disable Commit Log Archival

To disable commit log archival, you must reset the `archive_command` as `archive_command=` (must be set to empty).

Points to Note

- The modification on the `commitlog_archiving.properties` configuration file must be done on all the Cassandra nodes.
- The Cassandra process must be restarted after editing the configuration file.
- Commit logs are archived frequently because all writes to Cassandra are part of commit logs and this may consume more disk space. The disk on which the archive

location is mounted must have adequate disk space.

- For improved performance, it is recommended to have the Cassandra directory and commit log directory on separate disks.

Configure Log Replay

> **Note:** By default, Cohesity configures the log replay. Optionally, if you do not want Cohesity to make any changes on the Cassandra cluster you can manually configure log replay by following the procedure given below. Before you start with the procedure, contact Cohesity Support to update a setting on the Cohesity cluster thus allowing you to manually configure the log replay.

You need to configure the `restore_command`, `restore_directories`, and `restore_point_in_time` commands on the Cassandra cluster.

The command `restore_command` expects only a single command with arguments. The parameters must be entered verbatim. `STDOUT` and `STDIN` or multiple commands cannot be executed. To execute multiple commands you need to create a script with multiple commands and then call the script from the configuration file.

Enable Restore of Archived Commit Log

You must specify the parameters in the `restore_command` to enable the restoration of the archived commit log.

The following is a sample command:

`restore_command=/bin/cp -f %from %to` where

- `%from` is the fully qualified path of the archived commit log segment from the `restore_directories`.
- `%to` is the name of the live commit log directory.

You must also specify the parameters in the `restore_directories` command.

The following is a sample command:

```
restore_directories=/restore_directory
```

You must also specify the parameters in the `restore_point_in_time` command to restore to a point-in-time. Cassandra replays all the logs in the restore directory till the specified time.

The command accepts the time in the following format represented in GMT:

```
YYYY::MM::DD HH:MM:SS
```

The following is a sample command:

```
restore_point_in_time=2013:12:11 17:00:00
```

Disable Restore of Archived Commit Log

To disable the restore of commit log archival, the `restore_command`, `restore_directories`, and `restore_point_in_time` commands must be set to empty.

Register

To configure Cassandra for data protection on Cohesity, you need to register it as a source.

For more information, see Register a Cassandra Source.

Backup

A point-in-time backup is achieved by backing up the data and the commit logs of the Cassandra cluster. Data can be backed up at a lesser frequency than the logs. In Cassandra, commit logs are generated for the entire node and cannot be filtered based on keyspaces or tables. Therefore, the commit log backup and data backups cannot be tied together based on keyspace or table. You need to create two Protection Groups, one for data backup (keyspace and table) and the other for commit log backup. For more information, see Backup Cassandra.

Perform the following steps to create a Cassandra commit log backup:

- Configure the commit log archival in the Cassandra cluster. For more information, see Configure Commit Log Archival.

  > **Note:** The SSH user must have read and write permissions to the archive directory.

- Specify the archive log location while registering the Cassandra cluster as a source. For more information, see Register a Cassandra Source.
- Create a protection policy that has the commit log backup details. You can also use an existing policy that has the commit log backup details. For more information, see Create or Edit a Standard Policy.
- Create a Protection Group to backup the commit log. For more information, see Backup Cassandra.

  > **Note:** Cohesity removes the backed-up commit logs from the `/backup` location.

Recover

Point-in-time recovery is achieved by copying both the data and commit logs on the Cassandra primary. Data restore is a copy of the `ss` tables from Cohesity. Commit log restore included both copying and replaying the logs.

For more information, see Recover Cassandra.

Considerations

- Commit logs can be replayed only to the original Cassandra cluster and with the same schema.
- Archival of Log Backups is not supported.
- For PITR, object renaming is not supported.
- DDL changes are not restored during log replay.
- You cannot back up specific keyspace or table commit logs for point-in-time recovery. The commit log backup is for the entire Cassandra cluster.
- You must have performed at least one protection run after the first commit log backup for successful point-in-time recovery.
- For recovery from a remote cluster (archival or replication), you must specify the target cluster in the protection run policy, data backup, and commit log backup.
- Cohesity allows only one backup of the commit logs per source.
- You cannot recover objects using a Log Backup Protection Group. For recovery of objects you must create a regular backup Protection Group.
- If a backup job is edited to add objects, the recovery timeline may display those objects even before their first backup run. A PITR recovery attempt for such objects within this time range will fail.
- If a backup job is edited to remove objects, the recovery timeline may continue to display those objects even after they are removed. A PITR recovery attempt for such objects within this time range will also fail.
- To remove commit log–based PITR from a Cassandra source, you must first delete the local log Protection Group, and backup snapshots associated with the source. The replicated log Protection Groups can still exist and do not need to be deleted.
- For PITR:
    - Do not register Cassandra on multiple Cohesity clusters.
    - Do not register the same Cassandra cluster with different names or IPs on a Cohesity cluster.
    - If registering the same source on multiple Cohesity clusters, do not set the retention for Cassandra Log (PITR) Backups on a replicated Cohesity cluster higher than the retention on the Primary Cohesity cluster to avoid data loss.

Cassandra Hot Standby

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Cassandra Hot Standby allows you to recover the latest snapshots from a full or incremental backup. You can recover the incremental changes of the protected Cassandra objects using Cohesity Platform. The latest snapshots are recovered from the source cluster to a target cluster. You can create an incremental recovery of a protected Cassandra object and if there is an incremental change in the backup, the Cohesity Platform initiates an incremental recovery. This ensures that the source cluster and the target cluster are always in sync.

The first recovery is a full recovery and all subsequent recoveries are incremental.

Incremental Recovery of Cassandra Objects

You can search through the backups using the protection group, keyspaces, or table name for the objects you plan to recover.

Perform the following steps to initiate incremental recovery of Cassandra objects:

1. In the Cohesity Platform, select **Data Protection > Recoveries**.

2. Click **Recover** at the upper right and navigate to **Databases** > **Cassandra Hot Standby**. The **New Recovery** page appears.

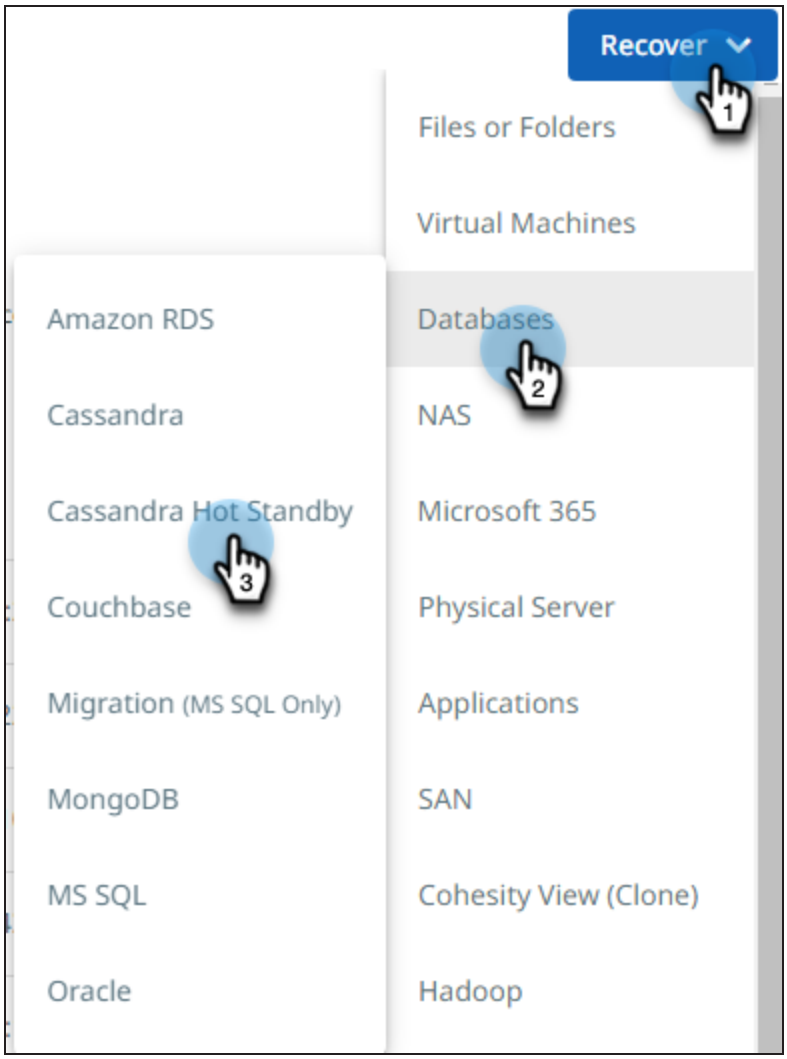

3. Search for the Cassandra database. Enter characters of the name and a list of item displays. You can optionally Use Wildcard * in Search for Recovery. You can also optionally narrow the search results by specifying filter criteria, for example, you can filter the search results by a specific **Protection Group** or **Source** or **Storage Domain.**
4. Select the object(s) or Protection Group and click **Cassandra Hot Standby**. The Cassandra Hot Standby page appears.

   If you need to select another Cassandra object, click the edit icon and select the required object and then click **Cassandra Hot Standby**.
5. In the **Hot Standby** section, do the following:
   a. **Target** - Specify the IP address of the target cluster to which you want to migrate the object.

b. **Overwrite existing objects with the same name** - Toggle this setting on if you want to overwrite existing objects with the same name. A warning is displayed, type **YES** to proceed.

c. **Data Centers** - Select the data center.

> **Note:** This field appears when the source and target Cassandra clusters are the same.

d. **Staging Directory List** - Specify the staging directory path to which you want to recover the files. You can specify one or more directories separated by a comma.

> **Note:** This field appears when the selected target is Cassandra DSE 6.x.

6. In the **Recovery Options** section, do the following:

   a. **Concurrency** - Maximum number of recovery streams that will be created to exchange data with the cluster.

   b. **Bandwidth Throttling** - Maximum network bandwidth that each recovery stream can use for exchanging data with the cluster. Turn on the **Enabling this will allow you to set bandwidth throttling for each worker** toggle to specify the bandwidth throttling. In the **MB/s per Recovery Stream** field, specify the value.

   c. **Task Name** - Enter the recovery task name.

7. Click **Start Hot Standby**.

The incremental recovery workflow appears on the **Recoveries** page and the status shows as **Migrating**.

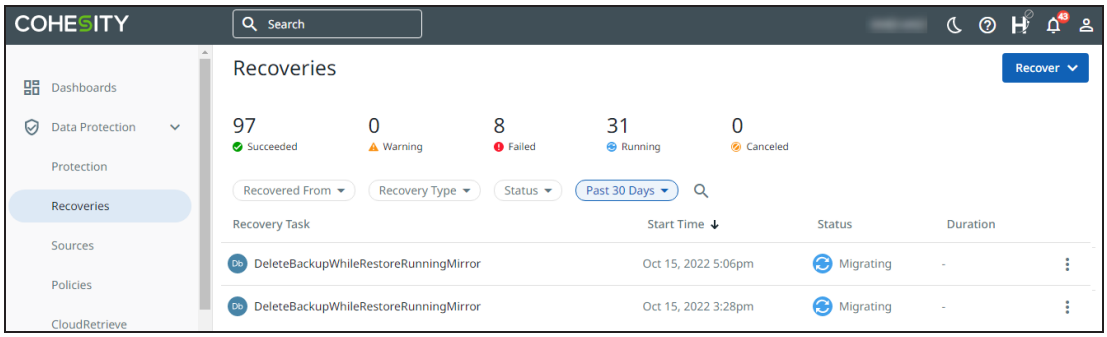

The following animated image shows an example of Cassandra incremental recovery workflow.

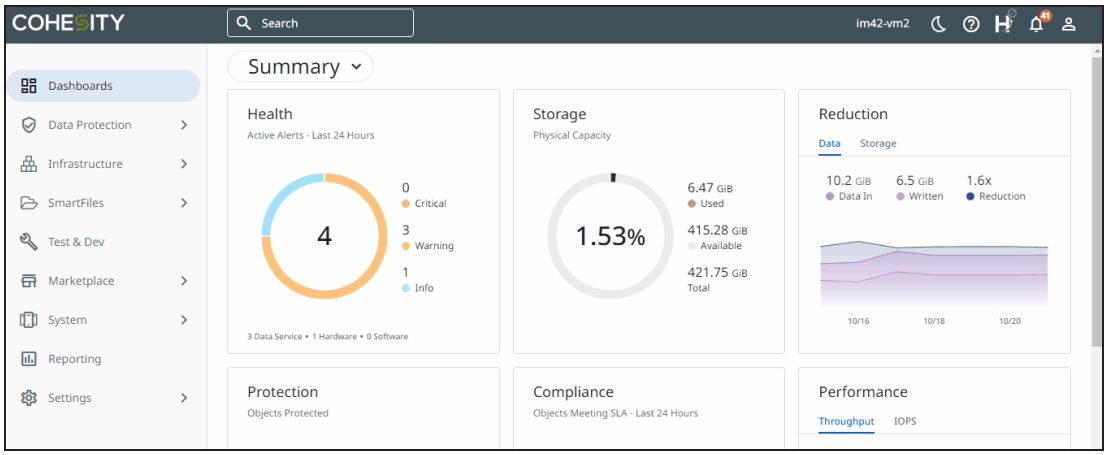

Finalize Migration

You can finalize an in-progress incremental recovery workflow. Finalizing an incremental recovery does not stop the current running recovery but new incremental recoveries are not initiated.

Perform the following steps to finalize an in-progress incremental recovery workflow:

1. Navigate to **Data Protection > Recoveries**. The Recoveries page appears.
2. Find the incremental recovery in the list and do any one of the following:
   - Click the task name to display its detail page, click ⋮ and select **Finalize Migration**.
   - Click ⋮ for the incremental recovery that you want to finalize and click **Finalize Migration**.

     The **Complete Migration** page appears.
3. Click **Complete Migration** if you want to finalize the incremental recovery or click **Cancel** to continue the incremental recovery.

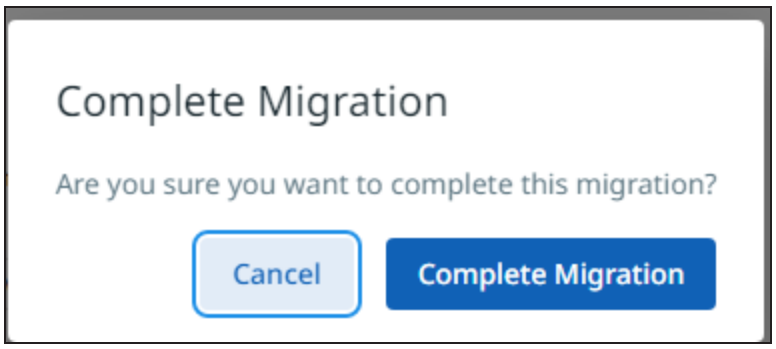

Once you finalize the incremental recovery the status changes from **Migrating** to **In Sync** on the **Recoveries** page.

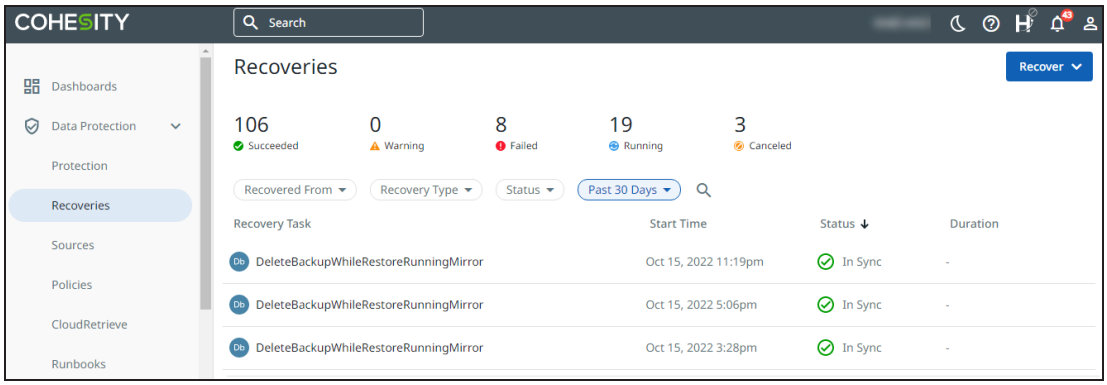

Cancel

You can cancel an in-progress incremental recovery workflow.

Perform the following steps to cancel the incremental recovery workflow:

1. Navigate to **Data Protection > Recoveries**. The Recoveries page appears.
2. Find the incremental recovery in the list and do any one of the following:
   - Click the task name to display its details page and click **Cancel**.
   - Click for the incremental recovery that you want to stop and click **Cancel**. The **Cancel Recovery** page appears.
3. Click **Cancel Recovery** if you want to cancel the incremental recovery or click **Continue Recovery** to continue the incremental recovery.

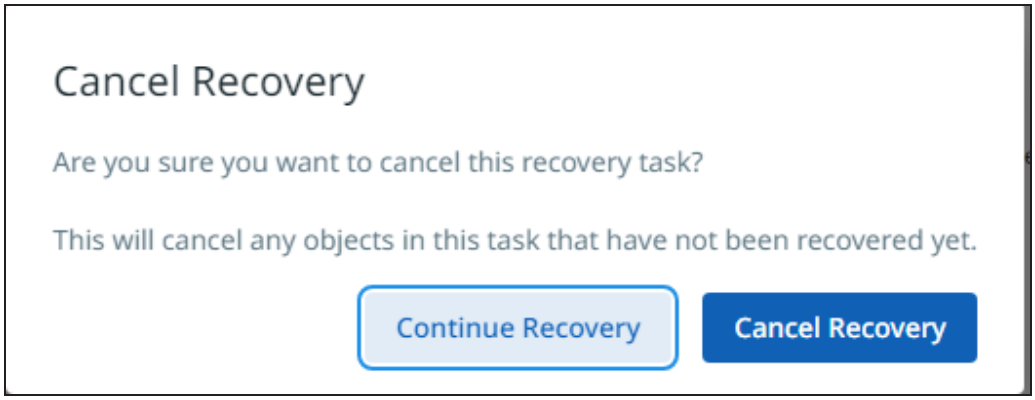

Once you cancel the incremental recovery the status changes from **Migrating** to **Canceled** on the **Recoveries** page.

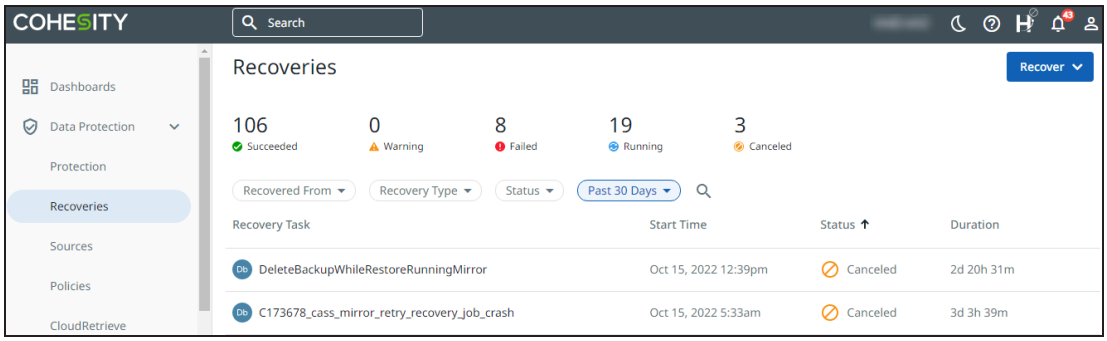

Considerations

- The Cassandra version on the source cluster and the target cluster must be the same.
- Two consecutive protection runs of the Cassandra objects must be less than the gc_grace_seconds of a table.
- Table truncation is not supported.
- Multiple schema alteration queries executed on the single column on the source Cassandra cluster between two backups is not restored to the target Cassandra cluster during a single invocation of incremental data recovery and results in an error.

  For example, if the following schema alteration commands are used on a table:

  - ALTER TABLE newyork.transit_map DROP events [where events column is dropped]
  - ALTER TABLE newyork.transit_map ADD events text [where events column of type text is added again to the schema]

  The preceding commands are performed on the Cassandra cluster and then the incremental data recovery workflow is executed which moves data objects to the destination Cassandra cluster.

  This results in an error with the incremental data recovery workflow as schema alteration queries executed on the source Cassandra cluster will not be propagated to the destination Cassandra cluster during a single invocation of data recovery workflow.

Couchbase Using Connector Agent

Cohesity natively integrates with Couchbase to provide backup and recovery for Couchbase deployments on Linux x86_64 platform.

Cohesity provides a connector agent that enables you to perform a full backup, incremental backup, and restore at the bucket level.

> **Note:** With Cohesity Connector Agent, you can now backup and restore Couchbase buckets and Ephemeral buckets.

Core Components

The following are the core components for protecting Couchbase buckets:

- Couchbase cluster
- Cohesity Couchbase Connector Agent

As illustrated, you can integrate Cohesity with your Couchbase deployments using the Cohesity Connector Agent.

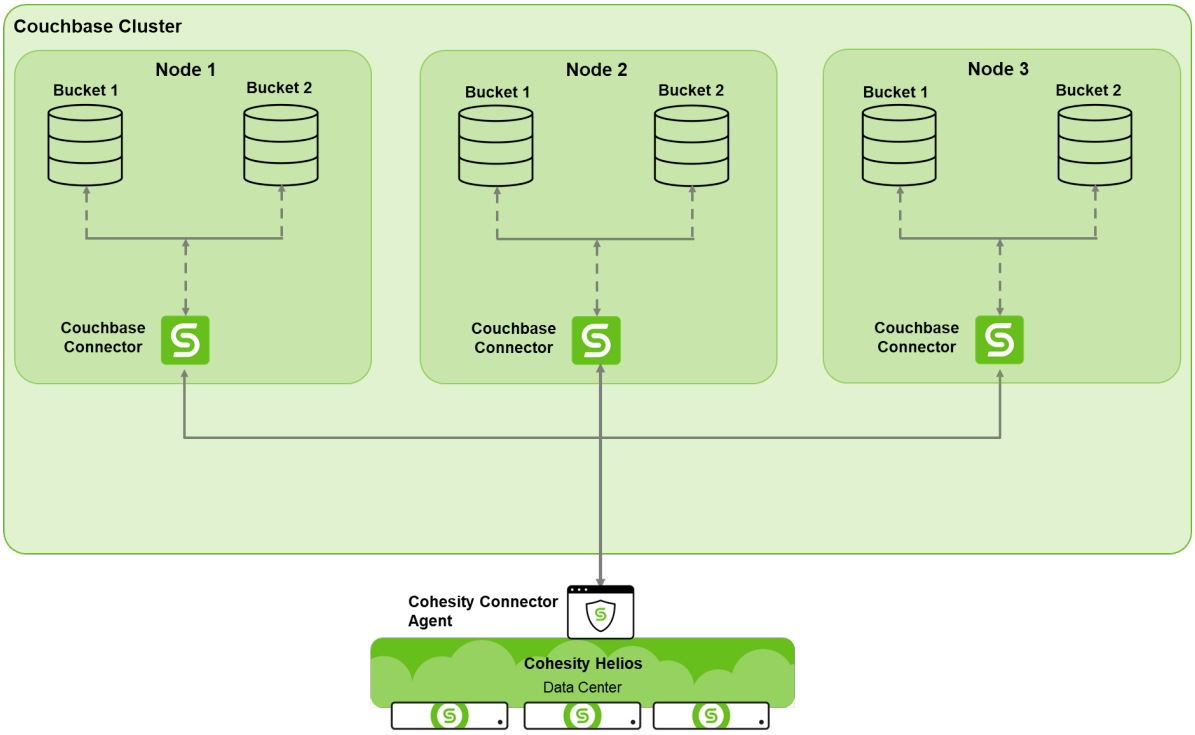

Plan and Prepare for Couchbase Protection Using Connector Agent

Before you register your Couchbase deployment as a source with Cohesity, ensure the following prerequisites:

- The user account of the instance must have bucket admin privileges.
- Download and install the Cohesity Linux Agent on the Couchbase nodes. A minimum of 3 nodes are recommended during registration.
- Download and install the Couchbase connector plug-in.
- Review and understand the prerequisites.

Supported Versions

For more information, see NoSQL and Hadoop Platforms.

Port Requirements

You must open certain ports in the firewall to allow the Cohesity cluster to transmit and receive data. Ensure that you open the following ports in the firewall to allow data transfer between Couchbase and the Cohesity cluster:

Port	Source	Target	Direction	Network Protocol
50051	Cohesity cluster	Couchbase	Unidirectional	TCP/IP
59999	Cohesity cluster	Couchbase	Unidirectional	TCP/IP
8091	Couchbase	Couchbase	Bidirectional	TCP/IP
18091	Couchbase	Couchbase	Bidirectional	TCP/IP
3000	Couchbase	Cohesity cluster	Unidirectional	TCP/IP

*Couchbase nodes must be able to communicate with each other on ports 8091 and 18091.

Minimum Permissions

Cohesity requires the following privileges enabled for the Couchbase user used during source registration for successful backup and recovery:

Permission	Description
Data Backup & Restore	Required to: • Run backup • Run DDL restore • Restore documents to existing bucket
Cluster Admin	Required to create or re-create a bucket during recovery.
Backup Full Admin	Required to: • Perform the backup-related task. • Access to the Couchbase Web Console.

Considerations

Review and understand the following before you protect your Couchbase cluster:

- Cohesity now supports bucket-level, scope-level, and collection-level of backup.
- Once a bucket or its objects such as scopes and collections are configured to be backed up using Cohesity, users should not initiate a manual backup on the same bucket outside of the Cohesity cluster.
- The same objects should not be part of multiple Protection Groups.
- Configure a protection policy to perform a periodic full backup at intervals less than the retention period.
- When restoring to the Couchbase 6.6 version, indexes won't be restored if the rename option is used. Index rebuild will be required after the restore operation, if the indexes get restored to a separate node from where they were backed up. For more information, see the description section on "cbbackupmgr restore" here.
- Cohesity recommends performing functional testing where the bucket size is not more than 250GB.
- Couchbase compaction is a resource intensive operation and can substantially increase the backup or recovery SLAs when run in parallel. It is recommended that the Cohesity workflows and Couchbase compaction are scheduled to run at different time periods.

Prerequisites

Ensure that:

- You use a root user or non-root user to install the Cohesity Linux agent and Couchbase Connector. If you choose to use a non-root user, you must add the `sudo` privileges for the non-root user using the syntax given below:

  ```
  <user_name> ALL=(ALL) NOPASSWD:SETENV: /path/to/command1,
  /path/to/command2, /path/to/command3
  ```

  The following is the `sudo` commands for Couchbase Connector:

  ```
  cohesityagent ALL=NOPASSWD:SETENV: /usr/bin/find,
  /usr/bin/stat,/usr/bin/test, /usr/bin/mv, /usr/bin/cat, /bin/mkdir,
  /bin/rm, /bin/cp, /bin/chown, /bin/chmod, /bin/java
  ```

  Where, `/opt/cohesity/couchbase` is the Couchbase connector installation directory.

  > **Note:**
  > - If the **cohesityagent** user entry is already present in **/etc/sudoers**, append the list of commands to the existing entry.
  > - The actual paths for each command might vary depending on the Platform and OS distribution. For detailed instructions, see Download and Install the Linux Agent.

- While you Register and Manage the Couchbase Source Using Connector Agent, provide the path of the certificate config file in the **Couchbase SSL Settings** field.

  Note:
  Cohesity Agent uses VIPs or FQDN for transmitting data while using S3 protocol. For Cohesity Connector Agents that use S3, the client fails with SSL errors, if x509_extensions is not present in the Cohesity SSL certificates. Ensure that you edit the sample file (**cert_req.conf**) to include the following x509_extensions and add all the VIPs to the `alt_names` list. For more information, see the Update SSL Certificates Using Cohesity CA or Customer CA certificate for Agent to Cluster Communication.

  ```
  #REQ FILE: cert_req.conf
  #--------------------------
  [ req ]
  default_bits    = 2048
  default_md      = sha384 (or sha256)
  prompt          = no
  encrypt_key     = no
  distinguished_name = req_distinguished_name
  req_extensions     = req_ext
  x509_extensions = req_ext
  ```

- Java 1.8 or higher version is installed on the Couchbase nodes.
- Java binary is available in the `$PATH` environment variable.
- The staging directory path must have a free space which is equal to 10% of the bucket size. For more information, see Couchbase documentation.

Download and Install the Cohesity Linux Agent

The Cohesity Linux Agent is available as an RPM-based installer. Based on your requirements, download the installer and install it accordingly on the node(s) of the Couchbase cluster. For detailed instructions, see Download and Install the Linux Agent.

Download and Install the Couchbase Connector

Cohesity provides Couchbase connector RPM for backup and restore of Couchbase deployments.

Download the Couchbase Connector

Cohesity provides a Couchbase Connector to backup and restore Couchbase buckets. The Couchbase Connector is available as a data source agent on the Cohesity Download portal.

Log in to the Cohesity Download portal, download the respective data source agent for Couchbase Connector, and copy the data source agent to the Couchbase host.

> **Note:** With Cohesity Connector Agent, you can now backup and restore Couchbase buckets and Ephemeral buckets.

Install the Cohesity Couchbase Connector RPM

Before you install the Cohesity Couchbase connector, ensure that you have `sudo` privileges.

To install the Cohesity Couchbase connector:

1. Log in to the Couchbase node where you have copied the Cohesity Couchbase connector RPM-based installer, `cohesity-couchbase-connector-<version>.x86_64.rpm`.
2. Run the following command to install the Cohesity Couchbase connector:

   Using `yum`:

   ```
   sudo sh -c 'yum localinstall /tmp/cohesity/cohesity-couchbase-connector-<version>.x86_64.rpm'
   ```

   Using `rpm -ivh`:

   ```
   sudo rpm -ivh cohesity-couchbase-connector-<version>.x86_64.rpm
   ```

Upgrade the Cohesity Couchbase Connector

To upgrade the Cohesity Couchbase connector, run the following command on Couchbase host:

```
sudo rpm -U cohesity-couchbase-connector-<version>.x86_64.rpm
```

Uninstall the Cohesity Couchbase Connector

To uninstall the Cohesity Couchbase connector, log in to the Couchbase node where you have installed the Cohesity Couchbase connector and run the following command:

Using `yum`:

```
rpm U --force <rpm-file name>
```

> **Note:** Once the connector is upgraded, edit and save the source to reflect the newly added groups, scopes, and collections. For more information, see Update the Couchbase Source Configuration.

Using `rpm`:

```
sudo rpm -e cohesity-couchbase-connector-<version>.x86_64.rpm
```

Register and Manage the Couchbase Source Using Connector Agent

To protect your Couchbase buckets, you need to register your Couchbase deployment as a source on the Cohesity cluster. After registering your Couchbase deployment as a source, you can:

- Update the Couchbase source configuration.
- Unregister the Couchbase source from the Cohesity cluster.

> **Note:** With Cohesity Connector Agent, you can now backup and restore Couchbase buckets and Ephemeral buckets.

Register Couchbase as a Source

Before you register Couchbase as a source, ensure that you:

- Review the considerations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for Couchbase Protection Using Connector Agent.

To register Couchbase as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
    a. From the **Source Type** drop-down menu, select **Couchbase**.
    b. From the **Host OS Type** drop-down menu, select **Linux**.
    c. In the **Hostnames/IP Addresses** field, enter the hostname or IP address of the node(s) that you have identified to run the Cohesity Couchbase connector. For more information, see Plan and Prepare for Couchbase Protection Using Connector Agent.
    d. In the **Datasource Agent Installation Path** field, enter the directory path on the Couchbase Linux machine, where you have installed the Cohesity Couchbase connector. For more information, see Install the Cohesity Couchbase Connector RPM section in Plan and Prepare for Couchbase Protection Using Connector Agent.
    e. Optional. In the **App Authentication** field, enter the username and password.
    f. Enable **Couchbase SSL Settings** to enable secure connection to the Couchbase cluster while reading or writing data. In the **SSL Certificate Path** field, enter the path to the SSL certificate on the Couchbase node which is used

to securely connect to the Couchbase cluster.

> **Note:**
> Cohesity Agent uses VIPs or FQDN for transmitting data while using S3 protocol. For Cohesity Connector Agents that use S3, the client fails with SSL errors, if x509_extensions is not present in the Cohesity SSL certificates. Ensure that you edit the sample file (**cert_req.conf**) to include the following x509_extensions and add all the VIPs to the `alt_names` list. For more information, see the Update SSL Certificates Using Cohesity CA or Customer CA certificate for Agent to Cluster Communication.
>
> ```
> #REQ FILE: cert_req.conf
> #-------------------------
> [ req ]
> default_bits    = 2048
> default_md      = sha384 (or sha256)
> prompt          = no
> encrypt_key     = no
> distinguished_name = req_distinguished_name
> req_extensions     = req_ext
> x509_extensions = req_ext
> ```

g. Under **Source Settings**, enter the details of the Couchbase source:

Field	Mandatory?	Description
Staging Directory Path	Yes	Path to the staging directory. Spaces are not supported in the staging directory path.
Couchbase Installation Directory	Yes	Path to the Couchbase installation directory. Spaces are not supported in the Couchbase installation directory path.
Couchbase Administration Port or Secured Administration Port	No	Port for Couchbase administration

h. Click **Register**.

Update the Couchbase Source Configuration

After registering Couchbase as a source on the Cohesity cluster, you can update the Couchbase source configuration.

> **Note:**
> - Ensure that the Cohesity connector is installed on all the node(s) in the Couchbase cluster.
> - You cannot modify the source name.

To update the Couchbase source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Couchbase source and select **Edit**.
3. In the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Unregister the Couchbase Source

If you plan to stop protecting your Couchbase buckets, you can unregister the Couchbase source from the Cohesity cluster.

To unregister the Couchbase source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Couchbase source and select **Unregister**.

Backup Couchbase Using Connector Agent

Cohesity provides a simple workflow to back up your Couchbase buckets. After registering Couchbase as a source on Cohesity, you can start configuring your backups.

> **Note:** With Cohesity Connector Agent, you can now backup and restore Couchbase buckets and Ephemeral buckets.

Backup Process

Cohesity provides a Couchbase Connector for backup and restore of Couchbase deployments. The Couchbase Connector leverages the native backup capabilities provided by Couchbase.

After registering the Couchbase source on Cohesity, you can create a Protection Group and specify the objects you want to backup.

Protect Couchbase Deployments Using Connector Agent

Couchbase deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects such as buckets, scopes, and collections that you want to back up.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection GroupFor more information, see Manage Protection Policy.

To protect Couchbase deployments:

1. Do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Couchbase source and then select **Protect**.
2. From the **Source** drop-down list, select the registered Couchbase source or click **Register Source** to register a Couchbase source. For more information, see Register Couchbase as a Source.
3. In the **Objects** section, select the object(s) that you plan to protect.

   > **Note:** You can select the objects using checkboxes or auto-protect. If you select the checkbox, all the collections inside the selected object will be backed up. If you select auto-protect, all the collections inside the selected object and any newly added collections will also be backed up.
   > For recovering buckets, the auto-protect has to be enabled at the bucket level and for scopes, the auto-protect has to be enabled at either scope or bucket level. The manual selection of scopes and buckets using the checkbox will only allow recovery at the scope level.

4. In the **Protection Group** section, enter a name for the Protection Group in the **Name** field.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the **Settings** section, enter the following details:

Field	Mandatory?	Description
Storage Domain	Yes	Select the appropriate storage domain. For more information, see Manage Storage Domains.
Start Time	No	Define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
Mounts/VIPs	No	N/A for Couchbase data protection.

7. In the **Additional Settings** section, enter the maximum backup streams in the **Concurrency** field. It defines the maximum number of backup streams that will be created to exchange data with the cluster. Couchbase servers do not efficiently handle more than 6 parallel connections. So, it is recommended not to increase the concurrency beyond 6 in the Backup settings.
8. Click **Protect**.

Manage the Couchbase Backup Configuration Using Connector Agent

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Couchbase.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Couchbase
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity UI.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Couchbase.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Couchbase Using Connector Agent

You can recover the objects such as buckets, scopes, and collections in Couchbase to the same Couchbase cluster or a different Couchbase cluster.

> **Note:** With Cohesity Connector Agent, you can now backup and restore Couchbase buckets and Ephemeral buckets.

Recover Couchbase Buckets

To recover Couchbase buckets:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the bucket that needs to be recovered by entering characters of the bucket name, or by specifying the wildcard character *.
4. From the list of search results, select the object and the desired snapshots to recover.
5. Click **Next: Recover Options**.
6. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	Overwrites the original object in your primary environment. Type **Yes**, if you want to overwrite the original object.
New Object	To recover to a different Couchbase cluster. Select the **Target** Couchbase cluster to which the data needs to be recovered. Also, enable **Overwrite existing object with the same name** and type **Yes**, if you want to overwrite the original object in your primary environment.

7. In the **Restore Settings** section, select **DDL Only Restore** or **Disable GSI Indexes** based on your requirements. **DDL Only Restore** restores only the bucket metadata, and **Disable GSI Indexes** skips GSI indexes during restore.
8. Toggle the button to enable the Filtering Options section.

   In the filter keys, enter the document key or document key regex.

   In the filter values, enter any document value regex or the value to match.

   > **Note:** The regular expression must be compatible with the RE2 syntax.

9. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts/VIPs	N/A
Concurrency	The maximum number of recovery streams that will be created to exchange data with the cluster. Couchbase servers do not efficiently handle more than 6 parallel connections. It is recommended not to increase the concurrency beyond 6 in the Recovery settings.

Options	Description
Rename	If you select **New Object** in the **Recover As** section, this option allows you to rename any one of the objects (bucket or its objects such as scopes and collections). The original name of the object is displayed in the **Original Name** field. Enter the new name for the object in the **New Name** field. For example, if a bucket has a scope with 2 collections, bucket1.scope1.col1 and bucket1.scope1.col2 then during recovery the scope can be renamed as follows: • bucket1.scope1Ret.col1 • bucket1.scope1Ret.col2 **Note:** Make sure to rename both collections according to the scope name.
Task Name	Change the default name of the recovery task.

10. Click **Recover**.

Cancel a Recovery Run

When you cancel a recovery run, the Couchbase connector cancels the job by terminating the recovery processes.

To cancel a recovery run:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click the action icon (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
   - **Continue Recovery** if you want to continue to recovery run.
   - **Cancel Recovery** if you want to cancel the recovery run.

MongoDB

27 July 2026

The Cohesity MongoDB DataProtect solution natively integrates with MongoDB to provide consistent backup and recovery for MongoDB deployments.

Based on the policy settings, the Cohesity NoSQL and Hadoop Service initiates a backup of the MongoDB Database. The first backup is a full backup and everything after that is an incremental backup only. Occasionally, a full backup will be required if the oplog rollsover between two incremental backups. This can be avoided by increasing the frequency of backups or by increasing the size of the oplogs. The service also automatically determines the number of backup channels, based on the compute resources on Cohesity and the

topology of the MongoDB cluster. The necessary Cohesity Views are automatically generated where the backups are going to be written. Communication between Cohesity and the MongoDB Databases is managed by the Cohesity NoSQL and Hadoop Service as described in the section below. See the following topics for details:

Core components

The following lists the core components of the Cohesity MongoDB solution:

- Physical or virtual MongoDB instance.
- Cohesity Data Cloud.
- Cohesity Compute Nodes.
- Cohesity NoSQL and Hadoop Service running on Cohesity Compute Nodes.
- Optional. Archive or tiering target at your data center or cloud provider of choice.

As illustrated in the figure below, you can integrate Cohesity with an existing MongoDB infrastructure by installing the Cohesity NoSQL and Hadoop Service on Compute Nodes. Once the service is installed and running, you can create protection policies, Protection Groups, and replication schedules for the MongoDB database. Once the data is protected, it can be archived, sent to tape, and or replicated to another instance running in the cloud or on-premises as shown in the image below.

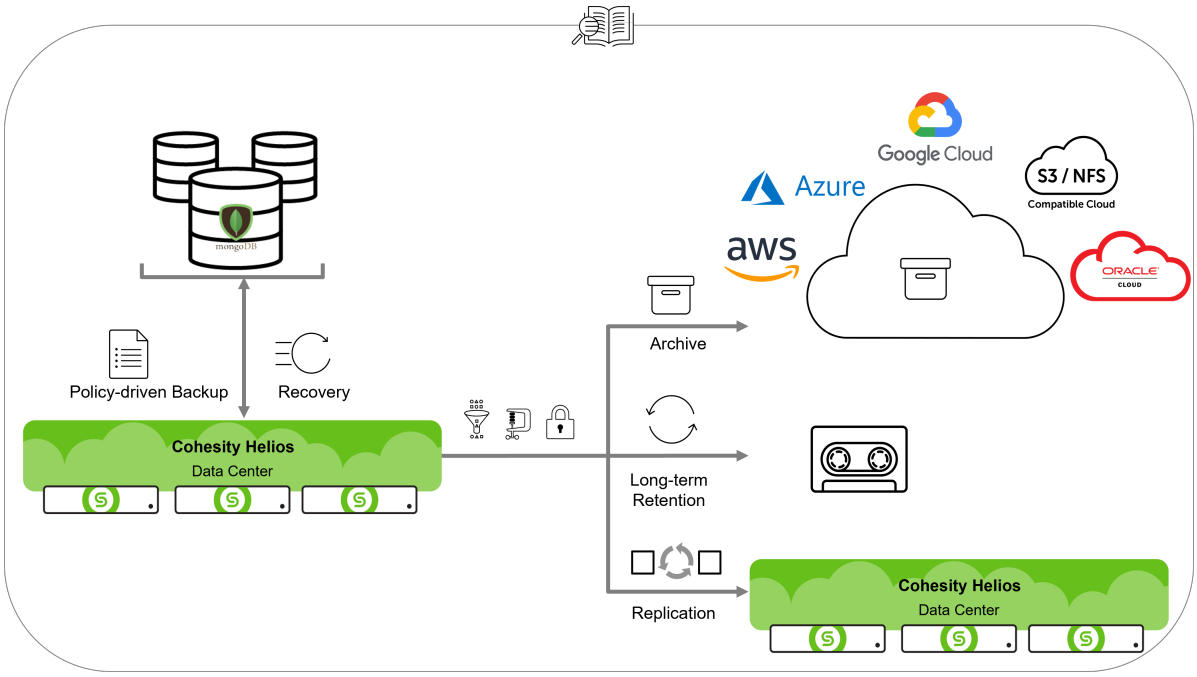

MongoDB Features

Following are the key MongoDB-specific features of the solution:

- Support replica sets and sharded clusters
- Granular backup and recovery - backup and recover individual collections
- Index definitions are backed up and restored
- Option to back up from secondary MongoDB nodes
- Handle topology changes seamlessly (node addition, removal, IP changes, etc.)
- Support Kerberos, LDAP, and SCRAM authenticated clusters
- Support SSL-enabled clusters.
- Support Continuous Data Protection (CDP).

Plan and Prepare for MongoDB Protection

27 July 2026

The following topics provide guidance that helps you get ready to implement the MongoDB DataProtect Solution.

Supported Versions

For more information, see NoSQL and Hadoop Platforms.

Prerequisites

Make sure the following prerequisites are met before backing up MongoDB databases:

- Configuring Continuous Data Protection (CDP) for MongoDB
- The NoSQL and Hadoop Service is in Running status. See NoSQL and Hadoop Service for details.
- Refer to the NoSQL and Hadoop Service - Post Installation Procedures to get a general idea about the following operations:
  - Manage SSL enabled NoSQL primary clusters
  - Work with Kerberos Security enabled primaries for data protection
- Ensuring all Nodes in MongoDB Cluster are Reachable
- Syncing Time Between MongoDB Cluster and Cohesity Cluster
- Ensuring Unique Value of _id in a Collection

Configuring Continuous Data Protection (CDP) for MongoDB

Ensure that the NoSQL and Hadoop Service app is installed and running.

To enable MongoDB CDP service, use the following command:

```
$ imanis_cli.sh -c mongocdp enable --cpu <cpu-cores> --memory <memory-GB> --
replica <num-replicas>
```

- **cpu-cores**: Number of CPU cores allocated to each MongoCDP service instance
- **memory-GB**: Amount of memory allocated to each MongoCDP service instance
- **num-replicas**: Number of MongoCDP service instances. Depending on the load on MongoDB primary and the recovery point objective, one or more MongoDB CDP services can be provisioned. The use of multiple smaller CDP services can help RPO in case of nodes going down.

For example, the below command enables MongoDB CDP with single replicas of 1core/2GB memory.

```
$ imanis_cli.sh -c mongocdp enable --cpu 1 --memory 2 --replica 1
INFO : Checking reachability of pod imanis-nodemanager-0
INFO : Log File : /tmp/imanis-cli-workdir.10686/cli/imanis-mongo-cdp.log
INFO : Enabling MongoCDP pods deployment
INFO : Reading allocatable cluster resources
INFO : Total allocatable Memory(GB) : 239
INFO : Total allocatable CPUs : 47
INFO : Number of MongoCDP replicas : 1
INFO : CPUs per MongoCDP replica : 1
INFO : Memory per MongoCDP replica(GB) : 2
INFO : Total CPUs for MongoCDP replicas : 1
INFO : Total Memory for MongoCDP replicas(GB): 2
INFO : Total Memory for Other Services(GB) : 15
INFO : Total CPUs for Other Services : 10
INFO : Small-Nodemanager instance config : 16GB/4 CPUs
INFO : Medium-Nodemanage instance config : 32GB/8 CPUs
INFO : Large-Nodemanager instance config : 40GB/10 CPUs
Please confirm above configurations for 'NoSQL & Hadoop Service' app
deployment with MongoCDP Pods
Do you want to continue (yes/no) yes ? yes
--------------------------------------------------------------------------
----
mongocdp : service(s) will be stopped (if running)
--------------------------------------------------------------------------
----
service : mongocdp is being stopped
```

```
Stopping MongoCDP services on imanis-mongocdp-0

MongoCDP services stopped successfully

Setting up MongoCDP services...

Configuring MongoCDP services...

MongoCDP services configured successfully

Deleting stale pod directory:
/cohesity/mount/sharedview/imanis/cdp/mongo/v1/pods/imanis-mongocdp-1

Starting MongoCDP service on imanis-mongocdp-0

MongoCDP service is running on [imanis-mongocdp-0]

MongoCDP service is running

Setting up MongoCDP service successfull

Requested number of pods with given memory and cpu cores is updated
successfully
```

To verify, once the app is Running, run the following command to check the pod status:

```
$ kubectl.sh  get pods  -A
```

Example:

```
$ kubectl.sh get pods -A
NAMESPACE                NAME
 READY   STATUS    RESTARTS   AGE
athena-6                 imanis-api-server-0
 1/1     Running   0          102m
athena-6                 imanis-api-server-1
 1/1     Running   0          102m
athena-6                 imanis-historyserver-0
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-0
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-1
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-2
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-3
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-4
 1/1     Running   0          102m
athena-6                 imanis-resourcemanager-0
 1/1     Running   0          102m
athena-6                 imanis-resourcemanager-1
 1/1     Running   0          102m
```

```
athena-6                  imanis-zookeeper-0
 1/1     Running   0          102m
athena-6                  imanis-zookeeper-1
 1/1     Running   0          102m
athena-6                  imanis-zookeeper-2
 1/1     Running   0          102m
cohesity-agentinstall     agentinstall-controller-0
 2/2     Running   0          12d
cohesity-coredns          coredns-8566b9b44-4hrrk
 1/1     Running   0          12d
cohesity-coredns          coredns-8566b9b44-97nvj
 1/1     Running   0          12d
cohesity-coredns          coredns-8566b9b44-sxrw5
 1/1     Running   0          12d
cohesity-infraoperator    infra-operator-controller-manager-787dbddd76-n7tcc
 2/2     Running   0          12d
cohesity-infraoperator    infra-operator-controller-manager-787dbddd76-qqd96
 2/2     Running   0          12d
cohesity-magnus           magnus-actor-f8b5fbcc-q6fgb
 1/1     Running   0          12d
cohesity-magnus           magnus-server-0
 1/1     Running   0          12d
```

To disable MongoDB CDP:

1. Run the following command:

   $ imanis_cli.sh -c mongocdp disable

2. Check the pod status from Cohesity node with command "kubectl.sh get pods -A" and verify that no mongocdp pods are running.

> **Note:** Once MongoCDP service is disabled, MongoDB data protection would no longer work.

Ensuring all Nodes in MongoDB Cluster are Reachable

Make sure that all nodes in the MongoDB cluster are reachable through hostname and IP address from all containers running the NoSQL and Hadoop Services. Generally this is not a problem if the MongoDB servers are registered in DNS. If this is not the case, then refer To Manage the NoSQL and Hadoop Services.

Sizing the Primary Oplog of MongoDB Primary

Ensure that the oplog of MongoDB primary is sized to hold at least 30 minutes of data even when there is heavy load. This means, MongoDB primary would need a larger oplog if it has a larger maximum ingestion rate.

Syncing Time Between MongoDB Cluster and Cohesity Cluster

Ensure that the time on MongoDB cluster and Cohesity cluster is in sync. This is a mandatory requirement as NoSQL and Hadoop Services access MongoDB Oplogs for incremental backups.

Ensuring Unique Value of _id in a Collection

This prerequisite is to verify that all collections have a unique value for `_id`.

The field name _id is reserved for use as a primary key for a document and its value must be unique in the collection.

> **Warning:** If a collection has more than one document with the same _id, it can cause data corruption as we cannot identify the documents uniquely. Even with data corruptions backup/restore will succeed without any errors.

Refer to the following MongoDB documentation for more information:

- Field Names
- Unique Indexes

Firewall Ports

The following table lists the required firewall ports that are to be open to enable communication between Cohesity and MongoDB DataProtect Solution.

MongoDB Connector connects to MongoS and MongoD processes for backup and recovery. The ports associated with these processes should be open and accessible from the Cohesity nodes.

Port	Source	Target	Direction (From Node)	Network Protocol	Usage Notes	Type of Traffic
27017	Cohesity cluster	MongoDB Nodes	Outgoing	TCP	Used for MongoDB backups and recovery.	Backup and Recovery
27018	Cohesity cluster	MongoDB Nodes	Outgoing	TCP	Used for MongoDB backups and recovery.	Backup and Recovery
27019	Cohesity cluster	MongoDB Nodes	Outgoing	TCP	Used for MongoDB backups and recovery.	Backup and Recovery

Minimum Permissions

This section describes the minimum set of permissions needed for the MongoDB DataProtect Solution to communicate with the source.

MongoDB User & Roles

The user specified in the MongoDB source registration page must have clusterAdmin and readWriteAnyDatabase roles. For example:

```
db.createUser(
{
user: "username",
pwd: "password",
roles : [
{ role: "clusterAdmin", db: "admin" },
{ role: "readWriteAnyDatabase", db: "admin" }
]
}
)
The user must exist in all replicas. For a sharded environment, the user must exist in all replicas across all shards.
```

For more information on the preceding roles, refer to the MongoDB documentation on Built-In Roles.

> **Note:** This Is applicable for LDAP authentication as well.

Best Practices

- Ensure that the oplog of MongoDB primary is sized to hold at least 30 minutes of data even when there is heavy load. If you notice that a full backup is occurring frequently, consider increasing the size of the oplogs.
- MongoDB Data Protection solution uses replica set oplog to capture incremental data. Standalone MongoDB clusters do not have an Oplog so they must be converted to replica sets using the process outlined in Convert a Standalone to a Replica Set.

Considerations

- Re-creating a collection between two successive backups (drop a collection and re-create a collection with the same name) is not supported. The Protection Group must be re-created in that case.

- MongoDB supports a special type of collection called 'Capped Collections'. A full backup is performed for capped collections in every run.
- Backup and recovery of MongoDB Views are not supported.
- Distributed transactions (introduced in MongoDB 4.2) are not supported.
- If a shard is added to a MongoDB cluster, the shard should not be removed before the next backup run. Removing the shard results in data loss and the user should re-create the workflow.
- If the shard key of a collection is modified using the refineCollectionShardKey command, the latest shard Key may not be recovered in PITR. PITR will recover the shardKey which was assigned to the collection at the time of the last incremental backup.
- If a collection is resharded, PITR is not supported after reshardCollection till the next backup. The next backup will automatically take a full backup of the resharded collection.
- The available PITR timeline starts from the oldest incremental or full backup run which is within the CDP retention period set in the policy.
- If you upgrade the NoSQL and Hadoop Service from a version older than 5.2, PITR will only be available after the next policy-based incremental run. To use PITR, the protection group policy should be modified to include CDP. For more information, see Using MongoDB CDP.
- Backup and recovery of collections with non-unique values in the _id field is not supported.
- Standalone MongoDB clusters are not supported because they do not have an Oplog. Standalone clusters must be converted to replica sets using the process outlined in the reference link - Convert a Standalone to a Replica Set.
- Meta operations like collection addition or deletion are captured in the backup run immediately following the operation. For this reason, the list of collections on a MongoDB database selected for point-in-time recovery is captured in the latest backup run before the selected time (called base backup). The point-in-time recovery recovers only those collections that are available in the base backup. If any such collection is deleted before the selected point in time, it is still recovered.
- Cohesity recommends taking a periodic full backup at least once per week or whenever more than 150 million documents are inserted, updated, or deleted in any of the collections being protected by the Protection Group. Additionally, in order to reclaim space, the interval between two full backups must be less than the retention period set in the policy.
- The support for MongoDB time series collections has the following considerations.

    - Sharded time series collections are not supported.
    - Secondary indexes on time series collections are not supported.
- The backup process might not work during MongoDB upgrades.
- The Pause Future Run function does not work for MongoDB protection groups if a run is in progress.
- Cohesity does not support registering a MongoDB cluster using a mix of IPv4 and IPv6 addresses.
- For IPv6, Kerberos authentication is not supported.

Register a MongoDB Source

27 July 2026

Register MongoDB as a Source

To configure MongoDB for data protection on Cohesity, you need to register it as a source.

Before you register MongoDB as a source, ensure that you have reviewed and completed the instructions provided in the Plan and Prepare for MongoDB Protection topic.

To register MongoDB as a source:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** on the top-right of the page and then select **Databases > MongoDB**. The **Register MongoDB** dialog is displayed.
3. Enter the required details. Refer the table below for the field descriptions.
4. Click **Register**.

If the registration is successful, the registered cluster appears in the list on the **Sources** page.

> **Note:** Registering a faulty source triggers an error. The user must Update MongoDB Configuration to correct the information and save it to resolve the errors.

Field Descriptions

Field	Description
MongoDB Seeds	Specify the MongoS hosts for a sharded cluster and the MongoD hosts for a non-sharded cluster. You can specify one host or multiple hosts in a comma separated list. It is recommended to specify a list of hosts in which at least one will be up at all times and MongoDB will discover the rest. Both IPv4 and IPv6 addresses are supported. **IPv4 Format**: `<IPv4 address>:<port_number>` **Example**: `10.14.58.2:27017,10.14.58.3:27017` **IPv6 Format**: `[<IPv6 address>]:<port_number>` **Example**: `[2400::9000:0:0:500:12]:27017, [2400::9000:0:0:500:13]:27017`
Authentication	Select one from the below authentication types: • None • LDAP • SCRAM • Kerberos
Username	Specifies the username of the MongoDB cluster. • In case of authentication type = LDAP, this will be the LDAP username/password. • In case of authentication type = SCRAM, this will be the MongoDB username/password. • No username/password required for Authentication type = NONE.
Kerberos Principal	In case of authentication type = Kerberos, specify the Kerberos principal to access the MongoDB primary.
Password	Specifies the password for the MongoDB cluster.
Authentication Database	Authenticating Database for the MongoDB cluster. Displayed only for LDAP and SCRAM authentications.
Backup from Primary node	Select to backup from the primary node.
Backup from Secondary Node	Select to backup from the secondary node.

Field	Description
Secondary Node Tag	The tag assigned to the intended MongoDB Secondary node. Required only if **Backup from Secondary Node** is selected. The system will use this to identify the secondary nodes for reading backup data. Exactly one secondary node in each replica set should have the tag. Once a unique tag is assigned to the secondary nodes that are planned to be used as a backup, the same need to be mentioned in the **Secondary Node Tag** field. For example: If the secondary nodes are marked with the tag: **"usefor": "cohesitybackup"**, then enter **usefor:cohesitybackup** in the **Secondary Node Tag** field. A tags document contains user-defined tag field and value pairs for the replica set member. For more information on how to configure tags, see Configure Replica Set Tag Sets.
SSL Requirement	Toggle **YES** if connection to MongoDB has to be over SSL.

Update MongoDB Configuration

You can update the MongoDB details that you provided during the registration process with the latest MongoDB configuration.

To edit the MongoDB Source Configuration:

1. Navigate to **Data Protection > Sources**.
2. Browse to the required MongoDB source you want to edit and click ⋮ > **Edit**. The **Edit MongoDB Registration** dialog is displayed.
3. Edit the required details. Refer Field Descriptions for more details.
4. To update the details, click **Update**.

Refresh the MongoDB Source

Refresh automatically happens in the background but you can also manually refresh the MongoDB source and fetch the latest changes on the source.

To refresh the MongoDB Source:

1. Navigate to **Data Protection > Sources**.
2. Click ⋮ next to the MongoDB source and select **Refresh**.

Unregister the MongoDB Source

If you want to stop taking backups of the MongoDB source, you can unregister the source from Cohesity.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the MongoDB source.

To unregister the MongoDB Source:

1. Navigate to **Data Protection > Sources**.
2. Click ⋮ next to the MongoDB source and select **Unregister**.

Backup MongoDB

27 July 2026

Cohesity provides a simple sequence of tasks that constitute the procedure for backing up your MongoDB source. Once MongoDB is registered as a source on Cohesity, you can start configuring your backups.

> **Note:** In case of the full backup, the Sharded Cluster Balancer is automatically disabled. Disabling the balancer for the duration of the backup procedure is a recommended practice. For more information, refer to the MongoDB documentation on Disabling Balancer During Backups.
> If the balancer is active when the full backup starts, then Cohesity disables the balancer. After the full backup is complete, Cohesity automatically reactivates the balancer if it was originally running when the full backup started. However, in a rare scenario, if MongoDB fails or if a running backup workflow is killed, the balancer does not get reactivated. In this case, it is recommended that users manually enable the balancer on the MongoDB cluster. Refer to the MongoDB documentation on Manage Sharded Cluster Balancer.

Continuous Data Protection (CDP)

The Cohesity Continuous Data Protection (CDP) provides near-zero RPOs for MongoDB environments. The CDP captures and protects the application data every time there is a change in data so that it can be restored to a point in time as opposed to periodic snapshots available with regular backups.

For MongoDB CDP, Cohesity uses a specialized CDP Service which runs on the Cohesity cluster. The CDP Service captures changes from MongoDB oplog in real time and continuously backs up the changes to the Cohesity cluster. These changes are saved on the Cohesity cluster in a write-ahead-log (WAL) which can be used to recover collections to any point-in-time as long as all the previous changes are available.

Cohesity manages rare instances of oplog rollover by doing a full backup instead of an incremental one.

Using MongoDB CDP

> **Note:** Ensure that the Prerequisites are met before using MongoDB CDP.

To enable CDP, create a new protection policy and attach the Continuous Data Protection option to the policy. Within the CDP option specify the retain point in time configuration, and associate this policy with the MongoDB source's Protection Group.

After associating the protection policy with the Protection Group, the protection run will first initiate a full backup. The initial full backup is followed by an incremental protection run and the continuous protection of the source will begin from this point on. You can view all the continuously protected sources from the Cohesity UI.

To create a new protection policy with the **Continuous Data Protection** option, see Create or Edit a Standard Policy.

Protect MongoDB Objects

MongoDB Objects are protected using a Protection Group and a Protection Policy. A Protection Group enables you to select the objects that you want to back up. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured, retained, archived, and replicated .

A Protection Policy is a reusable list of settings that define the type of backup, backup schedule and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

To protect a MongoDB object:

1. **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MongoDB**. The **New Protection** page is displayed.
3. Select the required **Data Source** from the drop-down menu.
4. Select the required objects to protect.
5. To auto protect an object, click adjacent to it. All the databases under this object are also auto protected.
6. Click **Save Selection**.
7. Under **Protection Group**, enter a protection group name. Existing protection groups cannot be used while creating a new MongoDB protection.

8. Under **Protection Policies**, select an existing protection policy or create a New Policy if the available policies do not meet your schedule and retention needs. For more information on creating policies with Continuous Data Protection, see Continuous Data Protection (CDP).
9. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains .
10. Enter the **Start time** and select the appropriate **Time Zone**. This option is available only if the selected policy is set to **Backup Daily** and indicates when the protection group should run. The default time zone is the browser's time zone.
11. After selecting a Storage Domain, if you need to change any of the additional settings on the MongoDB Protection page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
12. Click **Protect**.

Your new Protection Group is now active and running.

Additional Settings

Fields	Description
Storage Domain	Select default storage domain or add a storage domain.
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection group by selecting a different time zone.
Pause Future Runs	Once enabled, no protection runs will be scheduled.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Fields	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the cluster. By default, maximum scaling is capped at 16.
Bandwidth Throttling	Maximum network bandwidth that each backup stream can use for exchanging data with the cluster.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • Success: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts . You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Fields	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • Incremental: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • Full: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	Optional. Enter a brief description about the Protection Group.

Manage Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to an existing Protection Group

You can add new objects to the Protection Group created for MongoDB protection. The newly added objects are backed up from the next protection run.

To add objects to the Protection Group:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click ⋮ next to the Protection Group and select **Edit**.
4. On the Edit Protection page, Click ✎ next to **Objects**.
5. Make the required changes. For more information, refer Protect MongoDB Objects.
6. Click **Protect**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the objects that are added to the MongoDB source. You can auto protect either the entire cluster or a database.

When the Auto Protect option is enabled:

- All the objects that are added to the MongoDB source in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the objects that are removed from the MongoDB source in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backups are preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly added objects, in the MongoDB source, to the Protection Group.
- Remove existing objects, that are removed from the MongoDB source, from the Protection Group. The existing backups are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB. For more information, see Protect MongoDB Objects.
3. Click ⋮ next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** area, click **Edit**.
5. In the pop-up window , toggle the to enable or disable Auto Protect.
6. Click **Protect**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB. For more information, see Protect MongoDB Objects.
3. Click ⋮ next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Protect**.

Start an On-demand Protection run

You can manually start a run for the Protection Group that you have created to backup MongoDB objects. It provides you with options to:

- Back up all the objects in the Protection Group.
- Select an incremental or full backup.
- Add replication and archive.

To start a protection run:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click ⋮ next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. The **Backup** drop-down list is defaulted to **Backup all Objects in the Protection Group**.
    2. In the **Backup Type** drop-down list, select the backup type.
    3. If required add details for Replication and Archival.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group. However, if a protection run is in progress, it continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click ⋮ next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered.

To resume future protection run:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click ⋮ next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the snapshots (local, archived, and replicated) that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete the Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived snapshots generated by the Protection Group and reclaim the space utilized by the Protection Group.

> **Note:** The storage space reclaimed by deleting the Protection Group takes some time to reflect.

To delete the Protection Group:

1. **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click ⋮ next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover MongoDB

27 July 2026

You can recover the protected MongoDB Objects using Cohesity. You can recover to the same location or an alternate location.

Considerations

- Recovering from different protection groups in a single recovery is not supported for MongoDB.
- Recovering items in the same protection group but from different recovery points is not supported in the same recovery task. You need to create different recovery tasks for different recovery points.

Recover MongoDB Objects

To recover MongoDB objects you can search through the backups using the Protection Group, databases or collections you plan to recover, and then start the recovery.

To recover MongoDB objects:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Databases > MongoDB**. The **Databases | MongoDB** dialog is displayed.
3. In the **Recover** page, search for the MongoDB protection groups or objects that need to be recovered.

   > **Note:** Wildcards are supported in the search.

4. Accept the latest Recovery Point or click and select a different recovery point.

   By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page.

   To select a snapshot or a point in time from the Timeline mode:

   > **Note:** To select a point in time recovery, CDP should be enabled in the protection policy. For more information, see Using MongoDB CDP.

   a. From the **Choose a date** drop-down, select a date.

      The snapshots available on the selected date are displayed in blue dots on the timeline bar.

   b. For snapshot recovery, Select the snapshot to be recovered by clicking the corresponding blue dot.

   c. For Point in time recovery, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar or by entering the exact time in the Time input field. You can recover collections or databases to a point in time.

   d. By default, the snapshot is recovered from the local cluster. You can also select the snapshot archived in the cloud by clicking the cloud ( ) icon if available in the **Location** option.

   e. Click **Select Recovery Point**.

- Snapshots: Select a different recovery point. You can change the time zone before or after selecting a new restore time.

5. Click **Next: Recover Options**. The recovery details are displayed.
6. Under **Recover As** select **Replace Original** or **New Object**.
   - **Replace Original** - Replace the original object source.
   - **New Object** - Specify the target for the recovered object.
7. Toggle **Overwrite existing object with the same name**. A warning is displayed, enter **YES** to proceed.
8. Under **Recovery Options**, select the required options for the recovery. For more information, see Recovery Options.
9. Click **Recover** to start the recovery operation.

Recovery Options

Option	Description
Rename	Rename objects. Choose **Bulk Rename** to append a suffix to all the selected objects, or choose **Rename Individual Objects**.
Concurrency	Maximum number of recovery streams that will be created to exchange data with the cluster. By default, maximum scaling is capped at 8.
Bandwidth Throttling	Maximum network bandwidth that each recovery stream can use for exchanging data with the cluster.
Zones/Shard tags	If the toggle option Is enabled, the ZoneKeyRanges associated with the original collections will be preserved. **Note:** If this option is enabled, it is required that the association of a shard with the zone is already done in the destination MongoDB. Recovery may fail otherwise.
Task Name	Recovery task name.

MongoDB via Ops Manager

The Cohesity MongoDB Integrated Solution integrates with MongoDB Ops Manager to provide full backup, incremental backup and recovery for MongoDB deployments.

Core Components

The following are the core components for protecting MongoDB objects:

- Ops Manager
- MongoDB cluster

As illustrated, you can integrate Cohesity with your MongoDB deployments.

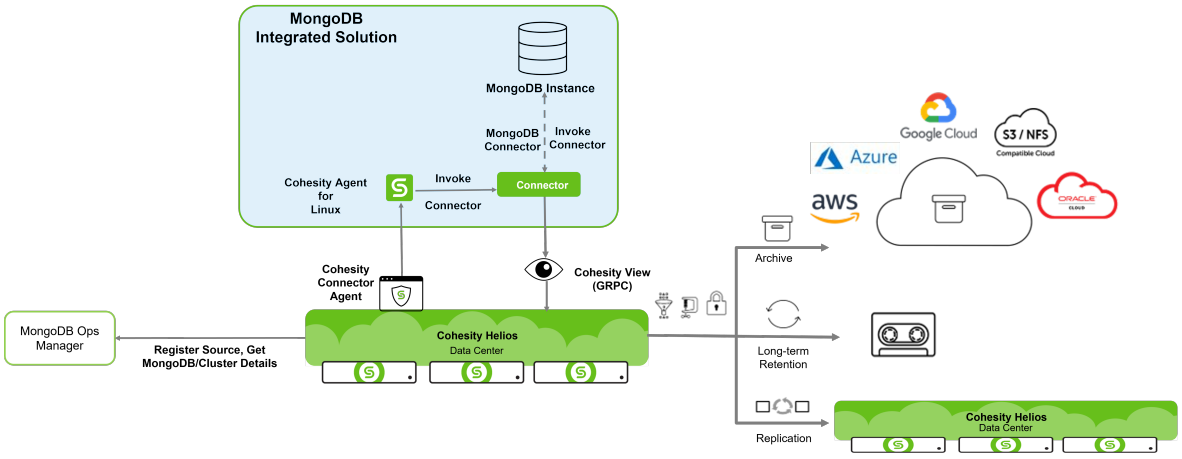

MongoDB Features

Following are the key MongoDB-specific features of the solution:

- Support replica sets and sharded clusters
- Automatically discover and protect all clusters deployed in a MongoDB OpsManager
- Backup and recovery - backup and recover individual clusters
- Schemas and indexes (including index data and index definitions) are backed up and recovered
- Option to back up from secondary, primary or preferred MongoDB nodes
- Support SSL-enabled MongoDB OpsManager
- Support Point In Time recovery using oplog backups
- Support OpsManager running in HA mode

Plan and Prepare for MongoDB Protection via Ops Manager

The following topics provide guidance that helps you get ready to implement the MongoDB DataProtect Solution via Ops Manager.

Supported MongoDB Versions

For more information, see MongoDB via Ops Manager.

Port Requirements

Ensure that you open the following ports in the firewall to allow data transfer between MongoDB and the Cohesity cluster:

Port	Source	Target	Direction	Network Protocol
50051	Cohesity	MongoDB Host	Omnidirectional	gRPC
8080/ 8443/ 80	Cohesity	Ops Manager	Cohesity to Ops Manager	HTTP/ HTTPS
11121	MongoDB Host	Cohesity	Omnidirectional	gRPC
11117	MongoDB Host	Cohesity	Omnidirectional	gRPC
11113	MongoDB Host	Cohesity	Omnidirectional	Insecure gRPC

Minimum Permissions

To use the Cohesity Connector for MongoDB backup, the Cohesity agent and MongoDB connector executables must be installed and run on each MongoDB host node.

Both the Cohesity agent and the MongoDB connector come as Installer RPM packages and need to be installed as root. Based on the options selected during installation, the service can run in one of the following modes:

1. Default: Run as a user with User ID 0 (root privileges).
2. Run as a service user (non-root) with sudo privileges.
3. Run as a service user (non-root) configured with specific capabilities using setcap (use this if sudo/root access is not permitted).

> **Note:** The sudo package is a required dependency for the Cohesity Linux Agent and must remain installed on the host even when the agent is configured to run as a non-root user without sudo privileges. For the full list of dependencies, see Install Cohesity Linux Agent.

Relevant command line options must be set while installing the RPMs. For more details, refer to Install Cohesity Linux Agent RPM Package.

Considerations

Review and understand the following recommendations before protecting your MongoDB cluster:

- MongoDB recommends that if any node in the cluster goes down, the backup process fails, requiring a full backup to be triggered again.
- MongoDB strongly recommends scheduling periodic full backups at least once per week to prevent data inconsistency scenarios.
- MongoDB backup protection jobs might fail during high-ingestion periods with significant application writes.
- The Cohesity solution does not support the use of Organization-level API keys. This is because the solution is designed to protect the entire OpsManager instance rather than individual organizations. You must generate and use a Global API Key for successful registration and integration.
- MongoDB recomments not running a resync on any node while a backup is active, as this will cause the snapshot and backup workflow to fail.
- Do not perform a primary stepdown during an active backup, as it may result in snapshot failure.
- The amount of data backed up in an incremental backup may not exactly match the data ingested in the source during that period. This amount depends on multiple factors such as source side compaction, the distribution of the changes, etc.
- Standalone MongoDB deployments are not supported and are filtered out by Cohesity.
- Multiple concurrent operations within the same project are not allowed. For example, if a project has two clusters, A and B, and a recovery operation is in progress on cluster A, then any backup or recovery operations for cluster B must wait until cluster A's recovery is complete.
- Non-ASCII or special characters in dbPath are not supported.
- The in-memory storage engine is not supported.
- MongoDB deployment configuration options must be the same for both the source (backup) and destination (recovery) clusters.
- Before recovery, make sure that the destination hosts have adequate storage available.
- If a protection group is edited to add a new collection or a new cluster gets protected as part of autoprotect, the immediate log backup for those collections is not enforced. Log backups for them would run at the next policy-defined time, and the restore timeline would be available from the first incremental backup that happens after at least one log backup is taken following a recovery.
- Once an Ops Manager source is created, the hostname and port cannot be edited. However, the authentication details can still be edited.

Download and Install Cohesity Agent and Connector

Install the Cohesity Agent on each MongoDB node in the cluster you want to protect.

Install Cohesity Linux Agent

The Cohesity Linux Agent is available in the bundle with support for RHEL based Linux distributions. You will need to install the RPM package for your Linux distribution.

> **Note:** Currently, only RHEL8 and RHEL9 are supported.

The installer packages and Linux distributions on which the installer package is supported are:

Installer Package	Linux Distribution
(Default) RPM	RHEL and its derivatives

The Cohesity Linux Agent has dependencies on the following packages, which must be installed on the Linux server:

Command/Package	RHEL and its derivatives
rsync	rsync
mount	util-linux
lvm2	lvm2
sudo	sudo
coreutils	coreutils
util-linux	util-linux
nfs client	nfs-utils
lsof	lsof
wget	wget

Install Cohesity Linux Agent RPM Package

1. Download the Cohesity Linux agent package (el-cohesity-agent) from the Cohesity Download portal and copy it to all MongoDB nodes.

2. Choose the execution mode. The Cohesity Linux agent supports the following execution modes:
   - Option 1: Run as a user with User ID 0 (root user).
   - Option 2: Run as a service user (non-root) with sudo privileges.
   - Option 3: Run as a service user (non-root) configured with specific capabilities using setcap.

   > **Note:**
   > - In all execution modes, the RPM installation must be done as the root user.
   > - The sudo package is a required dependency for the Cohesity Linux Agent and must remain installed on the host even when the agent is configured to run as a non-root user without sudo privileges.

   Option 1: For running the Linux agent as the root user, install using the following command:

   ```
   For example:
   [root@hostname.example.com ~]$ rpm -ivh el-cohesity-agent-xx-
   1.x86_64.rpm
   ```

   Option 2: For running the Linux agent as a non-root service user with sudo access, install using the following command:

   ```
   For example:
   [root@hostname.example.com ~]$ export COHESITYUSER=<service-user>
   ; ; rpm -ivh el-cohesity-agent-xx-1.x86_64.rpm
   ```

   > **Note:**
   > - Ensure the service user is an existing user on the host.
   > - In addition, it is necessary to create and preserve the following settings in the /etc/sudoers file for the service user:
   >
   >   ```
   >   <username> ALL=(ALL) NOPASSWD:ALL
   >   Defaults:<username> !requiretty
   >   ```

   Option 3: For running the Linux agent as a non-root service user without sudo access, install using the following commands:

```
For example:
[root@hostname.example.com ~]$ export COHESITYUSER=<service-user>
; export COHESITY_AGENT_SKIP_ALL_NO_SUDO=1 ; rpm -ivh el-cohesity-
agent-xx-1.x86_64.rpm
```

> **Note:**
> - Ensure the service user is an existing user on the host.
> - Setting capabilities using setcap is a one-time activity and is performed only during RPM installation or upgrade scenarios. Once capabilities are set using setcap, during regular execution of agent and connector, sudo access is not required.

Install Cohesity MongoDB Connector RPM package

1. Download the Cohesity MongoDB connector RPM package (e.g. cohesity-mongodb-connector-XXX-1.x86_64.rpm) from the Cohesity Download portal and copy it to all MongoDB host nodes.
2. Install the MongoDB connector RPM on each node of MongoDB cluster using the following command:

   ```
   [root@hostname.example.com ~]$ rpm -ivh cohesity-mongodb-connector-XXX-
   1.x86_64.rpm
   ```

Prerequisites

The following configurations are necessary prerequisites for the protection and recovery of MongoDB deployments via Ops Manager:

- Enable third-party backup management for the relevant projects in the Ops Manager.
- Enable the Ops Manager Backup daemon on the admin page.
- Install and configure the MongoDB Ops Manager agent on each server in the deployment to be protected. Also, make sure that the agent has read/write access to any of the configured directories. Make sure to configure the Oplog export path for MongoDB to store the Oplog Data.
- Synchronize system clocks to ensure consistent timestamps across machines using Network Time Protocol.
- Activate monitoring and backup services on the MongoDB deployment to be protected. Make sure the Monitoring, Backup, and automation are enabled on the Ops Manager agent for each host of the MongoDB Deployment to be protected.
- Create a Global API key with Global Backup Admin, Global User Admin, and Global Owner permissions if not already present. Make a note of the public and private keys,

as these credentials will be needed for registering MongoDB Ops Manager on Cohesity.

For more details, refer to MongoDB product documentation.

Configure VIP's & FQDN

Ensure that Cohesity VIPs are configured on the Cohesity Cluster and that an FQDN is assigned and resolvable from all MongoDB Replica Set sources, Ops Manager clusters, and Load Balancer nodes.

Register MongoDB Ops Manager Source

To protect your MongoDB objects, you need to register your MongoDB Ops Manager deployment as a source on the Cohesity cluster.

To register MongoDB Ops Manager as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Databases** > **MongoDB Ops Manager**.
3. On the **Register MongoDB Ops Manager** page, perform the following:
    a. In the **Hostname** field, enter the hostname.
    b. In the **Port** field, enter the port number configured for your Ops Manager instance.
    c. Enable **SSL Requirement** to ensure a secure connection between the Cohesity cluster and the MongoDB Ops Manager deployment.
    d. If the Ops Manager server uses custom CA, add the custom CA certificate in the text box or upload a local copy of the CA cert in pem format using the upload option.
    e. Enter the **Public Key** and **Private Key** for the Ops Manager API keys.
    f. Click **Register**.

Update MongoDB Ops Manager Source Configuration

After registering MongoDB as a source on the Cohesity cluster, you can update the MongoDB source configuration.

To update the MongoDB source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the MongoDB source and select **Edit**.
3. On the **Edit MongoDB Ops Manager** page, update the required configurations.

> **Note:** The host name and port cannot be changed.

4. Click **Update**.

Manage Agents – View Agent and Connector Installation Status for MongoDB Nodes

To protect your MongoDB clusters using Cohesity through Ops Manager, every MongoDB node must have both the Cohesity Agent and the Cohesity Connector installed. These components work together to enable application-consistent backups and ensure your data is fully protected.

The Manage Agents flow can be used to verify whether all the necessary components are installed across all nodes. Missing even one Agent or Connector on a node results in the following:

- Backups cannot run successfully
- Some parts of the cluster remain unprotected
- Troubleshooting becomes time-consuming

This feature provides complete visibility into the installation status of both the Cohesity Agent and Connector across all MongoDB hosts in Ops Manager.

With this feature, you can quickly:

- See which MongoDB hosts have both components installed
- Identify MongoDB hosts missing the Agent or Connector
- Check the version of the installed Cohesity Agent
- View health status for each node
- Filter large environments to find exactly what you need

This helps ensure that your MongoDB clusters are always fully prepared for successful backup operations.

To view the Cohesity Agent and Connector installation status:

1. Navigate to **Data Protection** > **Sources**.
2. Select the MongoDB source you want to review.
3. Click on the actions menu (⋮) and select **Manage Agents**.

   The **Manage Agents** dialog appears and displays all discovered hosts and their installation status. You can view the following information:

   1. **Hostname** - Displays the fully qualified domain name (FQDN) of the MongoDB host.

2. **Agent Version** - Displays the currently installed version of the Cohesity Agent on the host.
3. **Status** - Displays the installation status of the Agent and Connector. The available statuses are:
    1. **Healthy** - Both the Cohesity Agent and Connector are installed and operational.
    2. **Not Installed** - Neither the Agent nor the Connector is installed on the host.
    3. **Unhealthy** - At least one of the required components is missing.

   To view additional details about the status, click the info (i) icon next to the hostname.

Use Filters to Find the Right Hosts

Filters help you navigate environments with many organizations, projects, or clusters. You can filter by:

- **Organization** - Filters the hosts based on the Ops Manager organization to which they belong.
- **Project** - Displays hosts associated with a specific project within an organization.
- **Cluster** – Displays hosts belonging to a specific MongoDB cluster.
- **Agent Version** - Filters based on the version of the Cohesity Agent installed on the hosts.
- **Status** - Filters hosts based on the health or installation status of the Agent and Connector.

Last refreshed - Shows when the data was last updated.

Unregister MongoDB Ops Manager Source

If you plan to stop protecting your MongoDB objects, you can unregister the MongoDB source from the Cohesity cluster.

To unregister the MongoDB source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the MongoDB source and select **Unregister**.

> **Note:** For instructions on performing multi-cluster registration of the MongoDB Ops Manager source, contact Cohesity Support.

Backup MongoDB via Ops Manager

Cohesity provides a simple workflow to back up your MongoDB source. After registering MongoDB Ops Manager source on Cohesity, you can start configuring your backups.

Protect MongoDB Objects via Ops Manager

MongoDB deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects that you want to back up.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy .

To protect MongoDB deployments:

1. Select **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Databases** > **MongoDB via Ops Manager**.

   The **New Protection** dialog appears.
2. Under **Add Objects**, from the **Registered Source** drop-down list, select the registered MongoDB Ops Manager source.
3. Select the required objects to protect.
4. To auto protect an object, click adjacent to it. All the deployments under this object are also auto protected.
5. Click **Continue**.
6. Under **Protection Group**, enter a protection group name. Existing protection groups cannot be used while creating a new MongoDB protection.
7. Under **Policy**, select an existing protection policy or create a new policy if the available policies do not meet your schedule and retention needs.
   1. To enable oplog backup, select the policy that specifies the Log backup. Enabling oplog backup is a requirement for point-in-time recovery.
   2. As per MongoDB requirements, ensure that the policy has periodic full backup enabled with a frequency of at least once a week.
8. Click **More options** to provide additional settings.
9. Click **Protect**.

Your new Protection Group is now active and running.

Additional Settings

Fields	Description
Storage Domain	Select the default storage domain or add a storage domain.
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection group by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • Incremental: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • Full: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Pause Future Runs	Once enabled, no protection runs will be scheduled.
Convert to Full Backup on Node Failure	Enabling this option ensures that if the same node cannot be selected for an incremental backup, it is automatically converted to a full backup.
Preferred Backup Node	Choose your backup preference: Prioritize the primary or secondary node with automatic fallback, or restrict backups exclusively to the secondary node. • **Primary Preferred:** Backups prioritize the primary node but switch to the secondary node if needed. • **Secondary Preferred:** Backups prioritize the secondary node but switch to the primary node if needed. • **Secondary Only:** Backups are restricted exclusively to the secondary node, with no fallback to the primary node.
Preferred Backup Nodes List	To enable, click the I**Include Preferred backup node list** toggle. In the text box, enter a comma-separated list of host names to be used as preferred backup sources, in order of priority. These host names must exactly match those shown in the cluster status in Ops Manager. Once enabled, only the specified nodes will be considered for backups.

Fields	Description
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • Success: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Fields	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.
Description	Optional. Enter a brief description about the Protection Group.

Manage Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the objects that are added to the MongoDB source. You can auto-protect either the entire ops manager instance, the entire organization or an ops manager project.

When the Auto Protect option is enabled:

- All the objects that are added to the MongoDB Ops Manager entity in the future are automatically added to the Protection Group and are protected from the next protection run.

  > **Note:** The Cohesity agent must be installed on all hosts in the newly added clusters.

- All the objects that are removed from the MongoDB source in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backups are preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly added objects, in the MongoDB source, to the Protection Group.
- Remove existing objects, that are removed from the MongoDB source, from the Protection Group. The existing backups are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** area, click **Edit**.
5. In the pop-up window , toggle to enable or disable Auto Protect.
6. Click **Protect**.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Start On-demand Protection run

You can manually start a run for the Protection Group that you have created to back up MongoDB objects. It provides you with options to:

- Back up all the objects in the Protection Group.
- Select an incremental or full backup.
- Add replication and archive.

To start a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. The **Backup** drop-down list is defaulted to **Backup all Objects in the Protection Group**.
    2. In the **Backup Type** drop-down list, select the backup type.

3. If required add details for Replication and Archival.

5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group. However, if a protection run is in progress, it continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Delete Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the snapshots (local, archived, and replicated) that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete the Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived snapshots generated by the Protection Group and reclaim the space utilized by the Protection Group.

> **Note:** The storage space reclaimed by deleting the Protection Group takes some time to reflect.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MongoDB.
3. Click (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog, select any one of the following options:
    1. Delete Protection Group Only
    2. Delete Protection Group and Snapshots

Source-Side Deduplication for MongoDB

The source-side deduplication feature for MongoDB is designed to optimize backup efficiency by eliminating redundant data at the MongoDB source before transferring it to the Cohesity cluster. This optimization significantly reduces network bandwidth consumption, shortens backup windows, and optimizes storage consumption on the target, making it especially beneficial for large MongoDB deployments.

When source-side deduplication is enabled, only unique data blocks that have not been previously protected are transmitted to the Cohesity cluster. This capability integrates with Cohesity's built-in deduplication to further enhance storage efficiency by processing only new data blocks.

Advantages of Source-Side Deduplication

- **Reduced Network Traffic**: Minimizes data transferred between MongoDB hosts and the Cohesity cluster, helping conserve network bandwidth.
- **Faster Backups**: Shortens backup windows, particularly for large MongoDB databases that require periodic full backups.
- **Optimized Storage Consumption**: Reduces upfront storage usage on the Cohesity cluster by sending only unique data blocks.
- **Improved Scalability**: Offloads deduplication processing from the Cohesity cluster, enabling the system to support a higher number of concurrent backup workloads.

Enable Source-Side Deduplication for MongoDB

To enable source-side deduplication for MongoDB, contact Cohesity Support.

> **Note:** Source-side deduplication is effective only when deduplication is enabled on the Storage Domain associated with the Protection Group. If deduplication is disabled on the Storage Domain, source-side deduplication does not take effect.

Recover MongoDB via Ops Manager

You can recover the protected MongoDB Objects using Cohesity.

Recover MongoDB Objects

To recover MongoDB objects, you can search through the backups for a specific MongoDB cluster using filters such as Ops Manager source and Protection Group name. Once located, you can start the recovery process.

To recover MongoDB objects:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Databases** > **MongoDB via Ops Manager**. The **Clusters | MongoDB** dialog is displayed.
3. On the **Recover** page, search for the MongoDB protection groups or sources that need to be recovered.

   > **Note:** Wildcards are supported in the search.

4. Accept the latest Recovery Point or click ✎ and select a different recovery point.

   By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page.

   To select a snapshot or a point in time from the Timeline mode:

   > **Note:** To select a point in time recovery, log backup must be enabled in the protection policy.

   a. From the **Choose a date** drop-down, select a date.

      The snapshots available on the selected date are displayed in blue dots on the timeline bar.

   b. For snapshot recovery, Select the snapshot to be recovered by clicking the corresponding blue dot.

   c. For Point in time recovery, select the point-in-time restore. You can specify the exact point-in-time recovery by selecting a point the timeline bar or by entering the exact time in the Time input field.

   d. Click **Select Recovery Point**.

   - Snapshots: Select a different recovery point. You can change the time zone before or after selecting a new recovery time.

5. Click **Next: Recover Options**. The recovery details are displayed.
6. Under **Recover As** select **Replace Original** or **New Object**.

- **Replace Original** - Replace the original object source.
- **New Object** - Specify the target for the recovered object.

7. Under **Recovery Options**, select the required options for the recovery. For more information, see Recovery Options.
8. Click **Start Recovery** to start the recovery operation.

Recovery options

Field	Description
Staging Directory Path	Specify the full directory path where oplog files will be temporarily stored during point-in-time recovery. This path must exist and be writable by the Cohesity service user on every host in the destination MongoDB deployment. Ensure sufficient disk space is available. The staging directory must have enough free space to store the oplog data generated since the earliest snapshot backup within the selected point-in-time recovery (PITR) window. You can estimate the required space by multiplying the average data change rate by the duration of the PITR window. Note that the change rate may vary throughout the day due to spikes in data ingestion. For example, if the change rate is 1 MB/s and the PITR window spans 6,000 seconds, you will need at least 6 GB of free space. Cohesity recommends allocating additional space as a buffer to prevent potential issues. (This option is available only if a specific point in time is selected for recovery).
Task Name	Specify a name for the recovery task. A default name is pre-populated, which can be changed.

Cancel a Recovery Run

When you cancel a recovery run, MongoDB cancels the job by terminating the recovery processes.

To cancel a recovery run:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click the action icon (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
   - **Continue Recovery** if you want to continue to recovery run.
   - **Cancel Recovery** if you want to cancel the recovery run.

Pinecone

Cohesity Pinecone protection enables secure, policy-driven backup and recovery for Pinecone vector database indexes, which are commonly used in AI and machine learning workloads such as Retrieval-Augmented Generation (RAG), semantic search, and recommendation systems.

Cohesity integrates directly with Pinecone's native Backups API to protect serverless indexes without moving customer data out of the Pinecone cloud. Only backup metadata is stored on the Cohesity cluster, ensuring zero data movement and avoiding data egress costs.

Pinecone protection is implemented through a dedicated Universal Data Adapter v2 (UDAv2) adapter that communicates securely with Pinecone over HTTPS and integrates with existing Cohesity services and workflows. The result is a unified approach to protecting AI vector data alongside other enterprise workloads.

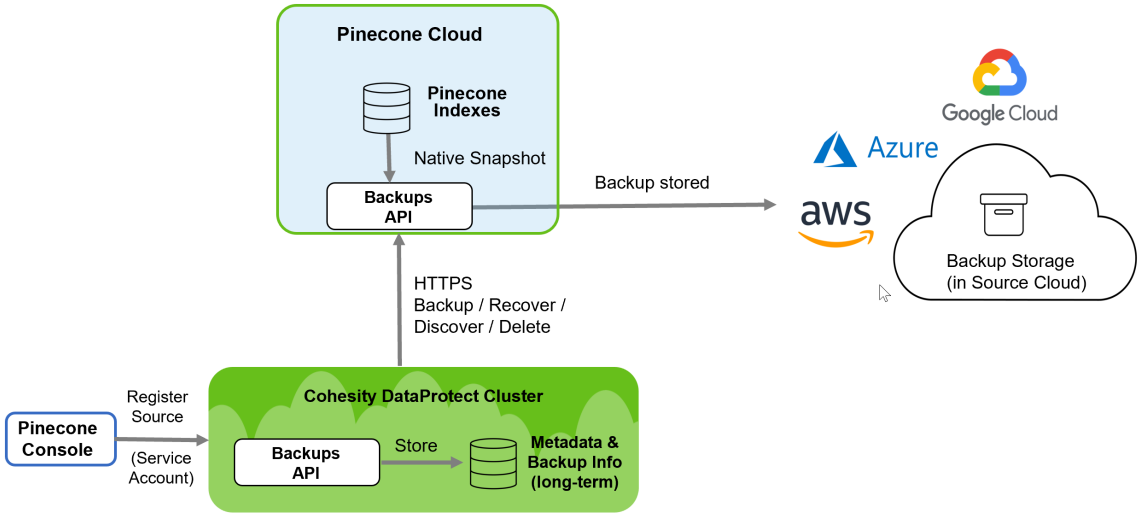

Pinecone Protection Core Components

The following are the core components for protecting Pinecone indexes:

- Pinecone cloud - external service where vector index data is stored, and native backup snapshots are created.
- Cohesity cluster - backup metadata such as snapshot identifiers and index details are stored.

Pinecone Features

Following are the key characteristics of Pinecone protection:

- **Backup Manager Style**: Cohesity does not move data but orchestrates the Pinecone's native backup mechanism of using Backups API for serverless indexes.
- **Zero Data Movement**: Backups remain stored in Pinecone's infrastructure.
- **Metadata Management**: Cohesity tracks backup metadata and manages lifecycle/retention policies.

- **Index Type Detection**: Automatically detects serverless indexes and uses the appropriate backup and recovery method.

Plan and Prepare for Pinecone Protection

To protect Pinecone with the Cohesity cluster, you must register your Pinecone account/organization as a source in the Cohesity cluster.

Prerequisites for Pinecone Protection

Following are the prerequisites for Pinecone Protection:

- Ensure that the service accounts have enterprise-level Pinecone subscription.
- Ensure outbound connectivity is supported for ***.pinecone.io** via HTTPS.

Considerations

- **Unsupported Functionalities** - The following functionalities are not supported for Pinecone:
  - Replication and archival of Pinecone snapshots
  - Auto-protection and Tag-based Auto-protection
- **Full back up support** - Only full backup is supported; incremental backup is not available for Pinecone.
- **New Location Recovery** - Recovering the Pinecone indexes to a new location are not supported.

Register and Manage Pinecone Source

To protect your Pinecone objects, you need to register Pinecone as a source on the Cohesity cluster. Pinecone organizations are registered using OAuth2 service account credentials.

> **Note:** The service account must have owner-level permissions to discover all the projects and indexes.

To register Pinecone as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Databases > Pinecone**.
3. On the **Register Pinecone** page, perform the following:
   1. In the **Client ID** field, enter the Pinecone OAuth client identifier used by Cohesity to authenticate API requests.

2. In the **Client Secret** field, enter the Pinecone OAuth client secret key used by Cohesity to authenticate API requests.
3. Click **Register**.

Update Pinecone Source Configuration

After registering Pinecone as a source on the Cohesity cluster, you can update the Pinecone source configuration.

To update the Pinecone source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Pinecone source and select **Edit**.
3. On the **Edit Pinecone** page, update the required configurations.

> **Note:** The Client ID cannot be changed.

4. Click **Update**.

Unregister Pinecone Source

If you plan to stop protecting your Pinecone objects, you can unregister the Pinecone source from the Cohesity cluster.

1. To unregister the Pinecone source:
2. Navigate to **Data Protection > Sources**.
3. Click the actions menu (⋮) next to the Pinecone source and select **Unregister**.

Backup Pinecone Indexes

After registering Pinecone source on Cohesity DataProtect, Cohesity automatically discovers all the organizations that the service account can access, as well as all supported workloads such as projects and indexes within those organizations. The discovered workloads appear under **Data Protection > Sources > Pinecone**.

Hierarchy View displays the entity hierarchy of the Pinecone objects, allowing you to navigate from the organization-level down to individual indexes.

Pinecone Entity Hierarchy

The Pinecone entity hierarchy follows this structure:

Structure	Description
Organization	Pinecone organization associated with the service account.
Project	Logical container for organizing and managing related Pinecone indexes.
Index	Individual entities that can be protected. You can view the index metadata details like type, project, index type, dimension, metric, record count, region, cloud, and size.

Supported Backup Types

Cohesity supports only full backups for Pinecone. Incremental backups are not supported.

Protect Pinecone Indexes

Pinecone deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects that you want to back up.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy .

To protect Pinecone deployments:

1. Select **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Databases** > **Pinecone**.

   The **New Protection** dialog appears.
2. Under **Add Objects**, from the **Registered Source** drop-down list, select the registered Pinecone source.
3. Select the required objects to be protected.
4. Click **Continue**.
5. Under **Protection Group**, enter a protection group name. Existing protection groups cannot be used while creating a new Pinecone protection.
6. Under **Policy**, select an existing protection policy or create a new policy if the available policies do not meet your schedule and retention needs.

   As per Pinecone requirements, ensure that the policy has periodic full backup enabled with a frequency of at least once a week.
7. Click **More options** to provide additional settings.
8. Click **Protect**.

Your new Protection Group is now active and running.

Additional Settings

Fields	Description
Storage Domain	Select the default storage domain or add a storage domain.
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection group by selecting a different time zone.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Pause Future Runs	Once enabled, no protection runs will be scheduled.
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Fields	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Fields	Description
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • Success: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.
Description	Optional. Enter a brief description about the Protection Group.

Manage Pinecone Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Start On-demand Protection run

You can manually start a run for the Protection Group that you have created to back up Pinecone objects. It provides you with options to:

- Back up all the objects in the Protection Group.
- Select an incremental or full backup.

- Add replication and archive.

To start a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you created for Pinecone.
3. Click (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. The **Backup** drop-down list is defaulted to **Backup all Objects in the Protection Group**.
    2. In the **Backup Type** drop-down list, select the backup type.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group. However, if a protection run is in progress, it continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Pinecone.
3. Click (⋮) next to the Protection Group and select **Pause Future Runs**.
4. In the **Pause Future Runs** dialog, select **Add Note** and enter the note.
5. Click **Confirm**.

Resume Paused Protection Runs

You can resume a paused protection run.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Delete Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the snapshots (local, archived, and replicated) that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete the Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved untill the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the snapshots generated by the Protection Group and reclaim the space utilized by the Protection Group.

> **Note:** The storage space reclaimed by deleting the Protection Group takes some time to reflect.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Pinecone.
3. Click (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog, perform the following:
    1. Select **Delete Snapshots** to delete the snapshots along with the Protection Group.
    2. Type **YES** to confirm.
    3. Click **Delete Group**.

> **Note:** Deleting all Snapshots will delete all the storage snapshots associated with the protection group. Any Cloudspin artifacts generated by the Protection Group must be deleted from the respective cloud portals. Snapshots under DataLock cannot be deleted.

Recover Pinecone Indexes

After you protect your Pinecone indexes, you can recover them from Cohesity DataProtect to their original location.

To recover the Pinecone indexes, you can search through the backups for a specific Pinecone cluster using filters such as Source and Protection Group name. Once located, you can start the recovery process.

To recover Pinecone indexes:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Databases** > **Pinecone**.

3. On the **New Recovery** page, search for the Pinecone protection groups or sources that need to be recovered.

   > **Note:** Wildcards are supported in the search.

4. You can click and select a different recovery point.

   To select a snapshot:

   a. From the **Date** drop-down, select a date. The snapshots available on the selected date are displayed.

   b. For snapshot recovery, select the snapshot to be recovered from the list.

   c. Click **Select Recovery Point**.

5. Click **Next: Recover Options**. The recovery details are displayed.
6. Under **Recover To** select **Original Location** to recover the Pinecone indexes to the original location. This will replace any existing index with the same name. Toggle On **Overwrite existing indexes** to allow Cohesity to overwrite the existing indexes in the target location. Type **YES** to confirm.

   > **Note:** Recovering the Pinecone indexes to a new location is not supported.

7. Select your **Recovery Options**:
   - **Rename**. Add a Prefix and/or Suffix to the full name or rename the recovered Pinecone index.
   - **Task Name**. Change the default name of the recovery task.
8. Click **Recover** to start the recovery operation.

Cancel a Recovery Run

When you cancel a recovery run, Pinecone cancels the job by terminating the recovery processes.

To cancel a recovery run:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click the action icon (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
   - **Continue Recovery** if you want to continue to recovery run.
   - **Cancel Recovery** if you want to cancel the recovery run.

Microsoft SQL Server

This chapter includes the following topics:

Planning and Preparation for Microsoft SQL Server Protection

The following topics provide guidance that helps you get ready to implement MS SQL Server:

Requirements for Microsoft SQL Server Protection

Supported Software for Microsoft SQL Server Protection

For more information, see Supported Software.

Agent Minimum Permissions for Microsoft SQL Server Protection

You can use either the host LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

It is recommended to run the Cohesity Agent service with an Active Directory domain user account that is a member of the local administrator of the SQL Server host. The AD domain user account must be a member of the SQL sysadmin server role. The user account must have log-on rights to the SQL Server host in the local security policy of the SQL Server host.

If you do not use the LOCAL SYSTEM account, ensure the following for the chosen account:

- The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
- The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.
- The account must have the sysadmin role in the MS SQL Server instance.

Ports Used for Communication

- On physical servers or VMs with an ephemeral or installed agent, open the ports 445, 11113, 11117, and 50051. For more information, see Manage Firewall Ports.
- If the Windows Firewall is used:
  - Inbound rules:
    - Add a rule to accept SQL Server traffic and TCP connections on local port 1433.
    - Set Remote Port to All Ports.
  - Outbound rules (for MS SQL Server 2016 running on Windows 2016):

    Update the "Block network access for R local user accounts in SQL server instance MSSQLSERVER" rule by going to General tab > Action window > select "Allow the connection".

Prerequisites for Microsoft SQL Server Protection

- Verify MS SQL Server services are running.
- The **Readable Secondary** field in the Availability Group Properties of the SQL Server Availability Group (AAG) must be set to **Yes** if you have selected **Secondary Only** as the **Backup Preference setting for AAG**.
- MS SQL database and instance names must meet the Microsoft SQL Server naming requirements. Names containing illegal Windows file characters, such as double quotes, pipe symbols, greater than and less than signs (" | > <), a space at the beginning or end of the database name (applicable only for VDI-based backup) are not supported. For more information, see Naming Files, Paths, and Namespaces.
- On the Server's Windows system, Cohesity recommends to set the Power Plan to High performance.

Credentials and Privileges for Microsoft SQL Server Protection

There are three accounts you must consider when installing the Cohesity Agent:

- **Installation Account**: The account you use to log in to the host and run the installer.
- **Service Account**: The account under which the Cohesity Agent service runs on the SQL Server host. (Selected at the time of installation.)
- **SQL Server login account**: The SQL Server account by which the Cohesity Agent has access to the databases. (Configured after installation.)

Installation Account

When installing the Cohesity Agent on a Microsoft Windows host, you must log onto the SQL Server host with an account that has local administrator privileges. This is because the Windows agent installer installs the Cohesity Agent service. Microsoft requires the installation account to have administrator privileges on the host.

Choose the Service Account

The installer gives you the option between two types of service accounts:

- The LocalSystem account.
- A user account. Assigning privileges to the service account will differ depending on which type of account you choose. Cohesity recommends using either the LocalSystem account or an Active Directory Domain User account. See the chart below for a comparison of service account types.

Comparison of Service Account Types

SERVICE ACCOUNT	TYPE	DEFINITION	RECOMMENDATION
LocalSystem (Default)	LocalSystem Account (Built-in Host Account)	The LocalSystem account is a predefined local account on the host computer.	Recommended: The Local System account is a predefined local account used by the service control manager. It has extensive privileges on the local computer, and acts as the computer on the network. Its token includes the NT AUTHORITY\SYSTEM and BUILTIN\Administrators SIDs; these accounts have access to most system objects.
User Account	Domain User Account (Active Directory)	A user whose username and password are stored in Active Directory.	Recommended: A domain user account enables the service to take full advantage of the service security features of Windows and Microsoft Active Directory Domain Services. The service has whatever local and network access is granted to the account, or to any groups of which the account is a member.
	Local User Account (Host account)	In Windows, a local user is one whose username and encrypted password are stored locally, only on the computer itself.	Not recommended: When you log in as a local user, the computer checks its own list of users and its own password file to see if you are allowed to log in to the computer. The computer itself then applies all the permissions and restrictions that are assigned to you for that computer.

> **Note:** If you are unsure which account to use, choose the LocalSystem account. This account already has most of the required permissions. The agent service account can be changed later if needed.

Microsoft SQL Server on Linux OS

Cohesity provides backup and recovery solutions for Microsoft SQL Server databases running on the Linux Operating System.

Requirements

Operating System	Database	Notes
Linux OS RHEL 8 and higher versions **Note:** This is an Early Access feature. Contact Cohesity Support to enable the feature.	Microsoft SQL Server 2019	• Only VDI-based backups are currently supported on Linux OS. • High Availability Microsoft SQL Server configurations like FCI and AG are not supported.

Considerations

- Cross-platform database restores between Windows and Linux are currently not supported.
- For VDI-based backups to work, Microsoft requires the backup application account (Cohesity Agent's service account) and the Microsoft SQL Server service account to be added to each other's primary groups using the usermod command.

  For example:

  sudo usermod -a -G mssql cohesityagent

  sudo usermod -a -G cohesityagent mssql

  where, mssql is the user account of the mssql-server service and cohesityagent is the user account of the Cohesity Agent service.

  Following this, the Microsoft SQL Server and the Cohesity Agent services must be restarted for the changes to take effect.

  > **Note:** The usermod -a -G command may not work for domain accounts.

- Currently, only password-based SQL Server authentication is supported and Active Directory-based authentication is not supported for Linux.

  For supported features of SQL Server 2019 on Linux, see Editions and supported features of SQL Server 2019 on Linux.

Register SQL Server on Linux OS

1. Navigate to **Data Protection > Sources**.
2. Click **Register > Microsoft SQL Server**.
   The **Register a Microsoft SQL Server** page is displayed.
3. In the **Hostname or IP Address** field, enter the IP address or FQDN of the server.
   We recommend you to provide the FQDN of the server.
4. Select the **Server Type** as Linux.
5. Provide the username and password for the **Database Authentication**.
6. Click **Register**.

Supported Database Types for Microsoft SQL Server

Before proceeding to deploy the Cohesity agent for SQL data protection, we recommend you correctly identify the type of SQL database you want to protect. Below is a table of SQL database types that are supported by the Cohesity SQL agent, and the recommended protection methods.

To understand the concept of each method, see About Microsoft SQL Server Backup.

Database Types	Definition	Adapter-based Protection Method		
		Volume-based Protection	File-based Protection	VDI-based Protection
Stand Alone	A SQL Server User database.	✓	✓	✓

Database Types	Definition	Adapter-based Protection Method		
		Volume-based Protection	File-based Protection	VDI-based Protection
AG	A database that belongs to an availability group (AG). For each availability database, the availability group maintains a single read-write copy (the primary replica) and one to four read-only copies (secondary replicas). For details, see Always On availability groups: a high-availability and disaster-recovery solution in the Microsoft documentation.	✅ **Note:** Volume-based Protection is not supported by AG FCC sources.	✅	✅
FCI	A database that belongs to a Failover Cluster Instance (FCI), which is a single instance of SQL Server that is installed across Windows Server Failover Clustering (WSFC) nodes.	✅	✅	✅
CDC	A change data capture (CDC) enabled SQL database records activity and is applied to an internal table. This makes the details of the changes available in an easily consumed relational format. For details, see About Change Data Capture (SQL Server) in the Microsoft documentation.	✅	✅	✅

Database Types	Definition	Adapter-based Protection Method		
		Volume-based Protection	File-based Protection	VDI-based Protection
TDE	A Transparent Data Encryption (TDE) enabled SQL database has its data files encrypted, known as encrypting data at rest. TDE performs real-time I/O encryption and decryption of the data and log files. For details, see Transparent Data Encryption (TDE) in the Microsoft documentation. **NOTE**: Cohesity does not manage the keys associated with the TDE database.	✓	✓	✓
Filestream	A Filestream enabled SQL Database is a database that contains a table that stores its data outside the database on the host file system. SQL Filestream leverages the OS file system's ability to handle large objects like images, PDFs, Binary, Bmp, Docs, and other unstructured data. For details, see Filestream (SQL Server) in the Microsoft documentation.	✓	✓	✓
SQL System	SQL Server maintains a set of four system-level databases (*master, model, msdb, tempdb*), which are essential for the operation of a server instance. System databases must be backed up after every significant update. The system databases that you must always back up include msdb, master, and model. System databases are not supported for log backups. For details, see Backup & restore: system databases (SQL Server) in the Microsoft documentation.	✓	✓	✓

Now that you can identify the type of SQL database you want to protect, you can register it as a Cohesity source.

Microsoft SQL Server Dashboard

Access the MS SQL Dashboard by going to **Dashboard** and from the drop-down, select **MS SQL**.

The page's upper portion displays information such as metrics, status and recovery history. In the page's lower portion, you can start the different actions from the Hosts tab and the Databases tab. The MS SQL Dashboard displays local MS SQL databases only. To restore or clone replicated or CloudRetrieved databases, use their respective restore or clone workflows.

From the Hosts tab, mouse over the actions menu for the following actions (some actions are available for physical servers only):

- **Protect Databases** ( )—Create a Protection Group using this server.
- **Unregister MS SQL** ( )—Unregister the MS SQL application. The server remains registered with the Cohesity cluster.
- **Edit** ( )—Configure the settings of a registered MS SQL server.
- **Unregister** ( )—Unregister this server from the Cohesity cluster.
- **Refresh** ( )—Manually force a refresh of the object hierarchy.
- **Upgrade Agents** ( )—Upgrade the Cohesity agent on the server if an upgrade is available.

From the Databases tab, mouse over the actions menu for the following actions (some actions are not available for system databases):

- **Protect** ( )—Create a Protection Group using this database.
- **Run Now** ( )—Start a run of this Protection Group. You have the option to back up one or more objects.

  If the policy associated with the Protection Group is configured with replication and archive targets, you can select those targets. If the policy supports multiple backup types, you can select an incremental, full or log backup.

  VM MS SQL Protection Groups do not support selecting individual databases for run now; only the entire VM is selectable.
- **Pause Future Runs/Resume Future Runs** ( )
  - Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the job is currently executing a run, the current run continues to execute and only future runs of this job are paused.

    - Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this job.

- **Cancel this Protection Group Run** ( )—Cancel the current run of this Protection Group. (This option is only available when a Protection Group run is executing.) This cancellation only affects the current run of the Protection Group and not future runs.
- **Clone** ( )—Make a copy of the database. The clone database is a unique copy of the original database taken from the backup.
- Follow the instructions provided in Clone Microsoft SQL Server Databases.
- **Recover** ( )—Restore a database. Follow the instructions provided in Recover Microsoft SQL Server.
- **Migrate** ( )—Migrate a database. Follow the instructions provided in Migrate Microsoft SQL Server Databases.

Database Recover and Clone Actions

In the Databases tab click a protected database to display available recover and clone actions.

To recover or clone the latest snapshot, click the arrow next to the **Protect** button at the top right corner and select an action.

By default, the latest snapshot is selected. You can optionally choose to restore from a different snapshot by clicking Edit ( ).

You can also provide a specific time to restore to if transaction logs are available.

Click a date on the calendar to select a snapshot from that date. The default point selected will be the latest log transaction or the latest full or incremental, whichever is the latest. The slider also displays available points in time and displays snapshots as circles. Manually select a point on the slider if you want to restore to a specific point in time. Selecting an invalid time from the slider automatically changes to the closest available snapshot. If your manual selection is within 10 minutes of a snapshot, your selection automatically snaps to it. You can also select a specific point in time by entering it in the time fields.

After selecting the desired recover point, click **Recover**. You can choose to clone the selected recover point by clicking the arrow adjacent to recover and selecting **Clone**.

In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

Register and Manage Microsoft SQL Server

To start protecting SQL Server databases, you need to register the SQL Server as a source in Cohesity. The following topics provide guidance that helps you register and manage SQL

Server:

Auto-discover Microsoft SQL Server Topology when Registering Microsoft SQL Server

To register an MS SQL Server, MS SQL FCI, or an MS SQL AG of a Windows cluster, you can provide the IP address of the server, FQDN of the server, or VIP of the SQL FCI. The Cohesity cluster discovers the entire MS SQL topology on the Windows cluster and displays the relevant subset based on the MS SQL address you provide while registering. From the topology list, you can select the servers you want to register with the Cohesity cluster. To display the entire MS SQL topology, use the Windows Cluster IP/FQDN.

Prerequisites

Before you register the MS SQL, ensure to meet the following:

- MS SQL prerequisites. For more information, see Requirements for Microsoft SQL Server Protection.
- Download and Install the Agent.

Download and Install the Agent

Install the Cohesity agent on each SQL server that you want Cohesity to protect.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation using the following steps:
4. **Volume CBT (Changed Block Tracker):** Install this component to leverage Cohesity's built-in CBT technology. Without this component, only full backups can be performed. Installing this component requires a one-time reboot to load the Cohesity CBT driver. Until a reboot, only full backups can be performed.

   File System CBT (Changed Block Tracker): Install this component to perform MS SQL file-based protection.

   > **Note:** Installing the Volume CBT or File System CBT is not necessary for MS SQL VDI-based protection.

5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

   - The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data

center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.

- The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.

6. Wait for the installation to finish.

   In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.

7. The agent starts automatically.
8. Reboot the machine if you have installed the Volume CBT component.
9. Repeat the agent installation process on each SQL server you want to protect. This includes any standalone MS SQL servers and Microsoft SQL Server nodes with AGs.

Register MS SQL Server

1. Select **Data Protection > Sources**.
2. Click **Register** > **Databases** > **MS SQL Server (early access: auto-detect)**.

   The Register an MS SQL Server page appears.

   > **Note:** Ensure that the agent is installed on all the nodes that are part of the Always on Availability Group (AG) and Failover Cluster Instance (FCI). For more details, Download and Install the Agent.

3. (*Applies only in multi-tenant environment with Data Source Connection enabled.*) From the **Data Source Connection** drop-down, select a Data Source Connection, and then click **Continue**.
4. In the **Hostname or IP Address** field, enter the IP address of the server, FQDN of the server, or VIP of the SQL FCI. Cohesity recommends you to provide the FQDN of the server.
5. Select the **Server Type**: Windows or Linux.

   For Linux, provide the username and password for the Database Authentication.

   > **Note:** This is an Early Access feature. Contact Cohesity Support to enable the feature.

6. Click **Register**.

   The Cohesity cluster auto-discovers the entire MS SQL topology on the cluster. Based on the MS SQL address you provide while registering, the Cohesity cluster displays

the relevant subset of that MS SQL server. The Cohesity cluster also displays which version of the Cohesity agent is running on each of the servers.

For AG and FCI, ensure to register all the nodes that are part of the AG and FCI.

7. From the topology list, select the entities you want to register with the Cohesity cluster and select **Complete Registration**.

   The selected entities are registered with the Cohesity cluster and are displayed on the source page of the Cohesity cluster UI.

Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs

This topic describes how to set up standalone MS SQL servers and MS SQL Server Always On Availability Groups (AGs) to be protected by a Cohesity cluster.

> **Note:** The **Readable Secondary** field in the Availability Group Properties of the SQL Server Availability Group (AG) must be set to "Yes" or "Read-intent only".

The following are the different methods to register your physical or virtual MS SQL Servers allowing you to protect your MS SQL server databases:

- **Registering a SQL Server using the Cohesity SQL Agent**:

  This option is the preferred method to protect your MS SQL server databases on physical or virtual hardware. This option can be used to protect any SQL server configurations such as MS MSQL standalone, FCI, and AG. For more information, see Set Up MS SQL on Virtual or Physical Servers using the Agent.

  To set up an MS SQL FCI, follow the steps described in Set Up Microsoft SQL Server FCI.

- **Registering a Hypervisor for SQL Server**:

  Cohesity provides an additional option to protect your MS SQL server at the hypervisor level. While this option is available, it is not the preferred method to protect your SQL server databases. Following are the two available registration methods:

  - **Persistent Agent**—For more information, see Set Up MS SQL on VMs as a VMware vCenter application with the Installed Agent.
  - **Auto-Deployed Agent (non-Persistent)**—For more information, see Set Up MS SQL on VMs as a VMware vCenter application with the Auto-deployed Agent.

Ensure Prerequisites are Met

Ensure the prerequisites listed in Prerequisites for Microsoft SQL Server Protection are met.

Set Up MS SQL on Virtual or Physical Servers using the Agent

This option can be used for both physical or virtual servers that are in standalone or AG configurations. Cohesity recommends you use this method to protect the MS SQL Servers

with the Cohesity cluster. For more information, see Set Up MS SQL on Virtual or Physical Servers using the Agent.

Download and Install the Agent

Install the Cohesity agent on each SQL server that you want Cohesity to protect.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation using the following steps:
4. **Volume CBT (Changed Block Tracker):** Install this component to leverage Cohesity's built-in CBT technology. Without this component, only full backups can be performed. Installing this component requires a one-time reboot to load the Cohesity CBT driver. Until a reboot, only full backups can be performed.

   File System CBT (Changed Block Tracker): Install this component to perform MS SQL file-based protection.

   > **Note:** Installing the Volume CBT or File System CBT is not necessary for MS SQL VDI-based protection.

5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

   - The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
   - The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.

6. Wait for the installation to finish.

   In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.
7. The agent starts automatically.
8. Reboot the machine if you have installed the Volume CBT component.
9. Repeat the agent installation process on each SQL server you want to protect. This includes any standalone MS SQL servers and Microsoft SQL Server nodes with AGs.

Register a Physical MS SQL Server

> **Important:** If the server contains one or more databases that belong to an Always On Availability Group (AAG), all remaining servers within the AAG must also be registered with the Cohesity cluster as MS SQL Server in order to be backed up. Navigate to **Data Protection > Sources > click on the required source > All Objects > Show AAG Nodes** to display the registered AAG nodes and any warnings about unregistered servers. If a replica is added to an existing AAG, the replica is automatically backed up by the next Protection Group run that protects that AAG, even if the AAG node list does not display it.

1. Select **Data Protection > Sources**.
2. Click **Register** > **Physical Server**.
3. (*Applies only in multi-tenant environement with Data Source Connection enabled.*) From the **Data Source Connection** drop-down, select a Data Source Connection, and then click **Continue**.
4. Specify the hostname (FQDN) or server IP of the server to register. Cohesity recommends using the FQDN.
5. Select the **Server Type**: Windows or Linux.

   For Linux, provide the username and password for the Database Authentication.

   > **Note:** This is an Early Access feature. Contact Cohesity Support to enable the feature.

6. Click **Register**.

   When registration finishes, a new entry appears in the physical section of the Sources page.

7. Next to the newly added physical server, mouse over the actions menu (⋮) and click **Register as MS SQL Server**.

   When registration finishes successfully, the MS SQL Server information is displayed when you click the 👁 icon.

> **Note:**
> You can throttle the CPU capacity on the MS SQL physical servers (running on Windows 2012R2 and later) you registered with the Cohesity cluster. Throttling optimizes the CPU capacity utilized when a Protection Group with source-side deduplication is run on the server. However, for MS SQL Protection Group, source-side deduplication is not enabled by default. You must contact Cohesity Support for enabling the source-side deduplication for a Protection Group. CPU throttling does not affect the MS SQL Server if source-side deduplication in a SQL Protection Group is not enabled. For more information on enabling CPU throttling, see Network Bandwidth and CPU Throttling for Physical Server.

AG FCC as a Source

Always On Availability Group (AG) First-Class Citizen (FCC) allows you to manage AG databases for backups using a single entity rather than multiple AG replicas.

Points to Remember Before Enabling AG FCC

- Once you switch to AG FCC, you must perform full backups for all the AG databases, since this is a new entity.
- Volume-based backup and VMware VSS-based backup are not supported by AG FCC sources.
- AG FCC is a cluster-level feature. On multi-tenant Cohesity clusters, AG FCC cannot be enabled for a single tenant; it must be enabled for all tenants.
- You can continue to protect the AG databases using the older method if there are concerns about modifying all AG SQL protection groups on the cluster.
- Log backups will not run until the initial full backup of the AG FCI is completed. Ensure to monitor the potential log growth during this period.
- AGs discovered on the registered standalone physical hosts or failover clusters are registered as sources internally.
- If AG FCC is enabled, after registration, databases belonging to an AG will appear only under the AG listed in the Sources page and not under the registered SQL host (AG replica). SQL host lists only the local databases.
- When AG sources are being protected in a protection group, other physical sources cannot be protected in the same protection group.
- Once you upgrade to AG FCC, protection runs for the existing AG databases displays the following warning:
  "*In order to enable advanced SQL Available Group functionality, please remove AG databases from Protection Groups, then refresh the physical hosts that are members*

of an AG. To disable this warning, gflagmagneto_sql_display_warning_to_enable_aag_fcc_on_prem_feature can be updated to false."

- AG FCC is enabled by default on new Cohesity clusters running version 7.3 or later. However, when an existing cluster with AG volume-based protection is upgraded to version 7.3 or later, AG FCC is not enabled automatically. Existing volume-based backups continue to run, and you can create new protection groups because AG FCC remains disabled.

To enable AG FCC:

1. Remove all AG databases from the protection groups.

   > **Note:** If you are unable to remove the AG databases due to legacy AAG blocking the migration, delete the AG/FCI protection groups. You can exclude snapshots.

2. Refresh the SQL sources that host AG databases. This identifies the sources as AG entities instead of individual SQL servers. You do not need to unregister these sources.
3. Create new backup jobs and start the backup.

Set Up MS SQL on VMs as a VMware vCenter application with the Installed Agent

This method must be used only when you are registering SQL Server at the hypervisor-level. You can use this option for standalone SQL or SQL servers in AG running on the VMware vCenter.

Download and Install the Agent

Before registering the MSSQL server at the hypervisor-level, you must first install the Cohesity agent on each SQL server that you want Cohesity to protect.

You can optionally use the auto-deployed Cohesity agent, which does not require agent installation. Cohesity, however, recommends installing the agent because this avoids additional configuration such as disabling Windows User Account Control (UAC). If you still want to use the auto-deployed agent, follow the steps described in Set Up MS SQL on VMs as a VMware vCenter application with the Auto-deployed Agent.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation wizard.
4. You do not need to install any components. Backups leverage VMware VADP CBT or Hyper-V 2016 RCT.

5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

   - The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
   - The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.

6. Wait for the installation to finish.

   In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.

7. The agent starts automatically.
8. Repeat the agent installation process on each SQL server you want to protect. This includes any standalone MS SQL servers and MS SQL Server nodes with AGs.

Register MS SQL

> **Important:** If the server contains one or more databases that belong to an Always On Availability Group (AAG), all remaining servers within the AAG must also be registered with the Cohesity cluster as MS SQL Server in order to be backed up. Navigate to **Data Protection > Sources > click on the required source > All Objects > Show AAG Nodes** to display the registered AAG nodes and any warnings about unregistered servers. If a replica is added to an existing AAG, the replica is automatically backed up by the next Protection Group run that protects that AAG, even if the AAG node list does not display it.

1. Select **Data Protection > Sources**.
2. Navigate the hierarchy to the server that is running MS SQL.
3. Mouse over the server's actions menu (⋮) and click **Register Applications**.
4. Select **Installed Agent**.

   If the agent is not already installed, registration fails.

5. Select MS SQL.
6. Click **Register**.

   The agent starts automatically.

   The server's MS SQL and agent information display when you click the icon.

Set Up MS SQL on VMs as a VMware vCenter application with the Auto-deployed Agent

> **Important:** If the server contains one or more databases that belong to an Always On Availability Group (AAG), all remaining servers within the AAG must also be registered with the Cohesity cluster as MS SQL Server in order to be backed up. Navigate to **Data Protection > Sources > click on the required source > All Objects > Show AAG Nodes** to display the registered AAG nodes and any warnings about unregistered servers. If a replica is added to an existing AAG, the replica is automatically backed up by the next Protection Group run that protects that AAG, even if the AAG node list does not display it.

1. Select **Data Protection > Sources**.
2. Navigate the hierarchy to the server that is running MS SQL.
3. Mouse over the server's actions menu (⋮) and click **Register Applications**.
4. Select **Auto-Deployed Agent**.

   Enter the Windows administrator **Username** and **Password** for the server that is running MS SQL. Windows administrator privileges are required. Ensure that the user account used for auto-deploying the agent is added to all the MS SQL server's SQL instances' **Security > Logins** with Windows Integrated Authentication.

   Disable the UAC property Admin Approval Mode

   1. Click **Start**, type secpol.msc in the Search programs and files box, and then press Enter.
   2. If the User Account Control dialog box appears, confirm that the action it displays is what you want, and then click **Yes**.
   3. In the console tree, expand Local Policies, and click **Security Options**.
   4. In the details pane, scroll down and double-click the **Group Policy** setting that you want to change. This opens the UAC policy settings that a local administrator can modify.
   5. Disable **User Account Control: Run all administrators in Admin Approval Mode** for the Administrator account.
   6. Ensure that **User Account Control: Admin Approval Mode for Built-in Administrator Account** is set to Disabled.
   7. Ensure that **User Account Control: Behavior of elevation prompt for administrators in Admin Approval Mode** is set to Elevate without Prompting.
   8. Restart/reboot the Windows server for the changes to take effect.
   9. Provide the Administrator/Domain Administrator user that has Admin Approval Mode disabled.

5. Select MS SQL.
6. Click **Register**.

   The server's MS SQL information displays when you click the icon.

Upgrade the Cohesity Agent and Add the File System CBT Component

If necessary, you can upgrade the agent and add the File System Changed Block Tracking (CBT) manually. Start the upgrade and select the option to include the File System CBT component in the upgrade.

Troubleshoot

- The log backup of AG MS SQL Server fails with the following error if there is a break in the log chain.

  ```
  Log chain break error: Discovered a break in the log chain for
  <Database Names>
  ```

  Resolution: To resolve this issue, ensure that no other third-party applications are running a log backup, and then perform a full backup run to reset the log chain.
- The log backup of the AG MS SQL Server fails with the following error if a database is added or removed from the AG MS SQL Server.

  ```
  AG relationship error: Discovered a AG relationship error for database
  <Database Names>
  ```

  Resolution: To resolve this error, perform a full backup run.

For more troubleshooting information, see the following KB article:

- Collecting troubleshooting information for MS SQL issues

Log in to the Cohesity Support Portal to see more KB articles.

Set Up Microsoft SQL Server FCI

This topic describes how to set up an MS SQL Failover Cluster Instance (FCI) to be protected by a Cohesity cluster. To set up a standalone MS SQL Server or MS SQL AG, follow the steps described in Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs.

MS SQL FCI - Overview

An MS SQL FCI is an MS SQL instance installed on a Windows server that is part of a Windows Server Failover Cluster (WSFC or simply Windows Cluster for short). The SQL instance physically exists on (and is accessible from) only one node of the WSFC at any given time. If the FCI's current node goes down, the FCI moves to another node in the WSFC. The FCI is addressed by a Virtual IP (VIP) which points to the physical IP of the node on which it currently resides.

Set Up MS SQL FCI Overview

MS SQL FCI setup has the following steps:

1. Ensure Prerequisites Are Met
2. Download and Install the Agent on Each Node
3. Register an MS SQL FCI as a Source
4. Register an MS SQL FCI as MS SQL

Ensure Prerequisites Are Met

Ensure the prerequisites listed in Prerequisites for Microsoft SQL Server Protection are met.

Download and Install the Agent on Each Node

Before registering an MS SQL Server FCI, you must download and install the Cohesity agent on each node within the MS SQL FCI.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation using the following steps:
4. **Volume CBT (Changed Block Tracker):** Install this component to leverage Cohesity's built-in CBT technology. Without this component, only full backups can be performed. Installing this component requires a one-time reboot to load the Cohesity CBT driver. Until a reboot, only full backups can be performed.

   File System CBT (Changed Block Tracker): Install this component to perform MS SQL file-based protection.

   > **Note:** Installing the Volume CBT or File System CBT is not necessary for MS SQL VDI-based protection.

5. **Service Account Credentials:** You can use either the LOCAL SYSTEM account or an account that meets the requirements to install the Cohesity agent.

   If you do not use the LOCAL SYSTEM account, ensure the following:

   - The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the SQL server.
   - The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the Cohesity agent.
6. Wait for the installation to finish.

In the SQL Server Management Studio (SSMS), ensure that the account used to install the Cohesity agent has **Server Roles: sysadmin** in the SQL server instances.

7. The agent starts automatically.
8. Reboot the machine if you have installed the Volume CBT component.
9. Repeat the agent installation process on each node within the MS SQL FCI.

   After you install the agent on each node within the FCI, the MS SQL FCI can be registered as a source with the Cohesity cluster.

Register an MS SQL FCI as a Source

This registration step makes an MS SQL FCI visible to the Cohesity cluster so that its content is eligible to be protected or backed up.

> **Note:**
> - An MS SQL FCI can be registered with only one Cohesity cluster at a time.
> - Each MS SQL FCI must be registered separately.

1. Select **Data Protection > Sources**.
2. Click **Register** > **Databases** > **MS SQL Cluster**.
3. Enter the **Hostname**.

   Cohesity recommends that you specify the fully qualified domain name of the MS SQL FCI to register. Or you can specify the virtual IP address (VIP) of the MS SQL FCI present under SQL Role.

Fully qualified domain name location:

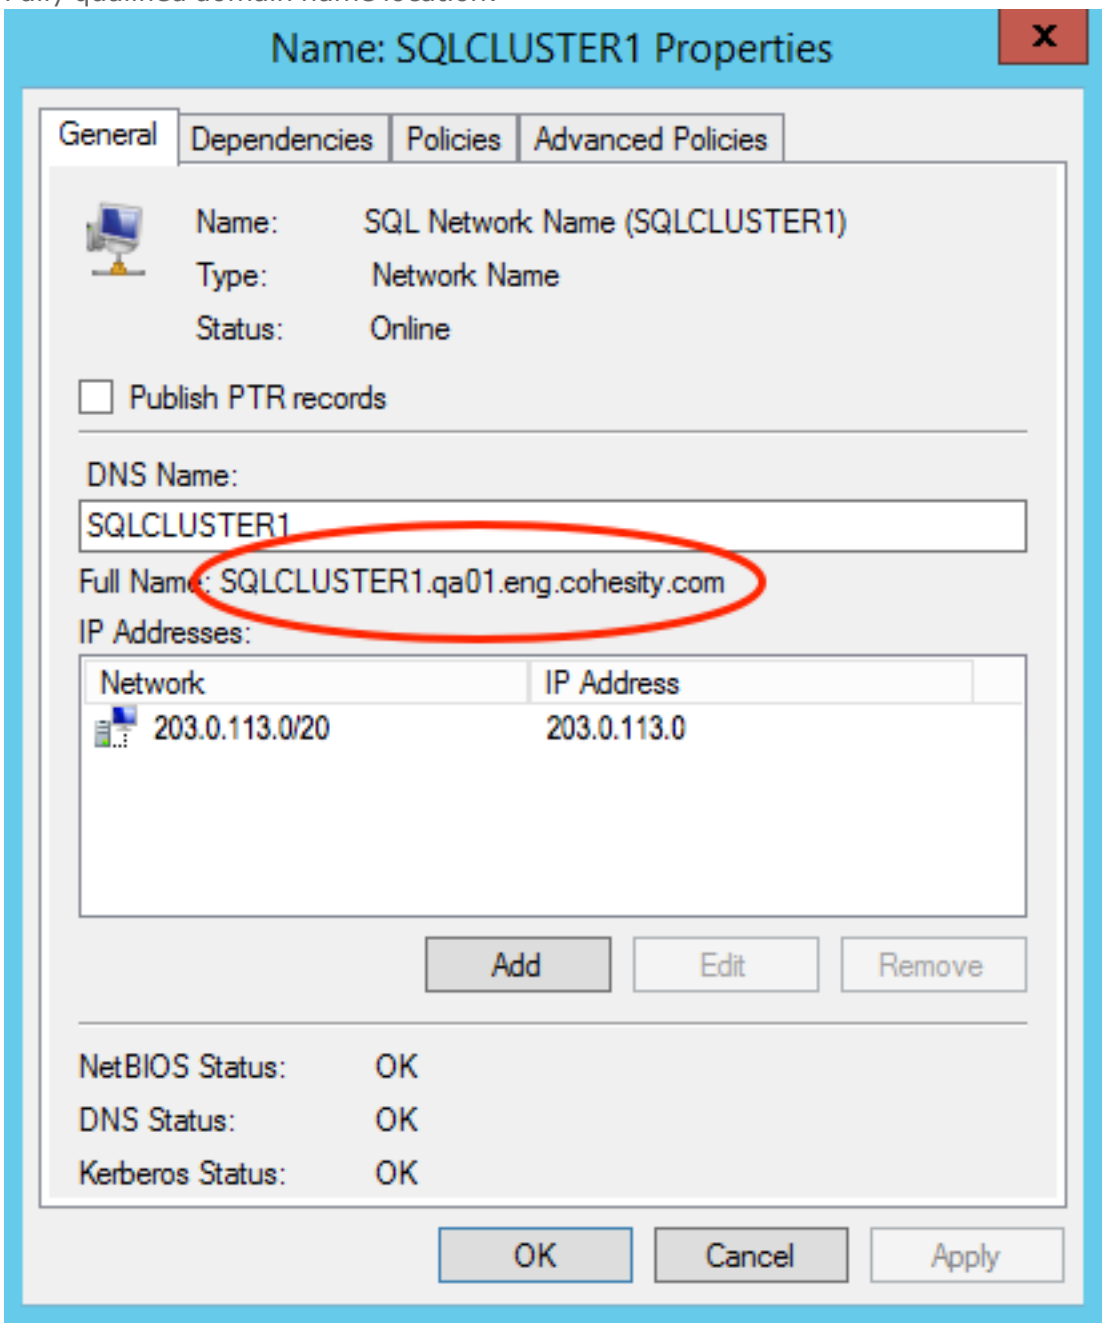

Virtual IP address location:

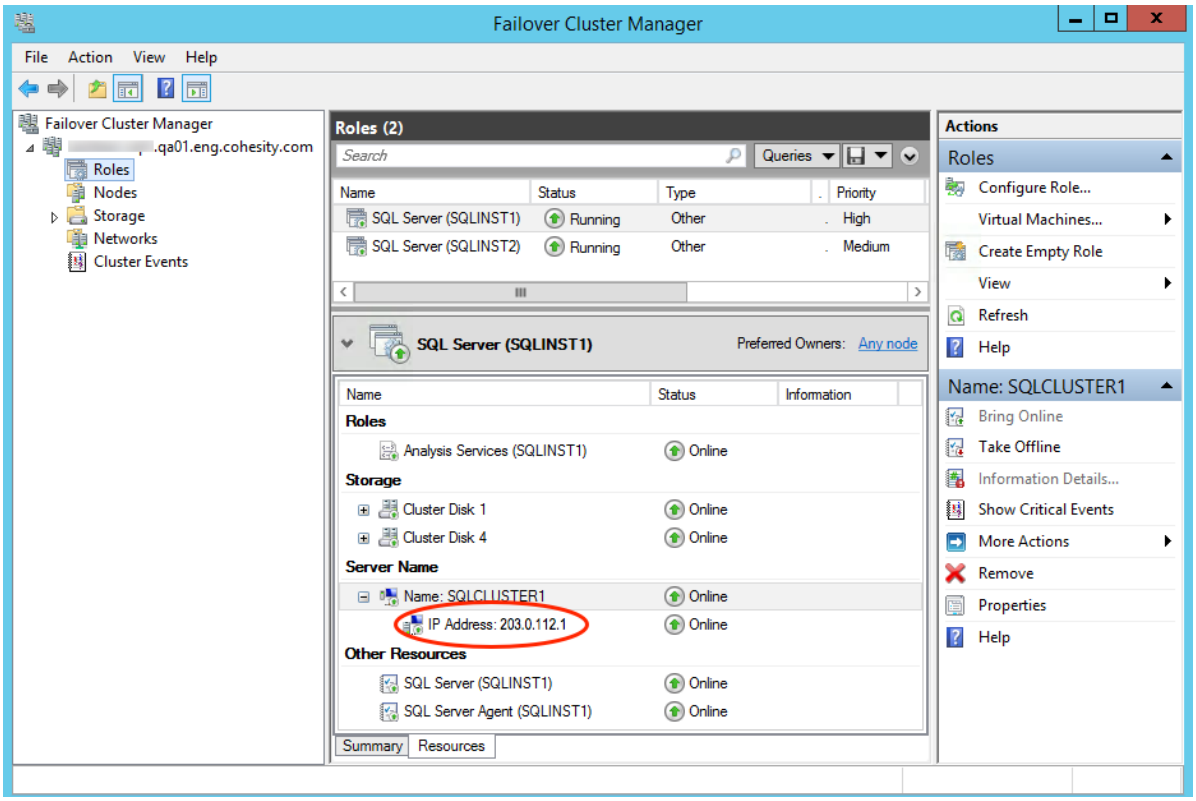

> **Note:** If the MS SQL FCI is not registered using an FQDN and all its nodes are not configured with FQDNs (for example, it is a test environment), changing the IP address of a registered node or adding/removing nodes may not update the FCI's registration configuration information, which could prevent some Cohesity SQL workflows. If you change the configuration information or add/remove nodes under these conditions, Cohesity recommends that you unregister and then reregister the MS SQL FCI. MS SQL FCIs that are registered using an FQDN do not have limitations when updating configuration information such as node IP address changes or adding/removing nodes.

4. Click **Register**.

   When registration finishes, a new entry appears in the physical section on the Sources page.

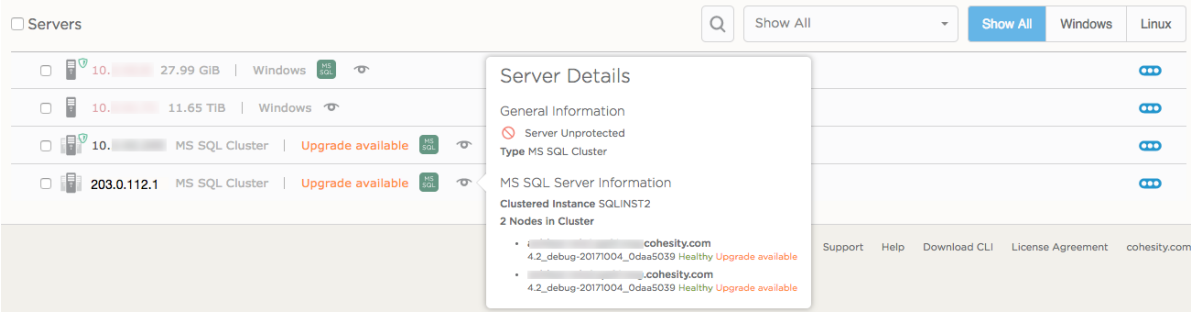

The MS SQL FCI is now registered as a physical server.

Register an MS SQL FCI as MS SQL

This registration step takes an MS SQL FCI source that is already visible (as a physical server) to the Cohesity cluster and registers it as MS SQL Server.

Next to the newly added MS SQL FCI (displayed as a physical server), mouse over the MS SQL FCI's actions menu and click **Register as MS SQL Server**.

When registration finishes successfully, MS SQL Server information displays when you click the icon.

FCI Standalone Multi-Registration

You can register and back up a Windows Server that participates in an MS SQL FCI. You can separately register the FCI clusters and their standalone nodes. This allows you to backup local drives for physical server backups.

Prerequisites

- Ensure that all the SQL agents are upgraded.
- Ensure that the following unsupported protection groups are not present:
  - A volume-based protection group that protects SQL servers participating in a WSFC and is either configured to back up all volumes.
  - A protection group that auto-protects a standalone SQL host participating in a WSFC.
  - A protection group that explicitly protects FCI databases on a SQL server participating in an MS SQL FCI.

Error Messages

- The following error message is displayed when unsupported protection groups are present:
  Error
  FCI-node multi-registration cannot proceed, please correct the following protection groups:

- <protectiongroupname1>, <protectiongroupname2> - volume-based jobs that back up all volumes of Windows Cluster node(s).
- <protectiongroupname1>, <protectiongroupname2> - jobs that autoprotect Windows Cluster node(s).
- <protectiongroupname1>, <protectiongroupname2> - jobs that explicitly back up some FCI databases from Windows Cluster node(s).

Solution

Either delete or fix the protection groups such that they no longer fit the unsupported criteria, and continue the procedure.

- The following error message is displayed when the agent on the Windows cluster node has not been upgraded to a version that supports this feature:
  Error
  Agents on the following sources do not support FCI-node multi-registration: <server1>, <server2>
  Solution
  Upgrade the agent on the Windows cluster node to a version that supports this feature.

Best Practices for Microsoft SQL Server

Cohesity recommends the following best practices when protecting MS SQL:

- Cohesity recommends triggering a full backup whenever a database file (.mdf, .ndf, .ldf) is added or removed from the database.
- Multiple Cohesity clusters concurrently backing up an MS SQL VM is not recommended because the jobs will interfere with the other's log chain.
- For MS SQL point in time restores, the Full Recovery Model is recommended as a best practice. For more information, see Choosing the Recovery Model for a Database.
- Set MS SQL Server to autostart.
- To accurately report the size of transaction log data read/written during log backups, the Windows local or domain Administrator account that is registered with Cohesity for backup must be added as a user to the 'model' system database. Use the following query:

```
-- Configures accurate reporting of the size of transaction log data
read/written during log backup
USE [model]
GO
CREATE USER [HOSTNAME\Administrator] FOR LOGIN [HOSTNAME\Administrator]

GO
```

  Where HOSTNAME is the hostname of the SQL server.
  This assumes that the Administrator account used to register MS SQL is already

registered in the SQL server instance's Security Logins and has the "sysadmin" role per documentation. If privileges are not correct, the logical bytes read stat ("Size" column on the Job's details page) for MS SQL log backups will be the same as the physical bytes read stat ("Read" column on the Job's details page). If privileges are correct, then the size reported will be the sum of the sizes of all the .ldf files for the database.

Considerations for Microsoft SQL Server

General

- Windows Physical and Microsoft SQL Server refresh for a host can fail with the following error when the agent process reaches a high memory watermark.

  "kAgentError: Excessive memory usage (spike) on the Agent."

  This watermark is based on the peak memory usage of the agent. As a result, the refresh can continue to fail even after the current memory usage of the agent decreases.

  As a workaround, restart the agent.
- For file-based Protection Group, object-level granularity is not supported for Role-based access control (RBAC), such as setting access to objects through Active Directory and roles.
- If using native or third-party backup software to back up an application (such as MS SQL), note that the following behavior breaks the differential chain for the application on that server:
  - Performing an app-consistent backup on a VMware SQL VM.
  - Backing up a physical server using a (non-SQL) volume-based physical server Protection Group. You are allowed to run a file-based physical server Protection Group with a SQL Protection Group or 3rd party SQL protection tools.
    This is a Microsoft limitation. For more information, see Various Causes of Log Chains Breaking Issue in SQL Server.
- The SQL Server VDI log backup and restore operations require Sysadmin privileges which is a Microsoft requirement. For more details, see Microsoft Documentation on SQL Server VDI backup.
- On-prem SharePoint databases can be backed up and recovered through Microsoft SQL Server. To recover granular items, you must use third-party software, Ontrack Power Controls, to mine the database.

Protection

- Ensure that all AG databases you want to protect are selected from the respective replica.

- For file-based backups, Cohesity operations rely on default MS SQL file extensions (.mdf, .ndf, .ldf).
- SQL Database and logs residing on SMB/NFS can be protected using the MS SQL Native method.
- SQL Database and logs residing on Clustered Shared Volumes (CSV) can be protected using the MS SQL Native and MS SQL VDI methods.
- CBT information does not persist after a reboot for the file-based backup method.
- As per Microsoft requirements, VDI backups executed on secondary replicas are copy-only full backups. Copy-only backups on secondary replicas do not impact the log chain. For more information, see Offload supported backups to secondary replicas of an availability group.
- Overwriting Filestream database is not supported; this may prevent the completion of recovery.
- Generally, volume-based backups over Resilient File System (ReFS) are not supported. But if the volume is presented through the Windows Cluster Shared Volume (CSV) layer, the ReFS file system detection fails causing the volume-based backups to succeed.
- You can run log backups in parallel with existing log backup runs. This helps reduce storage buildup on the primary SQL Server when multiple large backups are in progress.

  This feature is supported only for file-based, VDI-based, and Filestream (backed up using VDI-based method) backups.

  > **Note:** When you run a log backup for a database backup job, any preceding backup job for the same database may skip the log backup if the new backup job is completed before the previous backup job for the database is started.

  The resource manager reserves additional slots specifically for Run Now based log backups to start immediately. The databases in queue are skipped, and the log backup is initiated.

Recovery

- Performing a point-in-time restore of an Always On Availability Group replica requires full backups from that replica.

  Log backups are merged across all replicas; full backups are not. During restore, the user sees multiple results for the same database from different replicas. Each replica shows a set of Snapshots for full backups done on that particular replica. For point in

time restore to a specific replica, the user must select a full backup Snapshot from that replica only. If a full backup does not exist, point in time restore is not possible.

Restore the database to the desired MS SQL Server instance on a replica that has a full backup.

- Using the "Restore to Original SQL Server instance" option is not supported if any of the following occurs:
  - The Windows host name changed and you attempt to restore a snapshot that was taken before the host name change.
  - Granular DB restore performed from a higher SQL Server version to a lower version. This is a Microsoft limitation.

  You can, toggle off "**Restore to Original SQL Server Instance**" and perform a database restore.

  For more information, see Rename computer hosting instance - SQL Server.

  Restoring a snapshot taken after the name change or upgrade is supported.
- Restoring MS SQL Server from backups of dynamic disks to the same machine is not supported.
- Use File-based or VDI-based backup methods to protect TDE databases. Follow the Microsoft TDE best practices for key and cert management.
- If multiple restore tasks are run on the same target host, the restore tasks will be serialized.
- **For system database recovery**:
  - Recovering the system database to an alternate instance is not supported.
  - When recovering to the original database, the original database is deleted before the recovery. Therefore, a recovery failure will also lead to the loss of the original database.
  - If you are recovering the system databases individually, then the recovery of the system database must be in the following order:
    1. Master
    2. MSDB
    3. Model
  - You cannot recover a system database along with a user database.
  - If you are recovering more than one system databases in a recovery task, then the recovery point for the selected system databases must be the same.
  - For recovering a database to an Amazon RDS Custom SQL Server, the recovery destination path must be under the **D:\rdsdbdata** folder.
- Download and File-Level recovery of SQL files from FortKnox is not supported.

Clone

- Cloning an MS SQL database and renaming the database via SSMS can leave the cloned database in the "Recovery Pending" state during a clone teardown task. This is because the Cohesity cluster database clone operation assumes the cloned database name and files are associated with the original clone operation. To prevent a disconnect between the Cohesity cluster and SQL Server operation, it is recommended to clone a new database with the intended database name rather than rename a cloned database.

Migration

- Cohesity does not support the migration of MS SQL FILESTREAM databases.
- Once a migration task is configured and created, the target SQL Server or database paths cannot be modified for the existing migration task.
- Cohesity does not detect changes/deletions to the target database with a running migration task. Changes made external to the migration task before finalizing can cause the migration task to fail.
- Cohesity does not support the migration of MS SQL databases that are replicated to an alternate Cohesity cluster.

General Recommendations for Microsoft SQL Server Backup Methods

Filestream databases can be protected using the following approaches, listed in the order of recommendation:

- VDI-based
- File-based
- Volume-based

Normal databases can be protected using the following approaches, listed in the order of recommendation:

- File-based
- VDI-based
- Volume-based

The following table can help you select the appropriate backup type:

	File-based Backup	VDI-based Backup	Volume-based Backup
Microsoft Approach	Leverages Microsoft VSS approach	Leverages Microsoft SQL Server native backup and restore commands through API calls	Leverages Microsoft VSS approach

	File-based Backup	VDI-based Backup	Volume-based Backup
Built-in Change Block Tracking (CBT)	Leverages Cohesity's built-in CBT technology	Leverages SQL server's differential backups	Leverages Cohesity's built-in CBT technology
Database clone from backup	Supported	Not Supported	Supported
Recovery: Cross-domain AD without trust	Supported (GRPC is used to transfer data)	Supported (GRPC is used to transfer data)	Not Supported (Recovery uses SMB and cross-domain trust is required for SMB access)
Database clone: Cross-domain AD without trust	Not Supported (Clone uses SMB and cross-domain trust is required for SMB access)	-	Not Supported (Clone uses SMB and cross-domain trust is required for SMB access)
Implicit Full backup after SQL host restart	The first backup after the SQL server reboot will be a full backup (incremental backup automatically switches to full backup)	-	-
Memory Optimized Database	Not Supported	Supported	Not Supported
SQL server-side compression	Not Supported	Supported but not recommended to use, since it may cause low deduplication rate on the Cohesity clusters, which leads to high storage utilization on the Cohesity cluster.	Not Supported
Parallel Log Backup to Full/ Incremental Backup	Supported	Supported	Not Supported
Fast Backup (multi-node)	Supported	Supported	Supported

	File-based Backup	VDI-based Backup	Volume-based Backup
Fast Restore (multi-node)	Supported	Supported	Not Supported. The restore throughput is limited by the single cluster node throughput.
Database-level recovery from archived backup	Supported (Beta till 7.0)	Supported (Beta till 7.0)	Not supported via workflow; can be performed manually.
CAD Policy	Supported	Supported	Not Supported
Transparent Data Encryption (TDE) Databases	Supported	Supported	Not Supported
Database on ReFS volumes	Supported	Supported	Not Supported
Quiesce Support	SQL Database goes through quiesce (freeze/thaw)	Not required. The SQL server performs consistent backup	SQL database goes through quiesce (freeze/thaw)
Free Space on SQL Host (VSS shadow storage requirement)	At least 10% is recommended for each volume on the SQL host by MSFT. Large databases with high IO rates may require additional space.	-	At least 10% is recommended for each volume on the SQL host by MSFT. Large databases with high IO rates may require additional space.
Must-Have	-	Must have Full and incremental (differential). NOTE: Setting extended retention is required for full backup to avoid expiry before the incremental backup expiry.	-

	File-based Backup	VDI-based Backup	Volume-based Backup
Parallel run-now backups (same type) for the same job	When the host is already part of a previous run-now submitted, the backup is queued until the previous submitted run is complete.	Backup of the selected database runs in parallel with the existing run-now request.	Not Supported
Known Issues that lead to backup failure	Very High transaction rate on databases may occur resulting in backup failure. This is due to Microsoft's VSS system provider not being able to maintain the VSS snapshot on SQL host with very high transactions. This is very rare. Configuring the shadow storage on volume with no/less IO helps to mitigate the issue.	Agent/SQL host may restart during backup.	Very High transaction rate on databases may occur resulting in backup failure. This is due to Microsoft's VSS system provider not being able to maintain the VSS snapshot on SQL host with very high transactions. This is very rare. Configuring the shadow storage on volume with no/less IO helps to mitigate the issue.
Known Issues that lead to incremental backup converting to full backup	First backup after SQL server reboot will be a full backup (incremental backup automatically switches to full backup)	If no explicit policy is available for full backup, incremental backup is converted to full backup after 7 days.	-
Known issues that lead to PIT clone failure	-	-	PIT clone may fail if enough free space is not available (at least equal to the database being cloned) on the primary volume during backup.
FortKnox	Supported	Supported	Not supported

> **Note:** REFS is supported by VDI-based and File-based SQL server.

The following table can help you select the appropriate backup type for Filestream databases:

Workload	Operation	File-based Backup	VDI-based Backup	Volume-based Backup	Recommended Order
Large Number of Small Files - high file count, small individual file sizes	Backup	Diff streamer cost per file dominates.	VDI periodic full (weekly). No per-file diff streamer overhead.	Block-level incremental. No per-file cost.	Volume-based → VDI → File-based
	Recovery	Per-file overhead scales poorly. Slow at very high file counts.	SQL Server handles file enumeration natively. No diff streamer cost.	Per-file copy overhead from single copy node compounds at scale.	VDI → File-based → Volume-based
	Clone / Instant Access	Clone and flat file recovery not supported.	Clone not supported.	Only approach supporting clones / instant access.	Volume-based only
	Archival and Cloud Recovery	Slower archival throughput than VDI.	Faster archival throughput - fewer .bak files to archive compared to file-based.	Cloud recovery not supported.	VDI → File-based

Workload	Operation	File-based Backup	VDI-based Backup	Volume-based Backup	Recommended Order
Significant File Size - fewer files, large individual sizes (data transfer dominated)	Backup	Diff streamer overhead is minimal when file count is low. Data copy dominates and file-based is efficient.	VDI periodic full has higher SQL Server overhead when data transfer is the bottleneck.	Block-level incremental wins regardless of file size.	Volume-based → File-based → VDI
	Recovery	Data transfer dominates - copies files directly with minimal overhead.	Overhead adds cost when data transfer is the bottleneck.	Single copy node, per-file overhead still makes it slowest.	File-based → VDI → Volume-based
	Clone / Instant Access	Clone and flat file recovery not supported.	Clone not supported.	Only approach supporting clones / instant access.	Volume-based only
	Archival and Cloud Recovery	Slower archival throughput than VDI.	Faster archival throughput - fewer .bak files to archive compared to file-based.	Cloud recovery not supported.	VDI → File-based

Backup Microsoft SQL Server

The following topics provide guidance to help you protect MS SQL databases:

For more information on MS SQL deployments, See Protect SQL Server with Cohesity—Deployment Considerations Guide.

About Microsoft SQL Server Backup

Cohesity supports the following MS SQL Server backups:

- Volume-based Backup
- File-based Backup
- VDI-based Backup

Cohesity also supports VM-based Backup and SQL Native Backup.

Volume-based Backup

A Cohesity MS SQL volume-based backup captures all MS SQL databases running on a VMware or Physical Server, whose files reside on one or more volumes, as well as all other files that are stored on the corresponding OS volume. This approach contrasts with a file-based MS SQL backup, which includes only the selected databases. The Cohesity cluster captures full database server backups and can optionally back up database transaction logs, so you can roll forward to any point in time.

For information on performing MS SQL volume-based backup, see Backup Microsoft SQL Server (Volume-based).

All configurations are for illustrative purposes only. Cohesity recommends following MS SQL best practices.

An example of MS SQL running on a Windows server that contains the following:

- Multiple volumes, C, D and F in the figure below
- Volume D is running multiple MS SQL databases
- Volume F contains database logs

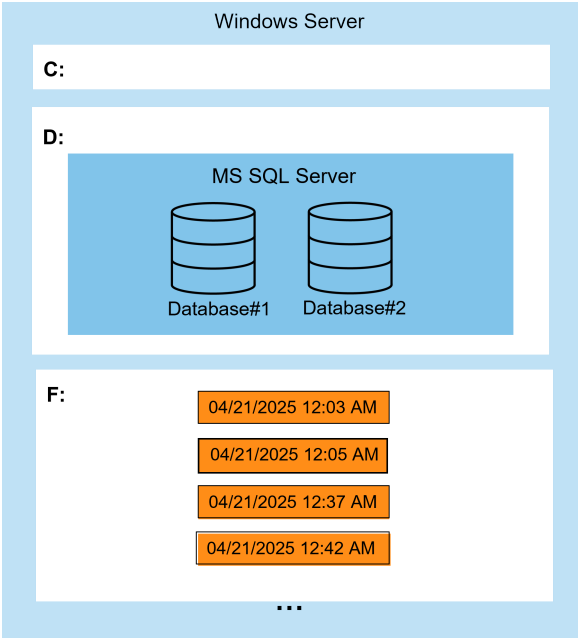

The MS SQL volume-based backup task consists of the following actions:

1. During an MS SQL protection run, a VSS snapshot of MS SQL is captured using Volume Shadow Copy Service (VSS). The VSS snapshot resides on the local drive as a special VSS volume. This capturing of the VSS snapshot is a full database server backup. This is the same for MS SQL running on VMware and physical servers.

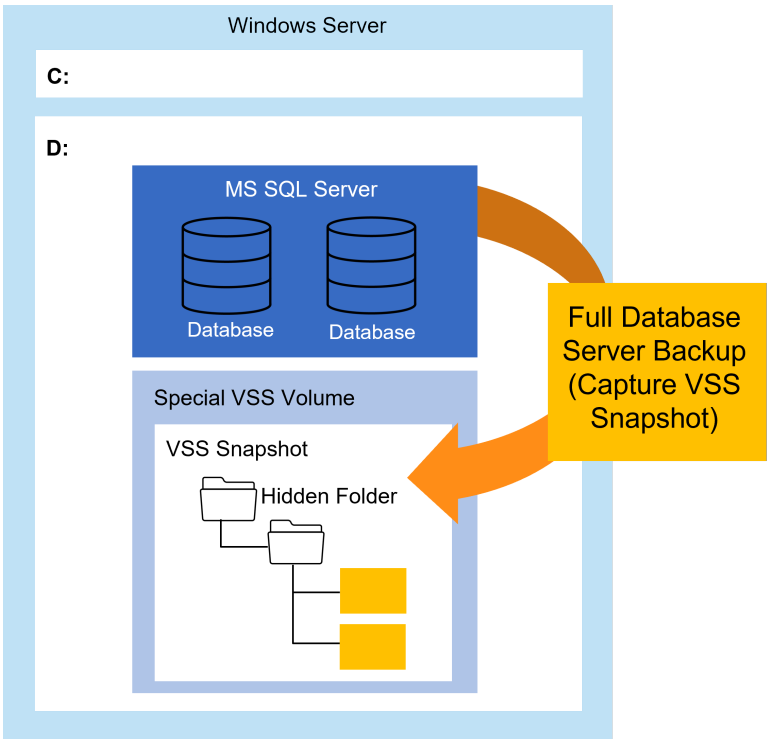

2. The snapshot is captured and saved on the Cohesity cluster with a date and time stamp. You can optionally choose for the Cohesity cluster to utilize Change Block Tracking (CBT) when backing up the snapshot. VMware VMs leverage VMware CBT to back up the snapshot. Physical servers leverage Cohesity CBT to back up the snapshot.

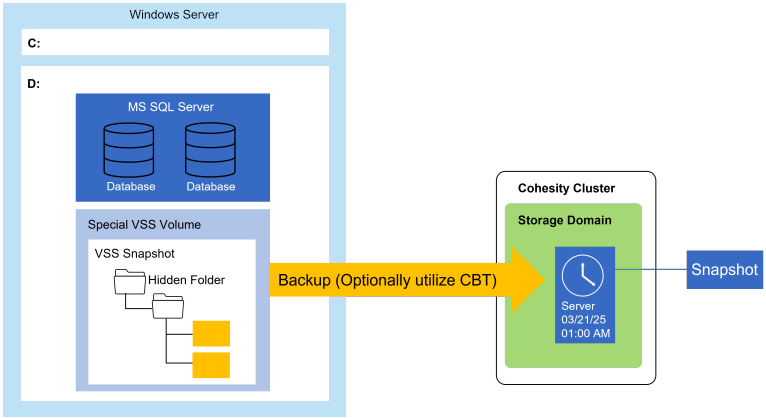

3. If specified in the policy, the database transaction logs are backed up to the Cohesity cluster using MS SQL Server Virtual Device Interface (VDI). The logs allow you to roll back to a specific point in time. MS SQL Server automatically truncates the SQL logs after a successful SQL log backup and marks the Virtual Log Files (VLFs) for reuse.

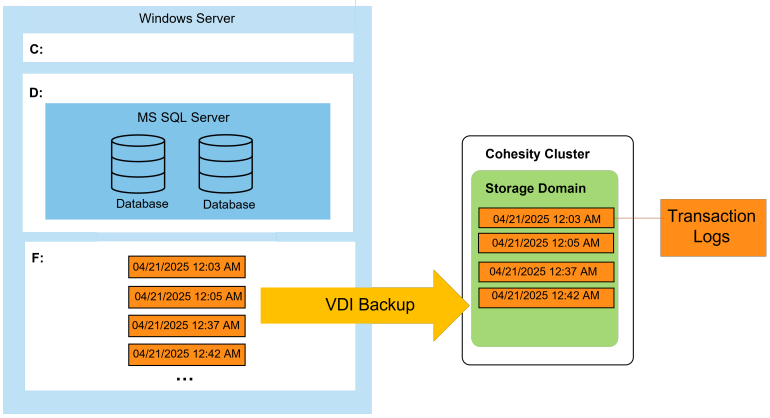

File-based Backup

A Cohesity cluster MS SQL file-based backup protects only the MS SQL databases you choose. It captures only the database files for those selected databases. This approach contrasts with a volume-based backup, which captures any and all the files contained on the volume. The Cohesity cluster captures full database server backups and can optionally back up database transaction logs, so you can roll forward to any point in time.

The advantage of a file-based backup is that a smaller amount of data is captured, which allows the Protection Group to run more quickly.

> **Note:** To enable file-based backup, you must install the File System CBT component during Cohesity agent installation.

For information on performing MS SQL file-based backup, see Backup Microsoft SQL Server (File-based).

All configurations are for illustrative purposes only. Cohesity recommends following MS SQL best practices.

An example of MS SQL running on a Windows server that contains the following:

- Multiple volumes, C, D, E and F in the figure below
- Volume D is running two MS SQL databases
- Volume E is running three MS SQL databases
- Volume F contains database logs

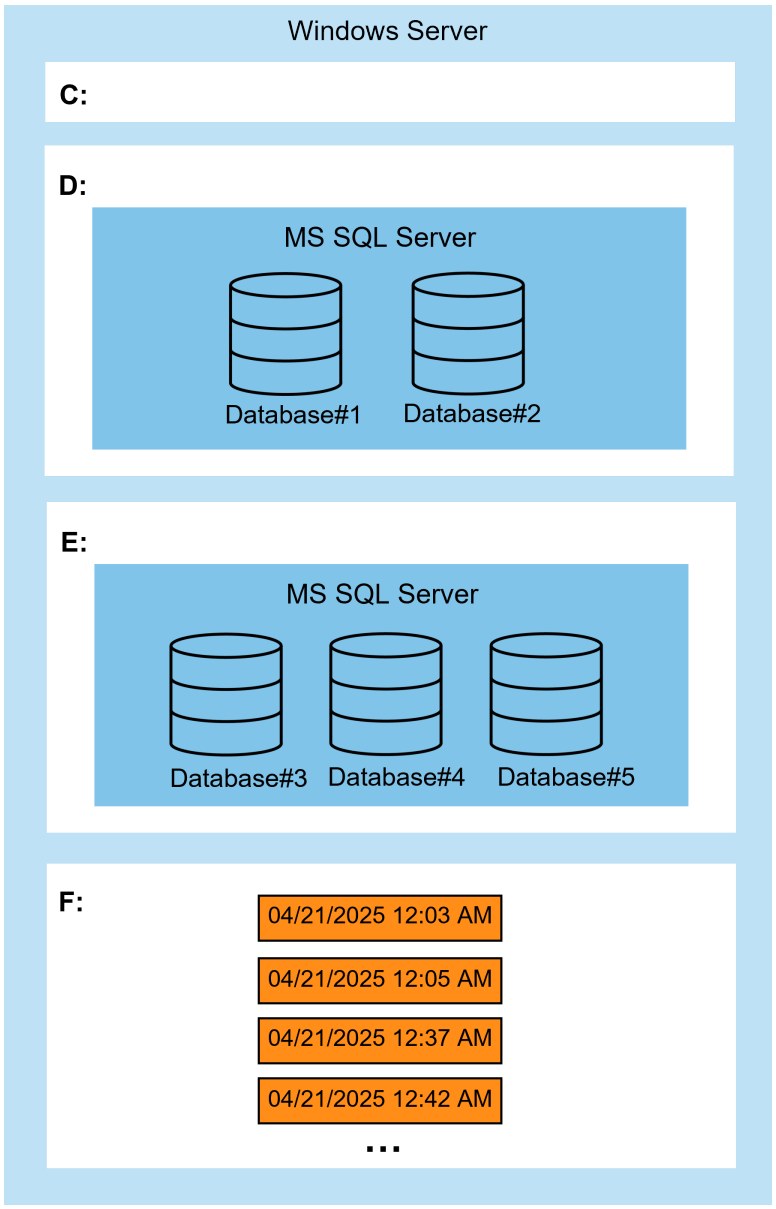

The MS SQL file-based backup task consists of the following actions:

1. During SQL database protection run, a VSS snapshot of the specified MS SQL databases is captured using Volume Shadow Copy Service (VSS). The VSS snapshot resides on the local drive as a special VSS volume. This VSS snapshot is a full database server backup.

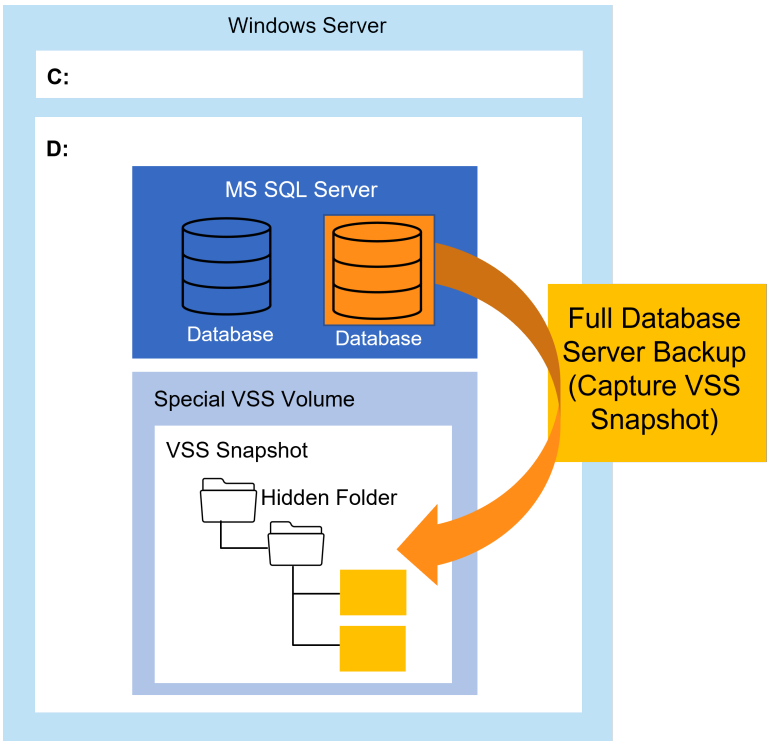

2. The snapshot is captured and saved on the Cohesity cluster with a date and time stamp. You can optionally choose for the Cohesity cluster to utilize Change Block Tracking (CBT) when backing up the snapshot. VMware VMs leverage VMware CBT to back up the snapshot. Physical servers leverage Cohesity CBT to back up the snapshot.

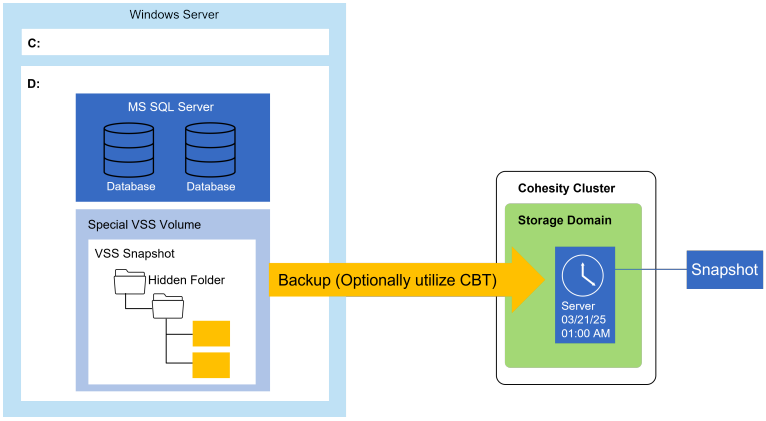

3. Backups of database transaction logs use MS SQL Server Virtual Device Interface (VDI) and are stored on the Cohesity cluster. The logs allow you to roll back to a specific point in time. MS SQL Server automatically truncates the SQL logs after a successful SQL log backup and marks the Virtual Log Files (VLFs) for reuse.

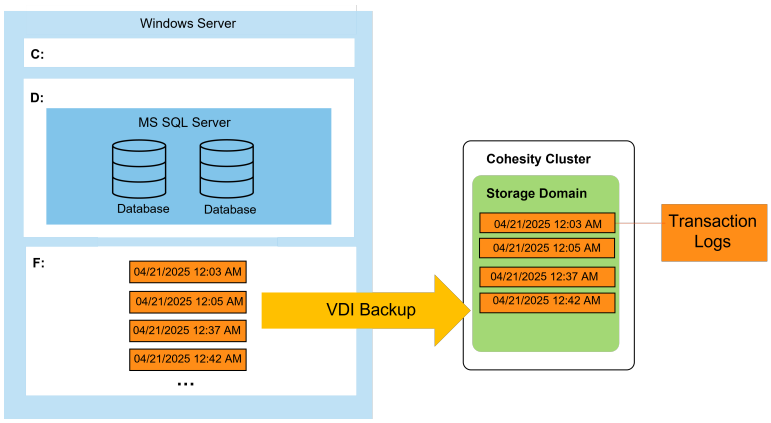

VDI-based Backup

A Cohesity cluster MS SQL VDI-based backup captures the native MS SQL server backups. It captures VDI-based backups of a specific database on a SQL Server instance running on a VMware or physical server.

The Virtual Device Interface (VDI) is a Microsoft interface that allows the Cohesity agent to execute SQL Server Native backup and restore commands. This backup type produces a standard SQL Server Native backup of type (.BAK), transaction logs, (.TRN), and differentials, (.DIFF). These backup files are stored on a Cohesity View and take advantage of View compression, deduplication, encryption, replication, and archiving features.

For information on performing MS SQL VDI-based backup, see, see Backup Microsoft SQL Server (VDI-based).

All configurations are for illustrative purposes only. Cohesity recommends following MS SQL best practices.

An example of MS SQL running on a Windows server that contains the following:

- Multiple volumes, C, D and F in the figure below
- Volume D is running multiple MS SQL databases
- Volume F contains database logs

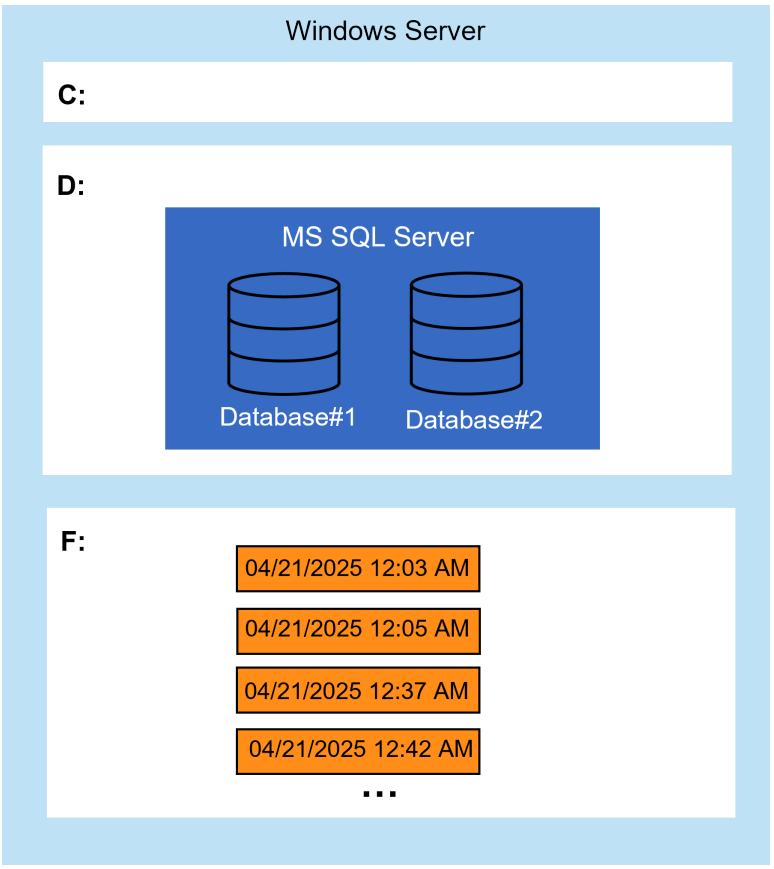

The MS SQL VDI-based backup task consists of the following actions:

1. During an MS SQL protection run, a "BACKUP DATABASE" command for the specified MS SQL databases is initiated using MS SQL Server Virtual Device Interface (VDI). The generated backup file is saved on the Cohesity cluster with a date and time stamp.

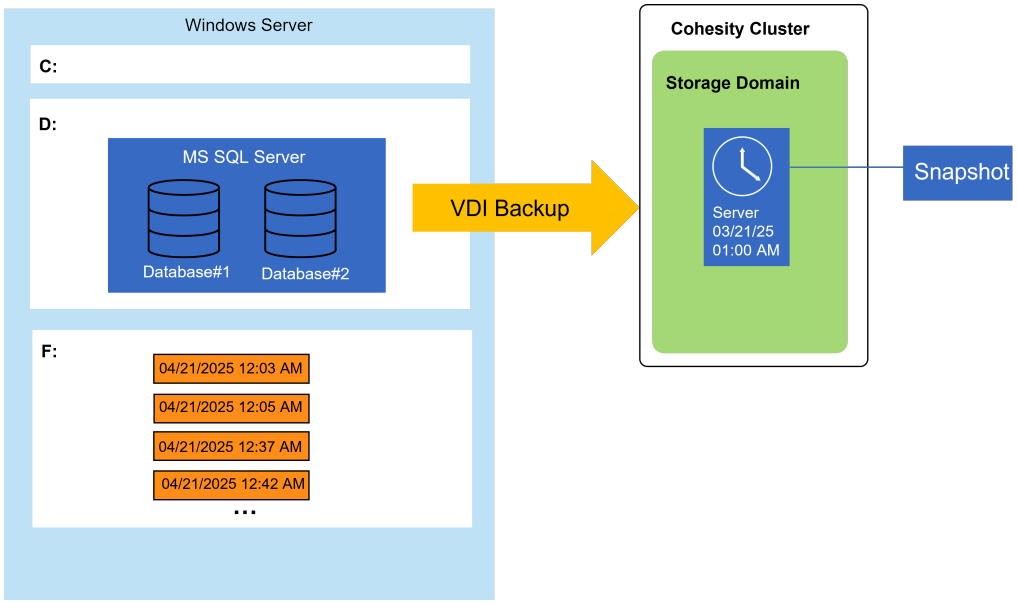

2. If the policy is configured to take incremental database backups, it translates to differential backups in case of VDI-based backup.
3. If transaction log backups are specified in the policy, the database transaction logs are backed up to the Cohesity cluster using MS SQL Server Virtual Device Interface (VDI). The logs allow you to roll back to a specific point in time. MS SQL Server automatically truncates the SQL logs after a successful SQL log backup and marks the Virtual Log Files (VLFs) for reuse.

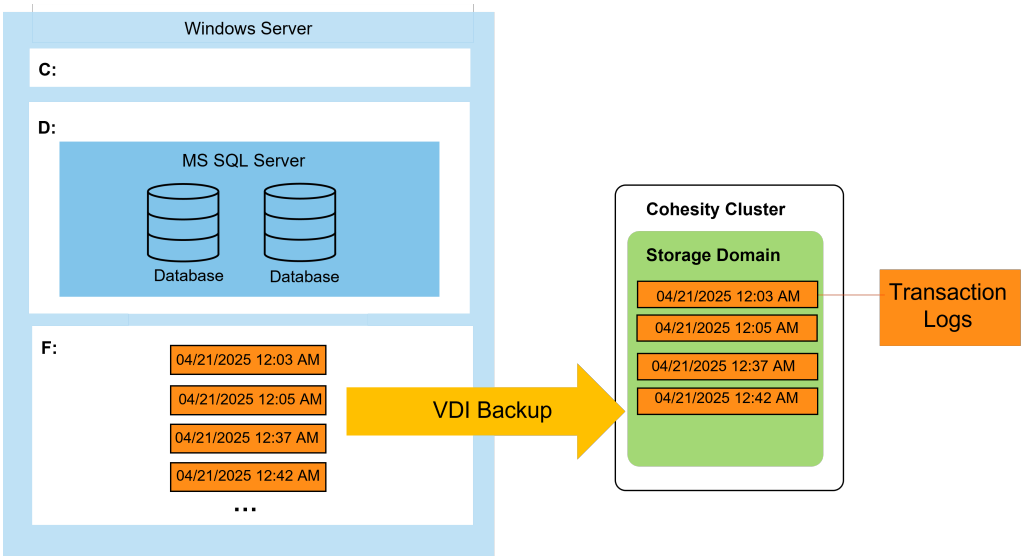

Compare File-Based, VDI-Based, Volume-Based Backup Methods

The following table compares the implications of each method.

Cohesity Facet	File-based backup	VDI-based backup	Volume-based backup
Technology Employed	Cohesity Agent uses Vss and Cohesity technology for file CBT, and VDI for T-logs.	Cohesity agent uses SQL Native backup commands.	Cohesity Agent uses Vss, for volume CBT, and VDI for T-logs.

Cohesity Facet	File-based backup	VDI-based backup	Volume-based backup
What does it capture?	Captures database files.	Captures the database in a SQL Server standard .BAK, .TRN, or .Diff file.	Captures all the volumes attached to the server and everything on those volumes. **Note:** Often, there are other files, like old .bak files, that sit on the volume along with the active database files. A large amount of other files will increase the duration of the protection group because these files will be captured along with the database files.
Logs	Yes. Captures a SQL log backup of the database.	Yes. Captures a SQL Server native log backup of the database.	Yes. Captures a SQL log backup of the database.
Cohesity Source Registration	Register as a SQL Server application.	Register as a SQL Server application.	Register as a SQL Server application.

<table>
<tr><th>Cohesity Facet</th><th>File-based backup</th><th>VDI-based backup</th><th>Volume-based backup</th></tr>
<tr><td>Compatibility with SQL Server High Availability Features</td><td colspan="3">Windows Failover Cluster: Yes, backup will succeed if a failover occurs. Must be registered as a Cohesity "SQL Cluster" source.

AG: Yes, backup on the correct replica will succeed after an AG failover occurs.

Mirroring: Yes Each SQL instance must have its own Protection Group.

Log Shipping: Yes. Each SQL instance must have its own Protection Group.

Note: While Cohesity currently does not include specific features for mirroring and log shipping, it does not conflict with these SQL Server features.

Note: Log shipping via SQL Server Management Studio supports only full/incremental backup and not log backup since it can break the log chain.</td></tr>
<tr><td>Method</td><td>This is an MS SQL Server database backup at the file level. Allows for FULL SQL Server backups with logs, but only the database files are captured. The Cohesity Agent uses Vss, the Cohesity File CBT, and VDI to take an application-consistent snapshot of the volumes on that server. Faster because it's not backing up extraneous files.

Note: Cannot browse files or folders as in a volume-based backup.</td><td>This is a native MS SQL Server database backup at the file level. Allows for FULL SQL Server backups with logs and Differentials, but only the database is captured.

Agent uses SQL Server native backup and restore commands , to take an application-consistent backup of the database.</td><td>This is an MS SQL Server database backup at the volume level. Allows for FULL SQL Server backups – all databases, logs, and files/folders on the volumes are backed up. MS SQL Server is registered as a Cohesity source and MS SQL application. The Cohesity Agent uses Vss, the Cohesity Volume CBT, and VDI to take an application-consistent snapshot of the volumes on that server.</td></tr>
<tr><td>Recovery/Restore</td><td colspan="3">The database host server should be backed up and restored via a different protection group.

Supports Point-in-Time (PIT) restore of the database from across the entire backup history.</td></tr>
</table>

Cohesity Facet	File-based backup	VDI-based backup	Volume-based backup
Cloning	**Yes**. Can produce a representative database. *Advantage*: Applications can be tested with data that is as close to the production environment as possible: Can reproduce a database. *Advantage*: Applications can be tested with databases that can be destroyed and re-created on demand. No load on Production servers. The backup is NOT modified and therefore always safe. A successful clone is immediate validation of the backup's integrity. Uses storage resources on Cohesity cluster.	**Cloning is scheduled as a future enhancement.**	**Yes**. Can produce a representative volume. *Advantage*: Applications can be tested with data that is as close to the production environment as possible. Can reproduce a volume. *Advantage*: Applications can be tested in an environment that can be destroyed and recreated on demand. No load on Production servers. The backup is NOT modified and therefore always safe. Uses storage resources on Cohesity cluster.
Source-side deduplication	Supported	Not supported	Not supported
Archival	**CloudArchive**: Cohesity supports the ability to send backups to public cloud providers like AWS, Azure, and Google Cloud Platform. It also supports search and restore from them, back to the local site when needed.		
Replication	**Managing Copies of backups on another Cohesity cluster**: Backups, files, and volumes can be replicated to a second Cohesity cluster.		

About Log and Full DB Backup Retention Periods

The retention periods of log files and full database server backups are independent of each other. Consider the following examples to help determine the best retention periods for your database server backups. These retention guidelines also apply to full database server and log replication schedules. Retention periods of log files and full database backups must adhere to Microsoft rules for restoring databases.

1. A point-in-time restore using SQL log backups is dependent on a full backup because a full database server backup must also be available to apply the logs.

   When the Cohesity cluster restores to a specific point in time, it must first restore a full database server backup and then apply the appropriate set of logs. If a full database server backup is not available because the retention period of the full database server backup is too short, the Cohesity cluster cannot restore to a point in time within the specified log retention period. For example if the log retention period is three weeks and the full database server backup captures snapshots weekly with a retention period of two weeks, the Cohesity cluster will not be able to restore to any point in time from three weeks ago. In addition, the logs from three weeks ago cannot be utilized and are using disk space.

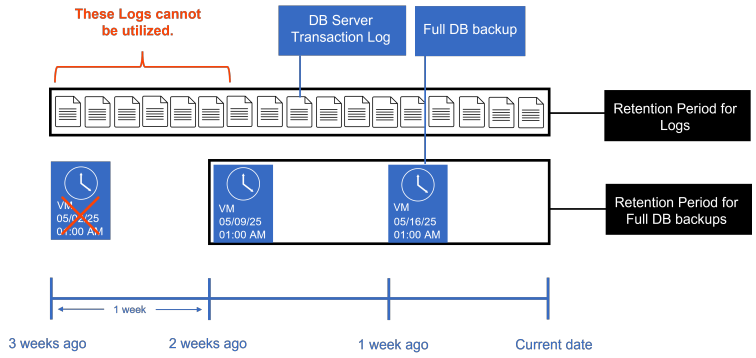

2. If the log retention period is shorter than the full database server backup retention period, the Cohesity cluster cannot restore to a point in time where the logs are not saved. For example, if the log retention period is two weeks and the full database server backup captures snapshots weekly with a retention period of three weeks, the Cohesity cluster will not be able to restore to any point in time between two and three weeks ago. The Cohesity cluster can, however, restore to the snapshot from three weeks ago.

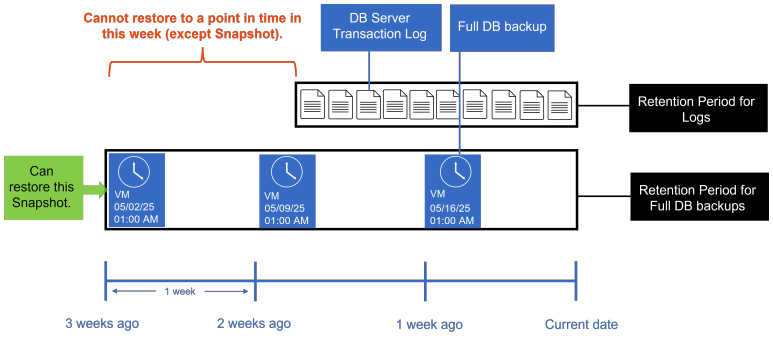

3. Cohesity recommends specifying a log retention period with extra time to allow for failed full database server backups as demonstrated by the following example. A full database server backup is scheduled weekly with a three-week retention period for both full database server and logs. If the full database server backup failed two weeks ago, the Cohesity cluster can still restore to the point in time from the log indicated by

the green arrow.

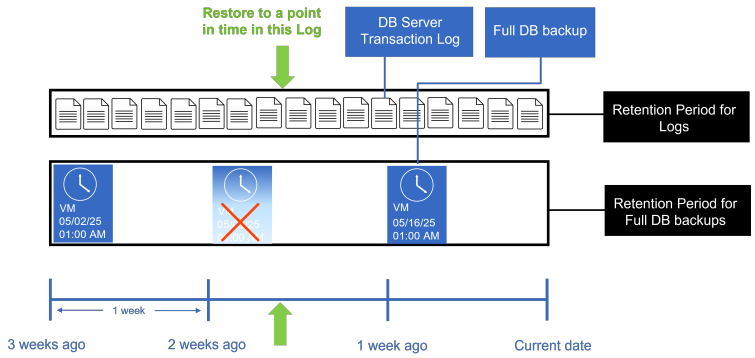

The Cohesity cluster restores full database server backup from three weeks ago and applies the appropriate logs to restore the state to a point.

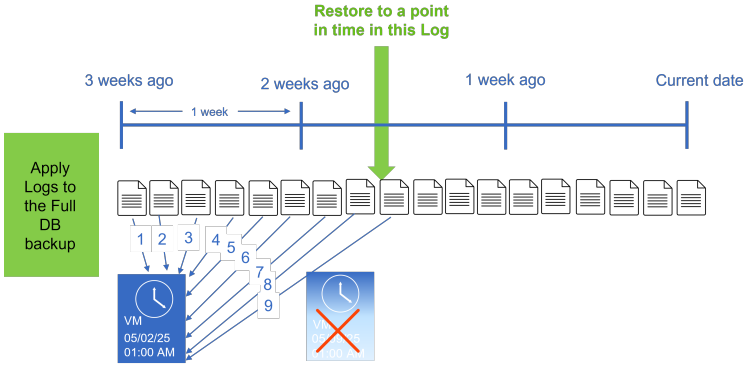

You can restore to a specific transaction in the log.

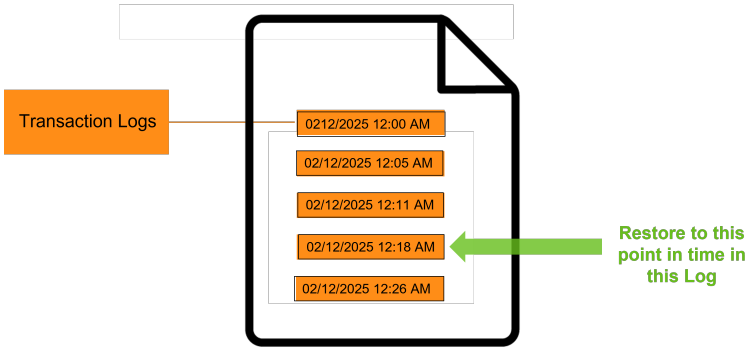

> **Note:** A full database backup is not a transaction log backup. Ensure that your SQL Server protection policy has transaction log backup for full and bulk-logged recovery model requirements.

Backup Microsoft SQL Server (Volume-based)

This topic describes how to protect MS SQL with a volume-based job. Volume-based jobs protect MS SQL at the server level, meaning all databases on a server. File-based jobs protect only the specific MS SQL databases that you select. To use a file-based job, follow the steps described in Backup Microsoft SQL Server (File-based).

MS SQL Protection Group Support

The following table describes supported job types for MS SQL running on physical servers and VMs.

	File-based Jobs	Volume-based Jobs	VDI-based Jobs	Both Job Types Using the Same Server but Protecting Distinct Databases
Physical MS SQL	Multiple file-based jobs protecting all or a subset of DBs	One volume-based job protecting all or a subset of DBs	Multiple VDI-based jobs protecting all or a subset of databases	Supported
SQL Server Virtual Machine Backup	NA	Volume-base with Hypervisor to provide application-consistent backup	NA	NA

Note:

- (*Applies to both standalone MS SQL and MS SQL FCI*.) If the VM SQL you registered has physical VMware RDM disks attached, then the MS SQL databases located on these disks will not be backed up and the following message is displayed in the Protection page:

  ```
  Identified <Number> physical RDM disks attached to this VM. Backup of physical RDM disks is not supported using VADP. Databases located on any of these disks are not backed up.
  ```

- (*Applies only to both standalone MS SQL and MS SQL FCI*.) If the VM SQL you registered has virtual VMware RDM disks attached, then the MS SQL databases located on these disks will not be backed up and the following message is displayed in the Protection page:

  ```
  Identified < Number> virtual RDM disks attached to this VM. Backup of virtual RDM disks is not supported. Databases located on any of these disks are not backed up.
  ```

- (*Applies only to standalone MS SQL*.) If the VM VM SQL you registered has iSCSI devices attached and if there are SQL databases having their database files on these iSCSI disks, then those SQL databases will not be backed up. The following message is displayed on the Protection page:

  ```
  Found < Number> databases with files on external storage. These databases cannot be backed up. Please use alternate backup methods like native dumps, Cohesity file base backup, or Cohesity's SQL VDI to protect the databases
  ```

Note: Converting an existing volume-based Protection Group to a file-based Protection Group is not supported. You must delete the existing Protection Group and create a new one.

Protect MS SQL Overview

Perform the following steps in the Cohesity cluster to protect SQL Server:

1. Register the SQL Server you want to protect as a source with the Cohesity cluster. See Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs or Set Up Microsoft SQL Server FCI.
2. Select **Dashboard** and from the drop-down, select **MS SQL**.

3. From the **Hosts** tab, you can protect all databases on a host. Select **Protect Databases** from the host's actions menu.

   From the **Databases** tab, you can protect a single database. Select **Protect** from the database's actions menu.

   You can change which objects you want to protect when creating the job.

You can also protect MS SQL by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. Select the objects to protect.

Create a Volume-based MS SQL Protection Group

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. In the **New Protection** pop-up, do one of the following:
   - Protect MS SQL by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered MS SQL source or click Register Source to register an MS SQL source. Select the objects to protect .
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. **Backup Type**: From the drop-down, select **Volume-based**.
     vi. Click **Protect**.
   - To protect MS SQL by specifying all the options, click **More Options**.
4. **Source:** Select the objects to protect.

   Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects

and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.

> **Note:** If you have registered a physical server that is a part of the registered MS SQL Failover Cluster, then the MS SQL Failover Cluster Instance will not appear as a source in the object list of the registered physical server.

If appropriate for your organization, enable Auto Protect.

Click to turn on **Auto Protect**. By enabling this option:

- All the SQL databases that are added to the SQL server in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the SQL databases that are removed from the SQL server in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the user is preserved until the snapshot expires.

For example, you create a Protection Group that backs up a SQL Server. The Protection Group runs for two weeks and then a new database is added to the SQL Server. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new database is also backed up even though the new database has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours.

To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Use this option to exclude SQL databases from the protection group. Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

  To exclude a database from the protection group, specify the absolute path of the database to be excluded from the backup at a protection job level. You can optionally use:

  - **Wildcard Exclusion** - You can optionally use the Wildcard * in any location in the search string to exclude all the matching databases from the backup at a protection job level.

- **Regex Exclusion** - You can use the RegEx string name match and manual user selection to exclude DBs from backup at a protection job level.

> **Note:** You can use the exclude SQL database feature only if Auto Protect is enabled for a backup job.

Protection of dismounted databases, recovery database types, and deleted databases are automatically skipped by the Protection Group.

After defining the filer click **Apply** to save the selection.

The following table lists the examples for excluding objects in an MS SQL database:

Sample Configuration		
10.2.157.20 (Host)		
*** SQL2012**	*** SQL2014**	*** MSSQLSERVER**
(SQL Server Instance)	(SQL Server Instance)	(SQL Server Instance)
- UserDB1	- UserDB3	- ProdDB1
- UserDB2	- UserDB4	- ProdDB2
- TestDB1	- TestDB3	
- TestDB2	- TestDB4	

Example of Absolute Path

- **10.2.157.20/SQL2012/UserDB1** - Excludes the database 'UserDB1' from the SQL server instance 'SQL2012' on the host 10.2.157.20.
- **10.2.157.20/SQL2012/** - Excludes the 'SQL2012' from the SQL server instance and its databases.

> **Note:** The absolute instance path must end with '/' and the plain text instance path must contain two or more slashes.

Example of Wildcard Exclusion

- **10.2.157.20/SQL2012/User*** - Excludes the databases 'UserDB1' and 'UserDB2' from the SQL server instance 'SQL2012' on the host 10.2.157.20.
- **10.2.157.20/SQL201?/*** - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.

Example of Regex Exclusion

10.2.157.20/SQL.*/.* - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.

Considerations

Consider the following when using the **Auto Protect** option on the backup job

- If new databases are added to the host, the existing filters are applied to the newly added databases. For example, if the exclude filter matches the newly added databases then these databases will be excluded in the backup job.
- If any system databases are excluded by the filters in an SQL server, it will exclude all the system databases belonging to that SQL server.
- When defining filters to exclude AG databases, ensure that the filters are configured on all the AG hosts. For example, to exclude an AG database and if the AG database has 4 hosts, then ensure that the filters are defined on all the four AG hosts.
- Auto protect functionality is supported only for physical SQL jobs (file-based, Native, and volume-based)
- To selectively exclude a SQL database from the protection job, you must upgrade all the cluster nodes to version 6.5.1 and later.

If MS SQL belongs to an AG network, selecting that server displays options for protecting the AG.

- Select all servers that belong to this AG network.

  > **Note:** If backing up VMs within an AG, protecting the AG databases requires that all servers within the AG are registered with the Cohesity cluster and are selected for protection. Otherwise AG databases are skipped and only local databases will be backed up. New servers added to AG are not automatically protected.

  When using this option when there are multiple AGs, only interlinked AG nodes are automatically selected. For example:

  - If there are AGs that use only Node1 and Node2 and other AGs that use Node3 and Node4, if you select Node2 as a source, only Node1 is added automatically.
  - If AG1 uses Node1 and Node2, AG2 uses Node2 and Node4, AG3 uses Node3 and Node5. If you select Node2, then these nodes are selected: Node1, Node2 and Node4. Node3 and Node5 are not selected.

  When adding nodes to an AG that is protected by Cohesity, ensure you register the node with Cohesity and add it to the Protection Group. Otherwise the AG databases will not be protected.

  When removing nodes from an AG that is protected by Cohesity, you can remove it from the Protection Group or keep it in the job; the AG databases remain protected in either case. If the removed node remains in the Protection

Group, only its local databases are backed up. If you do not want to back up the removed node, you can safely remove it.

- Select only this server.

Click **Select**.

Click **Add**.

5. If appropriate, set options for specific objects by clicking (**Edit Server Options**) next to the object. A successful SQL log backup automatically truncates the SQL logs.
6. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

   > **Note:** This page displays registered MS SQL objects only, so only MS SQL servers are selected regardless of the other servers within the source object hierarchy.

7. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

   The selected storage domain stores the snapshots and logs created by this Protection Group.
8. **Backup Method:** Select **Volume-based**. This method protects MS SQL at the server level, meaning all databases on a server.

   If you changed the backup method to volume-based after objects were selected, you must select a server object.
9. **Make Full Backup Copy-only:** Toggle on if you want full backups to be copy-only backups so they do not affect the differential base.

   Copy-only full backups do not take log backups even if they are scheduled by the policy.
10. **User Databases:** Select how to handle AG databases during backup.
11. **AG Backup Preferences:** Select this option if AG databases will be backed up.

- **Use Server Preferences** uses MS SQL preferences.
- **Override Preferences** enables you to override MS SQL's preferences with your selection.

AAG replica preference for MS SQL backups

Cohesity uses "replica priority" to select the best replica when more than one qualified replica matches the backup preference. For the following AAG Backup Preference settings, Cohesity uses replicas in the indicated order of preference to back up MS SQL databases:

Backup Preference Setting	Replica Used for MS SQL Backup
Prefer Secondary or Any	1. Sync Secondary Replica 2. Async Secondary Replica 3. Primary Replica
Secondary Only*	1. Sync Secondary Replica 2. Async Secondary Replica
Primary Only	1. Primary Replica

* If the AAG uses the **Secondary Only** Backup Preference setting, ensure the AAG replicas are set to "**Readable secondary**=**Yes**". If all replicas are set to "**Readable secondary**=**No**" with **Secondary Only** as the Backup Preference setting, then the backup of AAG databases fails with "*unavailable_for_vss*".

You can set the **Readable Secondary** field to **No** if the Backup Preference Setting is set to **Prefer Secondary or Any** or **Primary Only**.

> **Note:** You must apply the exclusion filter on each AAG replica to ensure the database is excluded when an SQL Server failover occurs.

12. **System Databases:** Select whether to back up or skip system databases.
13. **Volume Backup:** Optionally select **Back up only the volumes associated with selected databases** to back up only volumes that are associated with selected databases. This option is not used for MS SQL FCI; only volumes with databases are backed up. For MS SQL running on VMs, all volumes are always backed up.
14. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.

15. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
16. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
17. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
18. **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

    For configuration details, see Configure Pre & Post Scripts.
19. **Incremental After Restart:** For Windows physical servers, toggle on if you prefer incremental backup for the first scheduled backup job, which runs after an intentional server restart. If an incremental backup is not possible, a full backup is performed. For example, if the server restarts after a power failure, a full backup is performed.

> **Note:** Incremental backup consistency across reboots is not guaranteed if I/O operations are performed on the volume outside of the fully functional state of the server. This includes I/O operations that are performed by booting into the recovery mode. For example, if a server is shut down and its volume is attached to another server that does not have the Cohesity agent, ensure that the next backup is a full backup operation.

This option is not available for file-based backups.

20. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.

    Indexing is not supported for file-based Jobs. Without indexing, file/folder-level recovery and searching are not supported.

21. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:

    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.

22. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.

23. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

> **Note:**
> The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

24. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
25. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

In the Protection page, the following message will be displayed when you mouse over the run status of the MS SQL databases that are located on the attached disks:

```
<database_name> has files located on external storage. Cannot backup the database. Please use alternate backup methods like native dumps to protect the databases.
```

Edit an MS SQL Protection Group

1. Select **Data Protection > Protection**.
2. Locate the Protection Group you want to edit.
3. Mouse over the action icon (⋮) next to the job and select one of the following:
    1. **Run Now** - Start a run of this Protection Group. You have the option to back up one or more objects.
        1. **Backup all Objects in the Protection Group** - Select this option to backup all the databases in the protection group. Any databases excluded from the backup will not be included in the protection job run.
        2. **Backup only Selected Objects in the Protection Group** - you can manually select the databases to be backed up at the protection job level. However, the selection of the database that is excluded at the backup job level is not supported.

        You can run log backups in parallel with existing log backup runs. This helps reduce storage buildup on the primary SQL Server when multiple large backups are in progress.

        This feature is supported only for file-based, VDI-based, and Filestream (backed up using VDI-based method) backups.

> **Note:** When you run a log backup for a database backup job, any preceding backup job for the same database may skip the log backup if the new backup job is completed before the previous backup job for the database is started.

2. **Pause Future Runs**
    1. Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the job is currently executing a run, the current run continues to execute and only future runs of this job are paused.
    2. Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this job.
3. **Edit** - Edit job settings of the protection group and click Protect.
4. **Delete** - Deletes the protection group.

4. Save the job settings by clicking **Protect**.

Backup Microsoft SQL Server (File-based)

This topic describes how to protect MS SQL with a file-based Protection Group. File-based Protection Groups protect only the specific MS SQL databases that you select. Volume-based Protection Groups protect MS SQL at the server level, meaning all databases on a server. To use a volume-based Protection Group, follow the steps described in Backup Microsoft SQL Server (Volume-based).

MS SQL Protection Group Support

The following table describes supported job types for MS SQL running on physical servers and VMs.

	File-based Jobs	Volume-based Jobs	VDI-based Jobs	Both Job Types Using the Same Server but Protecting Distinct Databases
Physical MS SQL	Multiple file-based jobs protecting all or a subset of DBs	One volume-based job protecting all or a subset of DBs	Multiple VDI-based jobs protecting all or a subset of databases	Supported
SQL Server Virtual Machine Backup	NA	Volume-base with Hypervisor to provide application-consistent backup	NA	NA

Protect MS SQL Overview

Perform the following steps in the Cohesity cluster to protect SQL Server:

1. Register the SQL Server you want to protect as a source with the Cohesity cluster. See Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs or Set Up Microsoft SQL Server FCI.
2. Select **Dashboard** and from the drop-down, select **MS SQL**.
3. From the **Hosts** tab, you can protect all databases on a host. Select **Protect Databases** from the host's actions menu.

   From the **Databases** tab, you can protect a single database. Select **Protect** from the database's actions menu.

   You can change which objects you want to protect when creating the job.

You can also protect MS SQL by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. Select the objects to protect.

Create a File-based MS SQL Protection Group

Perform the following to create a file-based SQL Server Protection Group:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. In the **New Protection** pop-up, do one of the following:
   - Protect MS SQL by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered MS SQL source or click Register Source to register an MS SQL source. Select the objects to protect.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

v. **Backup Type**: From the drop-down, select **File-based**.

vi. Click **Protect**.

- To protect MS SQL by specifying all the options, click **More Options**.

4. **Source:** Select the objects to protect.

   Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.

   > **Note:** If you have registered a physical server that is a part of the registered MS SQL Failover Cluster, then the MS SQL Failover Cluster Instance will not appear as a source in the object list of the registered physical server.

   If appropriate for your organization, enable Auto Protect.

   Click to turn on **Auto Protect**. By enabling this option:

   - All the SQL databases that are added to the SQL server in the future are automatically added to the Protection Group and are protected from the next protection run.
   - All the SQL databases that are removed from the SQL server in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the user is preserved until the snapshot expires.

   For example, you create a Protection Group that backs up a SQL Server. The Protection Group runs for two weeks and then a new database is added to the SQL Server. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new database is also backed up even though the new database has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours.

   To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click .

   Considerations

   Consider the following when using the **Auto Protect** option on the backup job

   - If new databases are added to the host, the existing filters are applied to the newly added databases. For example, if the exclude filter matches the newly added databases then these databases will be excluded in the backup job.

- If any system databases are excluded by the filters in an SQL server, it will exclude all the system databases belonging to that SQL server.
- When defining filters to exclude AG databases, ensure that the filters are configured on all the AG hosts. For example, to exclude an AG database and if the AG database has 4 hosts, then ensure that the filters are defined on all the four AG hosts.
- Auto protect functionality is supported only for physical SQL jobs (file-based, Native, and volume-based)
- To selectively exclude a SQL database from the protection job, you must upgrade all the cluster nodes to version 6.5.1 and later.

If MS SQL belongs to an AG network, selecting that server displays options for protecting the AG.

- Select all servers that belong to this AG network.

  > **Note:** If backing up VMs within an AG, protecting the AG databases requires that all servers within the AG are registered with the Cohesity cluster and are selected for protection. Otherwise AG databases are skipped and only local databases will be backed up. New servers added to AG are not automatically protected.

  When using this option when there are multiple AGs, only interlinked AG nodes are automatically selected. For example:

  - If there are AGs that use only Node1 and Node2 and other AGs that use Node3 and Node4, if you select Node2 as a source, only Node1 is added automatically.
  - If AG1 uses Node1 and Node2, AG2 uses Node2 and Node4, AG3 uses Node3 and Node5. If you select Node2, then these nodes are selected: Node1, Node2 and Node4. Node3 and Node5 are not selected.

  When adding nodes to an AG that is protected by Cohesity, ensure you register the node with Cohesity and add it to the Protection Group. Otherwise the AG databases will not be protected.

  When removing nodes from an AG that is protected by Cohesity, you can remove it from the Protection Group or keep it; the AG databases remain protected in either case. If the removed node remains in the Protection Group, only its local databases are backed up. If you do not want to back up the removed node, you can safely remove it.

- Select only this server.

Click **Select**.

Click **Add**.

5. If appropriate, set options for specific objects by clicking (**Edit Server Options**) next to the object. A successful SQL log backup automatically truncates the SQL logs.
6. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

   > **Note:** This page displays registered MS SQL objects only, so only MS SQL servers are selected regardless of the other servers within the source object hierarchy.

7. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.
8. **Storage Domain:** Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

   The selected storage domain stores the snapshots and logs created by this Protection Group.
9. **Backup Method:** Select **File-based**. This method protects only the specific MS SQL databases that you select.

   If you want to change the backup method to volume-based after objects are selected, you must select a server object. A file-based Protection Group can protect any object.
10. **Make Full Backup Copy-only:** Toggle on if you want full backups to be copy-only backups so they do not affect the differential base.

    Copy-only full backups do not take log backups even if they are scheduled by the policy.
11. **Log Chain Break Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the log chain is broken.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.
12. **New Database Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the new database is detected.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.

13. **User Databases:** Does not apply to file-based Protection Groups. Only selected databases are backed up.
14. **AG Backup Preferences:** Select this option if AG databases will be backed up.
    - **Use Server Preferences** uses MS SQL preferences.
    - **Override Preferences** enables you to override MS SQL's preferences with your selection.

    AAG replica preference for MS SQL backups

    Cohesity uses "replica priority" to select the best replica when more than one qualified replica matches the backup preference. For the following AAG Backup Preference settings, Cohesity uses replicas in the indicated order of preference to back up MS SQL databases:

Backup Preference Setting	Replica Used for MS SQL Backup
Prefer Secondary or Any	1. Sync Secondary Replica 2. Async Secondary Replica 3. Primary Replica
Secondary Only*	1. Sync Secondary Replica 2. Async Secondary Replica
Primary Only	1. Primary Replica

    * If the AAG uses the **Secondary Only** Backup Preference setting, ensure the AAG replicas are set to "**Readable secondary**=**Yes**". If all replicas are set to "**Readable secondary**=**No**" with **Secondary Only** as the Backup Preference setting, then the backup of AAG databases fails with "*unavailable_for_vss*".

    You can set the **Readable Secondary** field to **No** if the Backup Preference Setting is set to **Prefer Secondary or Any** or **Primary Only**.

    > **Note:** You must apply the exclusion filter on each AAG replica to ensure the database is excluded when an SQL Server failover occurs.

15. **System Databases:** Does not apply to file-based Protection Groups. Only selected databases are backed up. When you select any system database, however, all system databases are selected automatically.
16. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up

and down arrows on your keyboard. Verify the AM or PM setting.

The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.

17. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
18. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
19. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
20. **Source Side Deduplication**: Toggle ON **Source Side Deduplication** and select the **Exclusions** from the drop-down, to optimize the full backup time when CBT data is lost. The exclusions are supported at the source level. When you select a source as exclusion, source side deduplication will not be applicable for that source in the protection group.
21. **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

    For configuration details, see Configure Pre & Post Scripts.

22. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.

    Indexing is not supported for file-based Protection Groups. Without indexing, file/folder-level recovery and searching are not supported.

23. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:

    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.

24. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.

25. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

    > **Note:**
    > The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

26. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.

    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).

27. Click **Protect**.

The Protection page is displayed and includes the new Protection Group.

Edit an MS SQL Protection Group

Perform the following to edit a SQL Server Protection Group:

1. Select **Data Protection > Protection**.
2. Locate the Protection Group you want to edit.
3. Mouse over the action icon (⋮) next to the Protection Group and select one of the following:
    1. **Run Now** - Start a run of this Protection Group. You have the option to back up one or more objects.
        1. **Backup all Objects in the Protection Group** - Select this option to backup all the databases in the protection group. Any databases excluded from the backup will not be included in the protection run.
        2. **Backup only Selected Objects in the Protection Group** - you can manually select the databases to be backed up at the Protection Group level.
        If the host (parent of the selected database) is already part of a previous run-now submitted, the backup will be queued until the previous run submitted is complete.

            > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

        You can run log backups in parallel with existing log backup runs. This helps reduce storage buildup on the primary SQL Server when multiple large backups are in progress.

        This feature is supported only for file-based, VDI-based, and Filestream (backed up using VDI-based method) backups.

        > **Note:** When you run a log backup for a database backup job, any preceding backup job for the same database may skip the log backup if the new backup job is completed before the previous backup job for the database is started.

    2. **Pause Future Runs**
        1. Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the Protection Group is currently executing a run, the current run continues to execute and only future runs of this

Protection Group are paused.

2. Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this Protection Group.

3. **Edit** - Edit the Protection Group settings of the protection group and click Protect.
4. **Delete** - Deletes the protection group.

4. Save the Protection Group settings by clicking **Protect**.

Deduplication in SQL Server File-based Backup

Cohesity supports an efficient Source-Side Deduplication mechanism (to detect changed blocks) for SQL Server File-based backups, that tracks the block-level changes and maintains CPU efficiency even when CBT (Change Block Tracking) is unavailable.

When a host reboot occurs during backup, the CBT map is cleared, and the CBT data is lost. This forces you to run a full backup instead of an incremental backup which can be very time-consuming for large SQL Server databases.

With this deduplication mechanism enabled for protection groups, if the CBT data is lost due to a Windows host server restart, the changed blocks are identified and transferred from the source. This reduces backup time and ensures efficient CPU utilization.

After the next full backup, CBT is re-established and used for subsequent backups.

> **Note:** Source-side deduplication functionality increases the CPU usage during backup operation (when CBT is lost).

The functionality is only supported for new protection groups. The backup of existing protection groups that already use source-side deduplication will complete with a warning which includes the instructions to use the efficient version of source-side deduplication.

You can set the deduplication for a protection group using the Source Side Deduplication option in the Additional Settings of a protection.

Backup Microsoft SQL Server (VDI-based)

This topic describes how to protect MS SQL with a VDI-based Protection Group.

MS SQL Protection Group Support

The following table describes supported job types for MS SQL running on physical servers and VMs.

	File-based Jobs	Volume-based Jobs	VDI-based Jobs	Both Job Types Using the Same Server but Protecting Distinct Databases
Physical MS SQL	Multiple file-based jobs protecting all or a subset of DBs	One volume-based job protecting all or a subset of DBs	Multiple VDI-based jobs protecting all or a subset of databases	Supported
SQL Server Virtual Machine Backup	NA	Volume-base with Hypervisor to provide application-consistent backup	NA	NA

Protect MS SQL Overview

Perform the following steps in the Cohesity cluster to protect SQL Server:

1. Register the SQL Server you want to protect as a source with the Cohesity cluster. See Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs or Set Up Microsoft SQL Server FCI.
2. Select **Dashboard** and from the drop-down, select **MS SQL**.
3. From the **Hosts** tab, you can protect all databases on a host. Select **Protect Databases** from the host's actions menu.

   From the **Databases** tab, you can protect a single database. Select **Protect** from the database's actions menu.

   You can change which objects you want to protect when creating the job.

You can also protect MS SQL by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. Select the objects to protect.

Create a VDI-based MS SQL Protection Group

Perform the following to create a VDI-based SQL Server Protection Group:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. In the **New Protection** pop-up, do one of the following:

- Protect MS SQL by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
  i. **Add Objects**: Select the registered MS SQL source or click Register Source to register an MS SQL source. Select the objects to protect.
  ii. **Protection Group**: Specify a name for the Protection Group.
  iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
  iv. **Storage Domain**: Select the appropriate storage domain.

  > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

  v. **Backup Type**: From the drop-down, select **VDI-based**.
  vi. Click **Protect**.
- To protect MS SQL by specifying all the options, click **More Options**.

4. **Source:** Select the objects to protect.

   Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.

   > **Note:** If you have registered a physical server that is a part of the registered MS SQL Failover Cluster, then the MS SQL Failover Cluster Instance will not appear as a source in the object list of the registered physical server.

   If appropriate for your organization, enable Auto Protect.

   Click to turn on **Auto Protect**. By enabling this option:

   - All the SQL databases that are added to the SQL server in the future are automatically added to the Protection Group and are protected from the next protection run.
   - All the SQL databases that are removed from the SQL server in the future are automatically removed from the Protection Group and are not protected from

the next protection run. The existing backup for the user is preserved until the snapshot expires.

For example, you create a Protection Group that backs up a SQL Server. The Protection Group runs for two weeks and then a new database is added to the SQL Server. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new database is also backed up even though the new database has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours.

To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Use this option to exclude SQL databases from the protection group. Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

  To exclude a database from the protection group, specify the absolute path of the database to be excluded from the backup at a protection job level. You can optionally use:

  - **Wildcard Exclusion** - You can optionally use the Wildcard * in any location in the search string to exclude all the matching databases from the backup at a protection job level.
  - **Regex Exclusion** - You can use the RegEx string name match and manual user selection to exclude DBs from backup at a protection job level.

> **Note:** You can use the exclude SQL database feature only if Auto Protect is enabled for a backup job.

Protection of dismounted databases, recovery database types, and deleted databases are automatically skipped by the Protection Group.

After defining the filer click **Apply** to save the selection.

The following table lists the examples for excluding objects in an MS SQL database:

Sample Configuration

10.2.157.20 (Host)

*** SQL2012**

(SQL Server Instance)

- UserDB1
- UserDB2
- TestDB1
- TestDB2

*** SQL2014**

(SQL Server Instance)

- UserDB3
- UserDB4
- TestDB3
- TestDB4

*** MSSQLSERVER**

(SQL Server Instance)

- ProdDB1
- ProdDB2

Example of Absolute Path

- **10.2.157.20/SQL2012/UserDB1** - Excludes the database 'UserDB1' from the SQL server instance 'SQL2012' on the host 10.2.157.20.
- **10.2.157.20/SQL2012/** - Excludes the 'SQL2012' from the SQL server instance and its databases.

> **Note:** The absolute instance path must end with '/' and the plain text instance path must contain two or more slashes.

Example of Wildcard Exclusion

- **10.2.157.20/SQL2012/User*** - Excludes the databases 'UserDB1' and 'UserDB2' from the SQL server instance 'SQL2012' on the host 10.2.157.20.
- **10.2.157.20/SQL201?/*** - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.

Example of Regex Exclusion

10.2.157.20/SQL.*/.* - Excludes the SQL server instances 'SQL2012' and 'SQL2014' from the host 10.2.157.20.

Considerations

Consider the following when using the **Auto Protect** option on the backup job

- If new databases are added to the host, the existing filters are applied to the newly added databases. For example, if the exclude filter matches the newly added databases then these databases will be excluded in the backup job.
- If any system databases are excluded by the filters in an SQL server, it will exclude all the system databases belonging to that SQL server.
- When defining filters to exclude AG databases, ensure that the filters are configured on all the AG hosts. For example, to exclude an AG database and if the AG database has 4 hosts, then ensure that the filters are defined on all the four AG hosts.
- Auto protect functionality is supported only for physical SQL jobs (file-based, Native, and volume-based)

- To selectively exclude a SQL database from the protection job, you must upgrade all the cluster nodes to version 6.5.1 and later.

If MS SQL belongs to an AG network, selecting that server displays options for protecting the AG.

- Select all servers that belong to this AG network.

  > **Note:** If backing up VMs within an AG, protecting the AG databases requires that all servers within the AG are registered with the Cohesity cluster and are selected for protection. Otherwise AG databases are skipped and only local databases will be backed up. New servers added to AG are not automatically protected.

  When using this option when there are multiple AGs, only interlinked AG nodes are automatically selected. For example:

  - If there are AGs that use only Node1 and Node2 and other AGs that use Node3 and Node4, if you select Node2 as a source, only Node1 is added automatically.
  - If AG1 uses Node1 and Node2, AG2 uses Node2 and Node4, AG3 uses Node3 and Node5. If you select Node2, then these nodes are selected: Node1, Node2 and Node4. Node3 and Node5 are not selected.

  When adding nodes to an AG that is protected by Cohesity, ensure you register the node with Cohesity and add it to the Protection Group. Otherwise the AG databases will not be protected.

  When removing nodes from an AG that is protected by Cohesity, you can remove it from the Protection Group or keep it; the AG databases remain protected in either case. If the removed node remains in the Protection Group, only its local databases are backed up. If you do not want to back up the removed node, you can safely remove it.

- Select only this server.

Click **Select**.

Click **Add**.

5. If appropriate, set options for specific objects by clicking (**Edit Server Options**) next to the object. A successful SQL log backup automatically truncates the SQL logs.
6. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

> **Note:** This page displays registered MS SQL objects only, so only MS SQL servers are selected regardless of the other servers within the source object hierarchy.

7. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**. Cohesity recommends you schedule periodic full backups for MS SQL VDI-based protection. For more information, see Create or Edit a Standard Policy.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

   The selected storage domain stores the snapshots and logs created by this Protection Group.
8. **Backup Method:** Select **VDI-based**. This method protects only the specific MS SQL databases that you select.
9. **WITH Clause:** Define the With clause that you want to use to customize the backup. For more information, see BACKUP (Transact-SQL).
10. **Number of Streams:** Define the number of Cohesity nodes that will be used for the backup process. Using multiple nodes provides more throughput for the backup process.
11. **Log Chain Break Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the log chain is broken.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.
12. **New Database Auto-Trigger Out of Band Incremental Backup:** When toggled ON, the out of band incremental backup will be triggered at the end of the log backup when the new database is detected.

    If a scheduled backup is already running for the protection group, the out of band incremental backup will be triggered after the scheduled backup.
13. **Make Full Backup Copy-only:** Toggle on if you want full backups to be copy-only backups so they do not affect the differential base.

    Copy-only full backups do not take log backups even if they are scheduled by the policy.
14. **User Databases:** Select how to handle AG databases during the backup.

15. **AG Backup Preferences:** Select this option if AG databases will be backed up.
    - **Use Server Preferences** uses MS SQL preferences.
    - **Override Preferences** enables you to override MS SQL's preferences with your selection.

    AAG replica preference for MS SQL backups

    Cohesity uses "replica priority" to select the best replica when more than one qualified replica matches the backup preference. For the following AAG Backup Preference settings, Cohesity uses replicas in the indicated order of preference to back up MS SQL databases:

Backup Preference Setting	Replica Used for MS SQL Backup
Prefer Secondary or Any	1. Sync Secondary Replica 2. Async Secondary Replica 3. Primary Replica
Secondary Only*	1. Sync Secondary Replica 2. Async Secondary Replica
Primary Only	1. Primary Replica

    * If the AAG uses the **Secondary Only** Backup Preference setting, ensure the AAG replicas are set to "**Readable secondary**=**Yes**". If all replicas are set to "**Readable secondary**=**No**" with **Secondary Only** as the Backup Preference setting, then the backup of AAG databases fails with "*unavailable_for_vss*".

    You can set the **Readable Secondary** field to **No** if the Backup Preference Setting is set to **Prefer Secondary or Any** or **Primary Only**.

    > **Note:** You must apply the exclusion filter on each AAG replica to ensure the database is excluded when an SQL Server failover occurs.

> **Note:** AG secondary replicas do not support regular full backups, and differential backups. If the policy is configured to take regular full backups, or differential backups, then select override preferences to backup only from the primary replica. All other AG backup preferences are only available if the policy is configured to take copy-only full backups, and log backups. Incremental backups are not supported on the AG secondary nodes. For more details, see Backup Types Supported on Secondary Replicas.

16. **System Databases:** Does not apply to VDI-based Protection Groups. Only selected databases are backed up. When you select any system database, however, all system databases are selected automatically.
17. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
18. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
19. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

> **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

- **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
- **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
- **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
- **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.

20. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
21. **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

    For configuration details, see Configure Pre & Post Scripts.
22. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.

    Indexing is not supported for VDI-based Protection Groups. Without indexing, file/folder-level recovery and searching are not supported.
23. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.
24. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
25. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

> **Note:**
> The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

26. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
27. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Edit an MS SQL Protection Group

Perform the following to edit a SQL Server Protection Group:

1. Select **Data Protection > Protection**.
2. Locate the Protection Group you want to edit.
3. Mouse over the action icon (⋮) next to the Protection Group and select one of the following:
    1. **Run Now** - Start a run of this Protection Group. You have the option to back up one or more objects.
        1. **Backup all Objects in the Protection Group** - Select this option to backup all the databases in the protection group. Any databases excluded from the backup will not be included in the protection run.
        2. **Backup only Selected Objects in the Protection Group** - you can manually select the check box for the databases to be backed up at the Protection Group level.
           This backup of the selected database runs in parallel with the existing run-now request. However, if a subset of the selected database includes a run-now/scheduled backup already active, these databases will be skipped from the run-now being submitted.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

You can run log backups in parallel with existing log backup runs. This helps reduce storage buildup on the primary SQL Server when multiple large backups are in progress.

This feature is supported only for file-based, VDI-based, and Filestream (backed up using VDI-based method) backups.

> **Note:** When you run a log backup for a database backup job, any preceding backup job for the same database may skip the log backup if the new backup job is completed before the previous backup job for the database is started.

2. **Pause Future Runs**
    1. Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the Protection Group is currently executing a run, the current run continues to execute and only future runs of this Protection Group are paused.
    2. Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this Protection Group.
3. **Edit** - Edit the Protection Group settings of the protection group and click Protect.
4. **Delete** - Deletes the protection group.
   Delete only the protection group and keep the snapshots or delete the snapshots also.

4. Save the Protection Group settings by clicking **Protect**.

Backup Microsoft SQL Server Filestream Databases

A Filestream database stores unstructured large objects in the Windows file system, controlled by SQL Server. This approach reduces database size and improves performance for large binary content compared with storing the files directly in the database tables.

This topic describes how to protect MS SQL Filestream databases.

MS SQL Protection Group Support

The following table describes supported job types for MS SQL running on physical servers and VMs.

	File-based Jobs	Volume-based Jobs	VDI-based Jobs	All Job Types Using the Same Server but Protecting Distinct Databases
Physical MS SQL	Multiple file-based jobs protecting all or a subset of Filestream and non-Filestream databases	One volume-based job protecting all or a subset of DBs	Multiple VDI-based jobs protecting all or a subset of databases	Supported
VM MS SQL	Not supported	Volume-based jobs protecting the entire VM	Not supported	Not supported

Protect MS SQL Overview

Perform the following steps in the Cohesity cluster to protect SQL Server:

1. Register the SQL Server you want to protect as a source with the Cohesity cluster. See Set Up Standalone Microsoft SQL Server or Microsoft SQL Server AGs or Set Up Microsoft SQL Server FCI.
2. Select **Dashboard** and from the drop-down, select **MS SQL**.
3. From the **Hosts** tab, you can protect all databases on a host. Select **Protect Databases** from the host's actions menu.

   From the **Databases** tab, you can protect a single database. Select **Protect** from the database's actions menu.

   You can change which objects you want to protect when creating the job.

You can also protect MS SQL by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Databases** > **MS SQL Server**.
3. Select the objects to protect.

Create an MS SQL Protection Group

1. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

> **Note:** This page displays registered MS SQL objects only, so only SQL servers are selected regardless of the other servers within the source object hierarchy.

2. **Source:** Select the objects to protect.

   Select the source objects to back up or exclude. By default, all objects in the source are excluded as indicated by the ☐ unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box (☑) indicates the object is selected to be backed up.

   If MS SQL belongs to an AG network, selecting that server displays options for protecting the AG.

   - Select all servers that belong to this AG network.

     > **Note:** If backing up VMs within an AG, protecting the AG databases requires that all servers within the AG are registered with the Cohesity cluster and are selected for protection. Otherwise AG databases are skipped and only local databases will be backed up. New servers added to AG are not automatically protected.

     When using this option when there are multiple AGs, only interlinked AG nodes are automatically selected. For example:

     - If there are AGs that use only Node1 and Node2 and other AGs that use Node3 and Node4, if you select Node2 as a source, only Node1 is added automatically.
     - If AG1 uses Node1 and Node2, AG2 uses Node2 and Node4, AG3 uses Node3 and Node5. If you select Node2, then these nodes are selected: Node1, Node2 and Node4. Node3 and Node5 are not selected.

     When adding nodes to an AG that is protected by Cohesity, ensure you register the node with Cohesity and add it to the Protection Group. Otherwise the AG databases will not be protected.

     When removing nodes from an AG that is protected by Cohesity, you can remove it from the Protection Group or keep it in the job; the AG databases remain protected in either case. If the removed node remains in the Protection Group, only its local databases are backed up. If you do not want to back up the removed node, you can safely remove it.

   - Select only this server.

Click **Select**.

Click **Add**.

3. If appropriate, set options for specific objects by clicking (**Edit Server Options**) next to the object. A successful SQL log backup automatically truncates the SQL logs.
4. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

   The selected storage domain stores the snapshots and logs created by this Protection Group.
5. **Backup Method:** Select **File-based**, **Volume-based** or **VDI-based**. File-based, volume-based and VDI-based jobs can backup Filestream databases.
6. **Make Full Backup Copy-only:** Ensure this is toggled off to ensure log backups are taken for Filestream databases.
7. **Support Filestream databases:** Toggle On the **Recommended if this protection group contains Filestream databases** option to protect Filestream databases for file-based backup.

   This option is displayed only when the **File-based** is selected as the **Backup Method**.
8. **User Databases:** Select how to handle AG databases during backup.
9. **System Databases:** Select whether to back up or skip system databases.
10. **Volume Backup:** Optionally select **Back up only the volumes associated with selected databases** to back up only volumes that are associated with selected databases. This option is not used for MS SQL FCI; only volumes with databases are backed up. For MS SQL running on VMs, all volumes are always backed up.
11. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
12. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

13. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
14. **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
15. **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

    For configuration details, see Configure Pre & Post Scripts.
16. **Indexing:** Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.

    For information about customizing indexing, see Customize Indexing.
17. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.

- **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

For more information, see Alerts.

18. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
19. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

> **Note:**
> The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

20. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
21. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Edit an MS SQL Protection Group

1. Select **Data Protection > Protection**.
2. Locate the Protection Group you want to edit.
3. Mouse over the action icon (⋮) next to the job and select one of the following:
    1. **Run Now** - Start a run of this Protection Group. You have the option to back up one or more objects.
        1. **Backup all Objects in the Protection Group** - Select this option to backup all the databases in the protection group. Any databases excluded from the backup will not be included in the protection job run.

2. **Backup only Selected Objects in the Protection Group** - you can manually select the databases to be backed up at the protection job level. However, the selection of the database that is excluded at the backup job level is not supported.

You can run log backups in parallel with existing log backup runs. This helps reduce storage buildup on the primary SQL Server when multiple large backups are in progress.

This feature is supported only for file-based, VDI-based, and Filestream (backed up using VDI-based method) backups.

> **Note:** When you run a log backup for a database backup job, any preceding backup job for the same database may skip the log backup if the new backup job is completed before the previous backup job for the database is started.

2. **Pause Future Runs**
    1. Select **Pause Future Runs** to stop new runs of this Protection Group from executing. However, if the job is currently executing a run, the current run continues to execute and only future runs of this job are paused.
    2. Select **Resume Future Runs** to restore the Protection Group to a running state. The runs resume as specified in the schedule defined in the policy associated with this job.
3. **Edit** - Edit job settings of the protection group and click Protect.
4. **Delete** - Deletes the protection group.

4. Save the job settings by clicking **Protect**.

Configure Microsoft SQL Server T-log Backup

Cohesity enhances the SQL Server T-log backup and restoration performance by splitting the T-log backup into multiple backup files and ingesting them in parallel through the Cohesity nodes.

This feature is applicable for File-based, Volume-based, and VDI-based backups.

Follow the below steps to configure the T-log backup settings for a Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect > Databases > MS SQL Server**.
3. Click **Objects** and select the required objects to be protected.
4. In the **Policy** drop-down list, select a policy that includes Log Backup (Database). The **Log Backup Settings** option will be enabled.

5. Optionally configure the following details:
   - **WITH Clause**: Define the WITH clause that you want to use to customize the backup. For more information, see BACKUP (Transact-SQL). The default WITH clause is "*WITH MAXTRANSFERSIZE = 4194304, BUFFERCOUNT = 16*".

     > **Important:** The WITH clause text box allows you to invoke the SQL Server general options during backups and restores. The text field accepts any SQL Server general backup or restore option like "ENCRYPTION". Note that specifying an option could add to the overall duration of the backup or restore process.

   - **Number of Streams**: Define the number of Cohesity nodes that will be used for the backup process. Using multiple nodes provides more throughput for the backup process. The default number of streams is *1*.

     For details on other settings, see Backup Microsoft SQL Server (VDI-based).
6. Click **Protect**.

> **Note:**
> - Ensure that the agents pre-installed on the SQL VMware source are upgraded to the latest version.
> - If a replication job is configured for the SQL VMware log backup, ensure that the replication cluster is upgraded to the latest version to enable seamless recovery.

Requirements

Ensure the following requirements are met for using SQL log backups:

- The Cohesity cluster's VIPs on **port 11117** must be reachable from the SQL source.
- In a multi-tenant environment, the Hybrid Extender IPs on **port 11117** must be reachable from the SQL source.

> **Note:** If the required port is unreachable, the log backup will fail to start, especially after upgrading from a version lower than 7.2.

Recover Microsoft SQL Server

Cohesity provides the ability to recover the MS SQL databases protected by a Protection Group.

- Recover Microsoft SQL Server Database Over the Original Database.
- Recover Microsoft SQL Server Database as a New Database in the Original Microsoft SQL Server Instance.
- Recover Microsoft SQL Server Database as a New Database in an Alternate Microsoft SQL Server Instance.

Prerequisites

- Before you can restore MS SQL databases, a full backup must exist.
- (*Applies for volume-based and VDI-based backed-up data.*) The cluster must be a member of the same or any trusted AD domain to which the MS SQL database is restored.
- Cohesity cluster VIPs must be reachable (DNS) and accessible (Ping) from the target MS SQL server.
- For recovery, the server and the cluster must be on the same VLAN.

The following topics provide guidance to help you restore MS SQL databases:

About Restoring Microsoft SQL Server Databases

A restore task copies a database from the backup copy and optionally applies transaction logs depending on the options you choose. Cohesity provides the ability to restore individual databases to a specific point in time.

The following options are supported when restoring an individual database:

- Restore to the original MS SQL server and original instance and overwrite the database.
- Restore to the original MS SQL server and original instance and rename the database.
- Restore to the original MS SQL server and an alternate instance as a specified database name.
- Restore to an alternate MS SQL server and instance as a specified database name.

Restoring physical server to physical server only is supported.

MS SQL DB Restore Example

The following example depicts restoring an individual MS SQL database running on a physical server to its original location and overwriting the existing database. The example assumes the following about the database:

- Failure has already occurred.
- Its MS SQL server is registered as MS SQL on the cluster.
- It was backed up by a Cohesity Protection Group that includes transaction logs.

The action steps in the following procedure are represented in the figure below with blue numbered time squares:

.

T1) The Cohesity user must select the following task options:

- Restore to the original SQL server and instance and overwrite the database
- An MS SQL database and snapshot
- Point in time (optional)
- Capture tail logs (optional)

T2) The Cohesity cluster copies the database files from the specified snapshot to the original location.

T3) If applicable, transaction logs are applied to the database for the specified point in time.

T4) The database is then brought online.

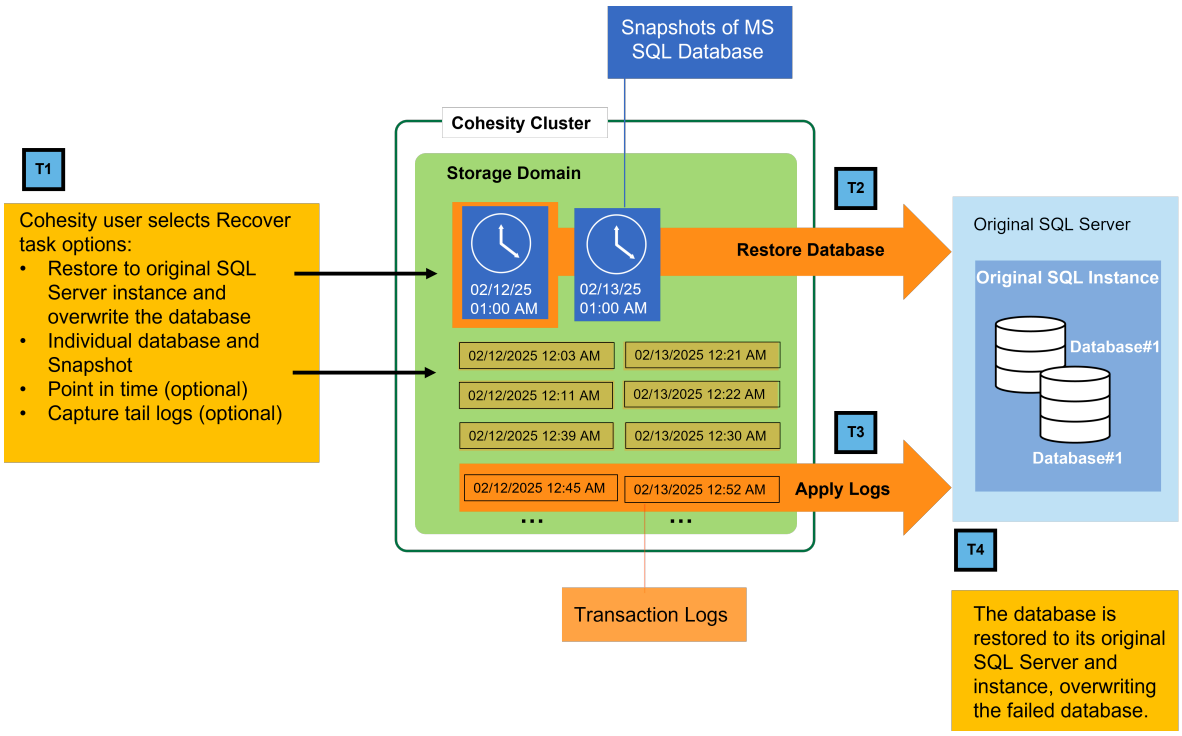

Resume MS SQL Database Recovery After Failure

You can resume a failed MS SQL Database recovery for databases in the following scenarios:

- If the databases are recovered using the WITH NORECOVERY option
- If the previous restore operation failed and the database remains in the restoring state.

This feature improves the RTO for database restore operations and is applicable only for point-in-time (PIT) restore operations.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Recover Microsoft SQL Server System Databases

System databases can be restored in case of system failures. The system databases are restored from the same instance and location where the system databases were backed up. You can also overwrite the existing system database files.

> **Note:**
> Restoring system databases to a different instance is not supported for the VM-based backup method.

The SQL Server system databases are recovered in the following order:

1. Master
2. MSDB
3. Model

Recovery Workflow

You can recover the system databases to the original SQL Server instance.

The following animated image shows how to recover the system databases to the original server instance:

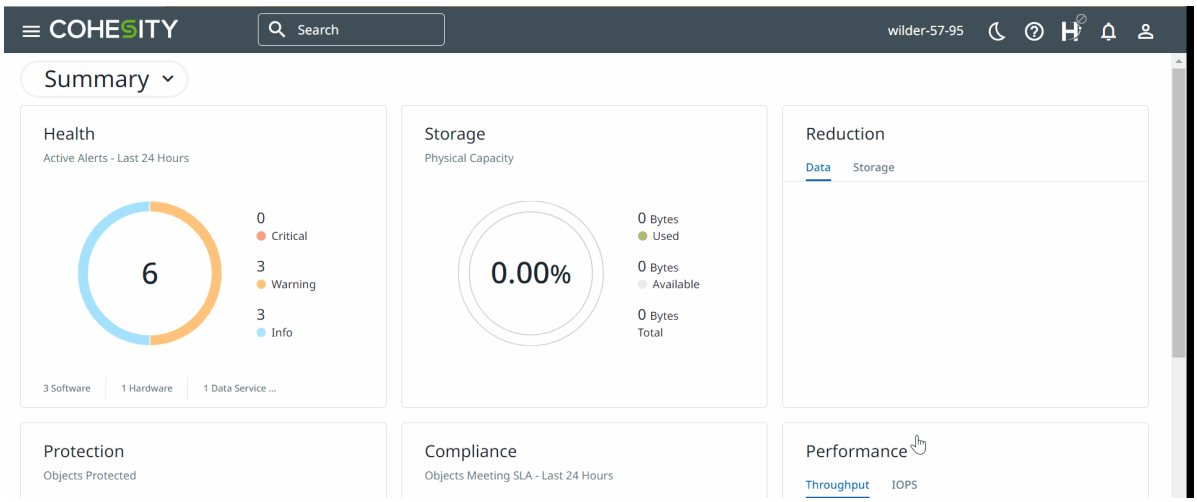

For the detailed steps on recovering system databases, see:

- Recover Microsoft SQL Server Database Over the Original Database

Consideration

- Recovering the system database to an alternate instance is not supported.
- When recovering to the original database, the original database is deleted before the recovery. Therefore, a recovery failure will also lead to the loss of the original database.
- If you are recovering the system databases individually, then the recovery of the system database must be in the following order:
  1. Master
  2. MSDB
  3. Model
- You cannot recover a system database along with a user database.
- If you are recovering more than one system database in a recovery task, then the recovery point for the selected system databases must be the same.

Recover Microsoft SQL Server Database Over the Original Database

Recovering an MS SQL database over the original database overwrites the original database. The recovered database will have the original database name.

> **Note:** The original database is deleted before the recovery. Therefore a recovery failure will also lead to the loss of the original database.

To recover the database to the original database:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Databases** > **MS SQL**.

   The Restore SQL page appears.
3. Search for the databases you want to recover from by entering the characters of the database name, Protection Group name, or by specifying the wildcard character *.

   A list of databases is displayed with details like Cohesity Cluster, Protection Group, Physical Server, and Create Date. The Create Date option indicates the date and timestamp when the database was created.

   > **Note:** Operations like rename may update the creation date at the SQL Server.
4. From the list of search results, select one or more databases you want to recover. You can select multiple MS SQL databases only if they are protected by the same

protection group.

By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the edit icon.

The Edit recovery point for <database_name> page appears.

By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

- To select a recovery point from the Timeline mode:
    a. From the **Choose a date** drop-down, select a date.

       The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.

    b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

    c. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    d. Click **Select Recovery Point** and then go to step 5.
- To select a recovery point from the List mode:
    a. Select the **List** tab from the top-right corner of the Edit recovery point for <database_name> page.

       By default, the page displays the recovery points available for the last 30 days.

    b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

    c. Select the recovery points to be recovered.

    d. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.

6. From the **Targets** option, select **Overwrite Original Database** to recover the database to the original location.

> **Note:** The original database is deleted before the recovery. Therefore a recovery failure will also lead to the loss of the original database.

7. Specify the following **Recovery Options**:

Recovery Options	Description
VDI Restore Settings	Enable this option to specify the WITH clause that you want to use for the restore. **Note:** • This option is available only for restoring the databases that are protected using Backup Microsoft SQL Server (VDI-based). • This option is not available for system database restore.
Keep CDC	Enable this option to restore a protected database that has the change data capture (CDC) enabled. If the protected database is not CDC enabled and the user tries to restore it with Keep CDC then, the database will be restored without CDC. **Note:** • This option is available only if the database is restored with recovery. • This option is not available for system database restore.
WITH Clause for Full/Incremental Restore	Specify the SQL Server WITH clause to customize the full or incremental recovery operation. This option is available only when recovering databases protected using VDI-based backup.
WITH Clause for Log Restore	Specify the SQL Server WITH clause to customize the transaction log restore operation. This option is available only when recovering databases protected using VDI-based backup and applies specifically to log backup recoveries.
Replay Entire Last Log	Enable this option to replay the entire last log backup for the specified point in time without using the STOPAT option. When disabled, only logs up to the selected point-in-time are replayed.

Recovery Options	Description
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Grou**p drop-down.
Task Name	Change the default name of the recovery task.
Capture Tail Logs	Enable this option to capture records that have not yet been protected. They are captured to ensure all transactions are protected before restoring the database. **Note:** This option is not available for system database restore.

8. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Microsoft SQL Server VDI-based Recovery Recommendations

The Virtual Device Interface (VDI) is a Microsoft interface that allows the Cohesity agent to execute SQL Server Native backup and restore commands. This backup type produces a standard SQL Server Native backup of type (.BAK), transaction logs, (.TRN), and differentials, (.DIFF). These backup files are stored on a Cohesity View and take advantage of View compression, deduplication, encryption, replication, and archiving features.

The aim of the SQL DBA is to have a combination of backups so that the database can be restored to any point in time. A good combination of backups consists of having a full backup, differential (in Cohesity, incremental) backups, and log backups. For example, all SQL database restores must begin with a full database backup, secondly, a differential can then be applied to the full, and finally, log backups can be applied in sequence to complete the database restore.

When the backups are applied during the database restore process, you are sequentially adding the captured changes to the database: FULL+DIFF+Log1+Log2+Log3 = Restored Database.

All Microsoft VDI backups, their differentials, and their logs are dependent on a full backup to perform a database restore. Microsoft requires that in order to restore a SQL database, you must start with a full backup, then its transaction logs can be applied. This means your backup retention policy must keep a full backup along with its log backups in order to successfully restore a database.

Simply put, a SQL VDI-based database restore requires a full backup to seed the database, then differential and/or log backups are applied to the specified point in time.

We recommend retaining two sets of full backups with their differential and log backups.

> **Note:** A good backup plan always includes a combination of full, differential, and log backups.

Recover Microsoft SQL Server Volume-based Backups from Archive

Follow one of the workflows described in the following sections to restore a volume-based MS SQL backups from the archive. Using standard MS SQL restore workflows are not supported when restoring from the archive. This is because the backup is volume-based which cannot extract the MDF, NDF, and LDF files from the archive workflow. The volumes that contain the database files must be retrieved from the archive.

> **Note:** When restoring from the archive for file-based MS SQL backups, you can use standard MS SQL restore from archive workflows. For file-based backup, the MDF, NDF, and LDF files can be directly extracted from the archive. For VDI-based backup, the .bak files can be directly extracted from the archive.

For file-based backup, the MDF, NDF, and LDF files can be directly extracted from the archive.

You can manually recover database files from a volume-based backup and then use them to manually create a new database in SQL Server.

> **Note:** This method is supported only for full and incremental backups and not log backups.

Files from volume-based backup can be recovered using the following workflows.

Restore to Original Cohesity Cluster (from where MS SQL was Archived)

Recover/Clone an Entire Snapshot (Protection Group run) or VM from the Archive to the Target vCenter

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** > **Virtual Machines** > **VMs**.
3. Select a Job or VM (crash consistent) to recover.
4. Select the Snapshot and recovery location.

Instant Mount a Volume of a Virtual or Physical Server from the Archive to the Target Server

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** > **Virtual Machines** > **Instant Volume Mount**.
3. Select the Server for instant volume mount.
4. Select the recover point, volumes to mount and the mount target.
5. Copy the database files to the desired MS SQL server.
6. Manually attach the database to the desired MS SQL instance.

Download Individual Files

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** > **Virtual Machines** > **Files or Folders**.
3. Select the database files to download.
4. Select the recover point and any additional recovery options.
5. Copy the database files to the desired MS SQL server.
6. Manually attach the database to the desired MS SQL instance.

Restore to Different Cohesity Cluster (CloudRetrieve from Remote Cluster)

1. In the Cohesity Dashboard, select **Data Protection > CloudRetrieve**.
2. Click **New Search**.

   Specify the archive to retrieve from, the Protection Group name and the Cluster where the Job originated.
3. After the search completes, download the Jobs to the Cluster. This may take some time to complete.
4. After the archived data is downloaded (an Inactive Job is created), you can use the standard restore/clone MS SQL database workflows. Do either of the following:
   - Click **Data Protection > CloudRetrieve > Downloaded Jobs**. Mouse over the actions menu and select an action.
   - Click **Data Protection > Recoveries Recover > MSSQL**.

Recover Microsoft SQL Server Database as a New Database in the Original Microsoft SQL Server Instance

To recover the database as a new database in the original MS SQL server instance:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Databases** > **MS SQL**.

   The Restore SQL page appears.

3. Search for the databases you want to recover from by entering characters of the database name, Protection Group name, or by specifying the wildcard character *. A list of databases is displayed with details like Cohesity Cluster, Protection Group, Physical Server, and Create Date. The Create Date option indicates the date and timestamp when the database was created.

   > **Note:** Operations like rename may update the creation date at the SQL Server.

4. From the list of search results, select one or more databases you want to recover. You can select multiple MS SQL databases only if they are protected by the same protection group.

   > **Note:** The option to select multiple databases in a recovery task is an Early Access feature. Contact Cohesity Support for enabling this feature.

   By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the edit icon.

   The Edit recovery point for <database_name> page appears.

   By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

   - To select a recovery point from the Timeline mode:
     a. From the **Choose a date** drop-down, select a date.

        The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.

     b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

     c. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     d. Click **Select Recovery Point** and then go to step 5.
   - To select a recovery point from the List mode:

a. Select the **List** tab from the top-right corner of the Edit recovery point for <database_name> page.

   By default, the page displays the recovery points available for the last 30 days.

b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

c. Select the recovery points to be recovered.

d. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. From the **Targets** option, select **Recover as a new Database** to recover the database as a new database.
7. Enable the **Restore to Original SQL Server Instance** option to restore the database in the original MS SQL Server instance.
8. In the **Data Files** field, specify the path to an existing directory. The newly created database files will reside in this path.

   > **Note:** If the database file names to be restored have a conflict with the existing database files of the database to which files are restored, then Cohesity resolves the conflict by appending a number to the file names.

9. In the **Log Files** field, provide the full path location to which the log files can be restored.
10. (Optional) Click **Add File Path Rules** to add custom rules for the data files.

    You can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

    Accept the file patterns and examples:

    - Wildcard character used for file matching such as *
      - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
      - *m:*.ndf* - Every ndf file on drive m:
      - *c:*1.ndf* - Every ndf file on drive c: that contains 1.ndf

- Regex pattern
  - *.*\.ndf* - All files with the ndf extension
  - *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf
- Full path to a source database file

Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

Click **Add File Path Rule** to add multiple custom rules.

11. Specify the following **Recovery Options**:

Recovery Options	Description
Rename	Select one of the following options: • **Bulk Rename**: Select this option if you want to rename the restored databases in bulk. You can specify a string in the **Suffix** field. The restored databases will be renamed by adding this suffix string with the original database names. • **Rename Individual Objects**: Select this option if you want to individually rename the restored databases. You can specify the new name of the databases in the **New Name** field.
WITH RECOVERY	Disable the option to perform a restore WITH NO RECOVERY (the database will be offline when restored). By default, this option is enabled and the database is restored with recovery (the database will be online when restored).
VDI Restore Settings	Enable this option to specify the WITH clause that you want to use for the restore. **Note:** This option is available only for restoring the databases that are protected using Backup Microsoft SQL Server (VDI-based).
Keep CDC	Enable this option to restore a protected database that has the change data capture (CDC) enabled. If the protected database is not CDC enabled and the user tries to restore it with Keep CDC then, the database will be restored without CDC. **Note:** This option is available only if the database is restored with recovery.
Overwrite Alternate Database	Enable this option to recover the database by overwriting the existing database.

Recovery Options	Description
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Grou**p drop-down.
Task Name	Change the default name of the recovery task.

12. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Microsoft SQL Server Database as a New Database in an Alternate Microsoft SQL Server Instance

To recover the database as a new database in an alternate MS SQL server instance:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Databases** > **MS SQL**.

   The Restore SQL page appears.
3. Search for the databases you want to recover from by entering characters of the database name, Protection Group name, or by specifying the wildcard character *. A list of databases is displayed with details like Cohesity Cluster, Protection Group, Physical Server, and Create Date. The Create Date option indicates the date and timestamp when the database was created.

   > **Note:** Operations like rename may update the creation date at the SQL Server.

4. From the list of search results, select one or more databases you want to recover. You can select multiple MS SQL databases only if they are protected by the same protection group.

   By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the edit icon.

   The Edit recovery point for <database_name> page appears.

   By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

- To select a recovery point from the Timeline mode:
    a. From the **Choose a date** drop-down, select a date.

       The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.

    b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

    c. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    d. Click **Select Recovery Point** and then go to step 5.
- To select a recovery point from the List mode:
    a. Select the **List** tab from the top-right corner of the Edit recovery point for <database_name> page.

       By default, the page displays the recovery points available for the last 30 days.

    b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

    c. Select the recovery points to be recovered.

    d. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. From the **Targets** option, select **Recover as a new Database** to recover the database as a new database.
7. Disable the **Restore to Original SQL Server Instance** option to restore the database to an alternate SQL Server instance.
8. From the **MS SQL Instance** drop-down, select an alternate SQL server instance.
9. In the **Data Files** field, specify the path to an existing directory. The newly created database files will reside in this path.

> **Note:**
> - If the database file names to be restored have a conflict with the existing database files of the database to which files are restored, then Cohesity resolves the conflict by appending a number to the file names.

10. In the **Log Files** field, provide the full path location to which the log files can be restored.
11. (Optional) Click **Add File Path Rules** add custom rules for the data files.

    You can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

    This is only used for the destination path in a custom rule and not the source. File path rules are case sensitive. Partial file name or extension (not only directory path) are required for a match to be identified.

    > **Note:** For VDI based backup, use two backslashes after the drive. For example, C:\\path.

    Accept the file patterns and examples:

    - Wildcard character used for file matching such as *
      - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
      - *m:*.ndf* - Every ndf file on drive m:
      - *c:*1.ndf* - Every ndf file on drive c: with a suffix of 1.ndf
    - Regex pattern
      - *.*\.ndf* - All files with the ndf extension
      - *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf
    - Full path to a source database file

    Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

    Click **Add File Path Rule** to add multiple custom rules.
12. Specify the following **Recovery Options**:

Recovery Options	Description
Rename	Optional. Enable the Rename option and select one of the following options to rename the restored databases: • **Bulk Rename**: Select this option if you want to rename the restored databases in bulk. You can specify a string in the **Suffix** field. The restored databases will be renamed by adding this suffix string with the original database names. • **Rename Individual Objects**: Select this option if you want to individually rename the restored databases. You can specify the new name of the databases in the **New Name** field.
WITH RECOVERY	Disable the option to perform a restore WITH NO RECOVERY (the database will be offline when restored). By default, this option is enabled and the database is restored with recovery (the database will be online when restored). **Note:** When restoring a database to an alternate MS SQL server instance with no recovery, a staging directory is created under the database files and logs file directories. To overcome the filename conflict with existing files in the database, the recovered files are saved in this staging directory.
VDI Restore Settings	Enable this option to specify the WITH clause that you want to use for the restore. **Note:** • This option is available only for restoring the databases that are protected using Backup Microsoft SQL Server (VDI-based).
Keep CDC	Enable this option to restore a protected database that has the change data capture (CDC) enabled. If the protected database is not CDC enabled and the user tries to restore it with Keep CDC then, the database will be restored without CDC. **Note:** • This option is available only if the database is restored with recovery. • This option is not available for system database restore.
WITH Clause for Full/Incremental Restore	Specify the SQL Server WITH clause to customize the full or incremental recovery operation. This option is available only when recovering databases protected using VDI-based backup.

Recovery Options	Description
WITH Clause for Log Restore	Specify the SQL Server WITH clause to customize the transaction log restore operation. This option is available only when recovering databases protected using VDI-based backup and applies specifically to log backup recoveries.
Replay Entire Last Log	Enable this option to replay the entire last log backup for the specified point in time without using the STOPAT option. When disabled, only logs up to the selected point-in-time are replayed.
Overwrite Alternate Database	Enable this option to recover the database by overwriting the existing database.
Cluster Interface	By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Grou**p drop-down.
Task Name	Change the default name of the recovery task.

13. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the recoveries page.

Recover Microsoft SQL Server Database as Flat Files

Cohesity allows you to recover the SQL database backup files to the specified location on the target host, rather than as a database. This provides flexibility for self-service recovery where direct access to the backup files is required.

The backup files are recovered to a subdirectory in the specified destination path, with the following file name format: `<OriginalHostName>_<OriginalInstanceName>_<OriginalDatabaseName>`

For example, the backup files are recovered to **C:\RestoreDir\SQLHost01_MSSQLSERVER_AdventureWorks**

where,

- **C:\RestoreDir** is the destination path,
- **SQLHost01** is the host,
- **MSSQLSERVER** is the instance,
- **AdventureWorks** is the database.

> **Note:** This is applicable only for the SQL Servers backed up using file-based and VDI-based backup and not applicable to volume-based backup.

When you recover databases as flat files, the following backup files are recovered based on the backup method and recovery point:

Backup Type	Recovery Point	Recovered Files
File-Based	Incremental Backup	Database files - .mdf, .ldf., .ndf, VSS metadata
	Point In Time	• Database files - .mdf, .ldf., ndf, VSS metadata • Log backup files - .trn
VDI	Incremental Backup	• Full backup file - .bak • Differential backup file - .bak
	Point In Time	• Full backup files - .bak • Differential backup files - .bak • Log backup files - .trn

To recover SQL Server databases as flat files:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Databases** > **Microsoft SQL**.

   The **New Recovery** page is displayed.
3. Search for the databases you want to recover from by entering characters of the database name, Protection Group name, or by specifying the wildcard character *. A list of databases is displayed with details like Cohesity Cluster, Protection Group, Physical Server, and Create Date. The Create Date option indicates the date and timestamp when the database was created.

   > **Note:** Operations like rename may update the creation date at the SQL Server.
4. From the list of search results, select one or more databases you want to recover. You can select multiple MS SQL databases only if they are protected by the same protection group.

   By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the edit icon.

   The Edit recovery point for <database_name> page appears.

By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

- To select a recovery point from the Timeline mode:
    a. From the **Choose a date** drop-down, select a date.

       The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.

    b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

    c. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    d. Click **Select Recovery Point** and then go to step 5.
- To select a recovery point from the List mode:
    a. Select the **List** tab from the top-right corner of the Edit recovery point for <database_name> page.

       By default, the page displays the recovery points available for the last 30 days.

    b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

    c. Select the recovery points to be recovered.

    d. By default, the data is recovered from the local cluster. You can also select a recovery point archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

    e. Click **Select Recovery Point** and then go to step 5.

5. Click **Next: Recover Options**.
6. From the **Targets** option, select Recover as a Flat File to recover the database files to the target location, rather than as a database.
7. Specify the following details:

Backup Type	Description
Physical Server	Windows server where the agent is installed and registered to the cluster.
Database File Path	Destination path in the physical server where the flat files will be recovered.
Overwrite Existing Files	Select to overwrite the existing files in the destination path. This is a destructive action and cannot be reverted.
Cluster Interface	By default, the **Auto Select** option is enabled, and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
Task Name	Change the default name of the recovery task.

8. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the **Recoveries** page.

Recover Microsoft SQL Server Filestream Databases

Before you can restore MS SQL Filestream databases, a snapshot must exist.

Restore Database Overview

1. Select **Dashboard** and from the drop-down, select **MS SQL**.
2. Click **Databases**.
3. Select **Recover** from the database's actions menu.
4. Set restore options.

You can also start a restore task by navigating to **Data Protection > Recoveries** > **Recover** > **Databases** > **MS SQL** and selecting a database to restore.

Set Restore Options

1. **Task Name:** Optionally change the default name for the task.
2. **Restore to Original SQL Server Instance:** Toggle on if you want to restore the database to its original instance.

   Overwriting a Filestream database is not supported and prevents the restore from completing.

   a. Select a **SQL Host** and **SQL Instance** if restoring to an alternate MS SQL instance.

> **Note:** Local MS SQL instances that are running within an MS SQL FCI do not display. You can, however, manually enter the local SQL instance name to restore to.

b. Provide the full path location to restore all database files.

   You can optionally restore data files and log files to different locations. Click **Add Alternative Log File Locations and Custom Rules** to provide locations for data and log files and any custom rules. For data files, you can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

   This is only used for the destination path in a custom rule and not the source.

   Accept the file patterns and examples:

   - Wildcard character used for file matching such as *
     - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
     - *m:*.ndf* - Every ndf file on drive m:
     - *c:*1.ndf* - Every ndf file on drive c: with a suffix of 1.ndf
   - Regex pattern
     - *.*\.ndf* - All files with the ndf extension
     - *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf
   - Full path to a source database file

   Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

   Click **Add Another Custom Rule** to add multiple custom rules.

   > **Note:** Ensure to follow the standard regular expression rules. Specify a dot '.' before asterisk '*' to match all characters, for example, L:.*.ndf.

c. Provide a **Database Name** for the restored database.

3. By default, the latest snapshot is selected. You can optionally choose to restore from a different snapshot by clicking Edit ( ).

   You can also provide a specific time to restore to if transaction logs are available.

Click a date on the calendar to select a snapshot from that date. The default point selected will be the latest log transaction or the latest full or incremental, whichever is the latest. The slider also displays available points in time and displays snapshots as circles. Manually select a point on the slider if you want to restore to a specific point in time. Selecting an invalid time from the slider automatically changes to the closest available snapshot. If your manual selection is within 10 minutes of a snapshot, your selection automatically snaps to it. You can also select a specific point in time by entering it in the time fields.

After selecting the desired recover point, click **Recover**. You can choose to clone the selected recover point by clicking the arrow adjacent to recover and selecting **Clone**.

In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

4. By default, an MS SQL restore WITH RECOVERY is performed. You can optionally toggle this off to perform a restore WITH NORECOVERY.
5. **Auto Select:** When enabled, the recovery task automatically selects the correct VLAN.

   > **Note:** If using VLANs and the cluster is joined to an AD domain—One of the machine accounts must be a DNS hostname for a virtual IP (VIP) of a VLAN configured on the cluster in order to recover to servers within that VLAN.

6. **Capture tail logs:** This option is available only when overwriting a database in its original location. This option is unavailable if the database's Protection Group has been paused or deleted.

   Tail logs capture records that have not yet been backed up. They are captured to ensure all transactions are backed up before restoring the database.
7. Click **Recover** to start the task.

SQL Maintenance Mode

You can enable maintenance mode for SQL Server instances and databases. Enabling maintenance mode skips any scheduled activities or manually initiated protection runs.

To place a SQL Server instance or database in maintenance mode:

1. Navigate to **Data Protection > Sources**.
2. Click on the action menu (⋮) of the source and select **Maintenance**.
3. You can initiate the maintenance mode immediately or later.

   To initiate the maintenance mode immediately:

1. Select **Start Now**. The **Start Maintenance Now** dialog displays the list of Protection Groups available in the source.
2. Review the list of Protection Groups and enter **Yes** to confirm.
3. Optionally, you can enter a note.
4. Click **Confirm**.

To schedule the maintenance mode:

1. Select **Schedule**.
2. On the **Schedule Maintenance** dialog, enter the start date and time and the end date and time.
3. Optionally, you can enter a note.
4. Click **Save**.

The SQL Server instance or database will be set in maintenance mode accordingly.

AG in Maintenance Mode

When maintenance mode is enabled for a SQL instance hosting the AG database, the replica hosting the instance is not considered for the AG backup plan. This behavior is common to both legacy (non-first-class citizen) and first-class citizen features.

For legacy AG backup, only the replica for which maintenance mode is enabled will be skipped. You must enable maintenance mode on all the replicas to skip the entire backup.

Clone Microsoft SQL Server Data

The following topics provide guidance to help you clone MS SQL databases:

About Cloning Microsoft SQL Server Databases

A clone task instantiates a backup of an MS SQL database so it can run directly off the Cohesity. While cloning an MS SQL database, you can choose a destination server that is registered with Cohesity as MS SQL. Cloning physical servers to physical servers only is supported.

MS SQL DB Clone Example

The following example depicts cloning a database from a physical SQL server to a SQL server instance running on a different physical server. The action steps in the following procedure are represented in the figure below with blue numbered time squares: T1 .

The process consists of the following actions:

T1) The Cohesity user must select the following task options:

- Clone to a specified SQL server and SQL instance
- An MS SQL database and snapshot

T2) The cloned database files are then brought online as a new database.

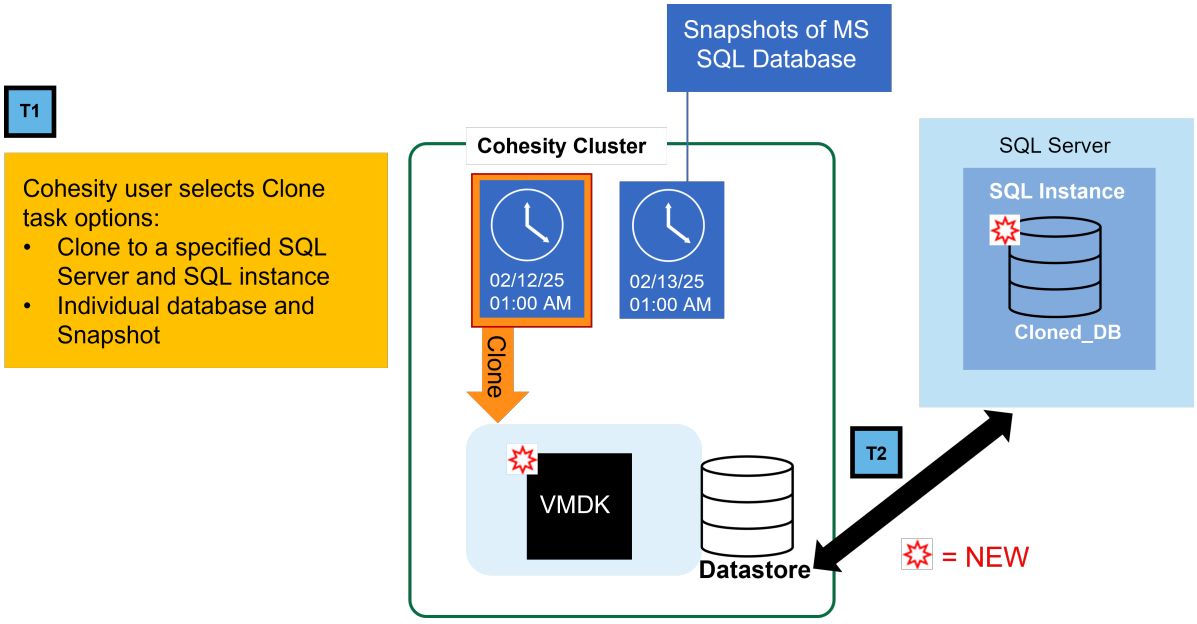

Clone Microsoft SQL Server Databases

Cohesity now provides the ability to clone an MS SQL database to a Standalone SQL Server, FCI, and AG. A major advantage is that the cloned MS SQL database files reside on the Storage Domain and require no additional storage from the MS SQL host. Multiple clones from a single source of MS SQL Database backup require little or no additional storage space in the Cohesity cluster.

> **Note:** Cloning an MS SQL database is supported with Volume-based and file-based backup methods.

Clone an MS SQL Database

1. Navigate to **Test & Dev** > **Clone** > **Databases**.
2. Search and select the MS SQL database. You can optionally perform a search by specifying the wildcard character *. You can also optionally narrow the search results by specifying filter criteria.
3. Provide the following information listed in the table, and then click **Clone**.

Options	Description
Task Name	(*Optional*) Change the default name for the Clone task.
Objects Being Cloned	By default, the database you selected for cloning is displayed. You can click the change icon to select a different database.
Clone Point	By default, the latest snapshot is cloned. Optionally click the Change icon to select a different snapshot to clone. The selected Restore Point is listed
SQL Host	Select the host to which you want to clone the database. By default, you can select only the physical servers as the host. If you want to select MS SQL FCI as the SQL host, Contact Cohesity Support. Cloning the MS SQL database to an MS SQL FCI is supported only if the following criteria are met: • Microsoft SQL Server version must be 2012 and above, supported on SMB file server. • The MS SQL database to be cloned is backed up with a file-based protection group.
SQL Instance	Select the instance for the cloned database.
Rename Database?	Enable this option if you want to rename the cloned database created by this task.
Cohesity Network interface	Select one of the following options: • **Auto Select**: By default, the **Auto Select** option is enabled and the clone task automatically selects the correct VLAN. • **Interface Group**: select a configured interface group from the **Interface Group** drop-down.

Cancel a Database Clone

You can cancel a running clone task for an individual database. Canceling the task may leave the database in an unusable state. You can delete the database manually and then start another clone task as needed.

1. Navigate to **Test & Dev**.
2. Click a Task Name link in the table to display the details page.
3. Click the **Cancel** button. (The Cancel button is available only if the task is running.)
4. Respond to the confirmation prompt.

Tear Down a Database Clone

To delete an MS SQL database clone from the Cohesity Storage Domain, perform the following steps:

1. Navigate to **Test & Dev**.
2. Click the cloned task or hover your mouse over the actions menu located next to the cloned task you want to delete.
3. Click **Tear Down Clone** > **Yes, Tear Down**.

The MS SQL database clone will be deleted from the Cohesity cluster Storage Domain.

> **Note:** The teardown of an MS SQL clone database removes it from the SQL server; the user does not have to do this beforehand. However, in the rare instance you created other additional databases with files on the mounted drive that the cloned database is on, those additional databases should be deleted first through the SQL Server Management Studio.

Migrate Microsoft SQL Server

The following topics provide guidance to help you migrate MS SQL databases:

About Migrating Microsoft SQL Server Databases

You can use MS SQL database migration for the following use cases:

- Back up and migrate a database from a single pane of glass.

  Leverage database backup to achieve database migration. No more shipping logs across multiple sites or multiple data centers.

  Upgrade the MS SQL Server version while minimizing application downtime and maintaining the ability to restore to any point in time.
- Staging a database on the target server for faster recovery with Cohesity's database migration.

  Stage a database on the target SQL server and perform only incremental updates at the file level to improve RPO.

Database migration for file-based MS SQL backups consists of the following actions, which must be initiated manually:

1. Perform the initial copy of database files to the target MS SQL server.
2. Sync forward to the latest incremental/full snapshot available on the cluster. You can trigger multiple syncs.

For each sync, Cohesity performs incremental file updates to database files on the target MS SQL server. Syncs ensure files are up to date. The files are not attached to the database until you complete the finalized workflow. This ensures that the database is up to date when finalization occurs.

3. Finalize the migration. Finalizing is available only if the target database is synced to the latest snapshot.

   This brings the database online. After the database is online and the finalize step succeeds, the recovery task is completed successfully. You cannot perform syncs after this point because the migration workflow is complete.

Migrate Microsoft SQL Server Databases

Before you migrate an MS SQL database, ensure that a file-based snapshot exists on the cluster.

Migrate MS SQL Database

This section lists the steps to migrate an MS SQL database.

1. Select **Dashboard** and from the drop-down, select **MS SQL**.
2. Select **Migrate** from the database's actions menu.

   You can also start a restore task by navigating to **Data Protection > Recoveries> Recover > Migrate Database (MS SQL only)**.
3. Search for an MS SQL server or database. Enter characters of the name and a list of items displays. You can optionally specify the wildcard character *. For more information, see Use Wildcard * in Search for Recovery. You can also optionally narrow the search results by specifying filter criteria, for example, you can filter the search results by a specific Protection Group. Click the **Add Filters** icon, specify the filter and click **Add**.
4. Select a database and set the migration options for the MS SQL database you want to migrate.
   1. **Task Name:** Optionally change the default name for the task.
   2. **Objects Being Recovered**: This displays the database to be migrated.
   3. **Sync Options**: Specify the sync options for the database to be migrated.

      > **Note:** Transaction log backup is not synced. Only the incremental and full backup is synced.

      - **Auto Sync Databases**: This option automatically keeps the database in sync with the latest available snapshot until you finalize the database and bring it online.

- **Manually Sync Databases**: Once the initial database migration is completed, you can sync the migrated database files with newer snapshots available until you finalize the database and bring it online.

4. **Recover Point:** This displays the selected recover point for the migration.
5. **Settings:** Specify the **SQL Host** and **SQL Instance** to migrate the database to.

   Provide the full path location to restore all database files.

   You can optionally restore data files and log files to different locations. Click **Add Alternative Log File Locations and Custom Rules** to provide locations for data and log files and any custom rules. For data files, you can provide a full path or pattern for the original files and a location to restore to. When providing a location, use the standard path format. For example, use C:\path for Windows instead of the extended path format \\?\c:\path.

   This is only used for the destination path in a custom rule and not the source.

   Accept the file patterns and examples:

   - Wildcard character used for file matching such as *
     - **.ndf* - All ndf files. Useful if you want every ndf file to go to the same folder
     - *m:*.ndf* - Every ndf file on drive m:
     - *c:*1.ndf* - Every ndf file on drive c: with a suffix of 1.ndf
   - Regex pattern
     - *.*\.ndf* - All files with the ndf extension
     - *^.*foo1.*\.ndf$* - All files with the word foo1 and the extension ndf
   - Full path to a source database file

   Database files that have non-default file extensions or that do not match a specified filename or pattern are restored to the data file restore location.

   Click **Add Another Custom Rule** to add multiple custom rules.

   > **Note:** Ensure to follow the standard regular expression rules. Specify a dot '.' before asterisk '*' to match all characters, for example, L:.*.ndf.

5. By default, an MS SQL restore WITH RECOVERY is performed. You can optionally toggle this off to perform a restore WITH NO RECOVERY.
6. Click **Start Migration**.

   The initial file copy task starts. This is the first step in the migration process.

7. After the file copy task is completed, the options available for the next migration step are displayed.
    - **Finalize** - Available if the target database is synchronized to the latest snapshot.
    - **Sync** - Available if there is a snapshot that is more recent than the target database.

    If you select the **Auto Sync Databases** option in Step 4, the **Auto Sync to latest** toggle is enabled by default. You can manually disable this option if required.

> **Note:** Database migration is not supported for AG databases.

Cancel Migration

You can cancel an in-progress migration task. Canceling the migration prevents syncing or finalizing the database and may leave it in an unusable state. Delete the database files manually and then start another migration task as needed.

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Find the migration task in the list and click the task name to display its detail page.
3. Click the **Cancel** button located at the top right of the page.
4. Respond to the confirmation prompt.

Oracle Database

Cohesity's software-defined platform provides a simple, fast, and cost-effective backup, recovery, and data management solution for growing Oracle database environments. Cohesity integrates with Oracle Recovery Manager (RMAN) to deliver an integrated and centralized approach to manage Oracle backup jobs. It provides a single location to register and schedule all RMAN scripts across the Oracle environment. See the following topics for details:

Our platform provides several integrations for protecting the Oracle Databases:

- **Cohesity Oracle Adapter**: Cohesity Oracle Adapter integrates with Oracle Recovery Manager (RMAN) to provide application-consistent backup and recovery for Oracle databases. The adapter provides a unified user interface allowing users to manage and provision all the Oracle database protection services, supporting Oracle single instance, RAC or Oracle Multitenant databases. For details, see Cohesity Oracle Adapter.
- **Cohesity Remote Adapter**: Cohesity Oracle Remote Adapter is a built-in component of Cohesity that offers flexible management of scripts running on remote hosts. This allows the Cohesity cluster to manage data protection processes and schedules as well as provide consolidated logs of all related activities. User-supplied

Oracle RMAN scripts are hosted on the Oracle server with one extra step to register the scripts with the Cohesity cluster. Once registered, the RMAN scripts are remotely triggered and monitored for completion and activity logged on the cluster. For details, see Oracle Database Remote Adapter.

- **RMAN NFS Target**: Cohesity can be presented as an NFS volume to be used as inline-deduplication Oracle databases backup target. Oracle RMAN can use the Cohesity Data Cloud as a deduplication target and automatically leverage Cohesity's QoS policy for optimal data tiering as well as global deduplication and compression to reduce secondary storage consumption. For details, see RMAN NFS Target for Oracle Database.

  For detailed information on Cohesity's integration with RMAN, see Oracle Data Protection Deployment Using RMAN NFS Target.

- **Cohesity Oracle SBT Plug-in**: The Cohesity Oracle SBT Plug-in allows you to use Cohesity Data Cloud as target storage for backing up Oracle databases. For details, see Cohesity Oracle SBT Plug-in.

Supported Features

The following table shows how each protection method supports different features.

Features	Cohesity Oracle Adapter	Remote Adapter	RMAN NFS Target	Cohesity Oracle SBT Plug-in
Database backup	Supported	Supported	Supported	Supported
SLA-driven policies	Supported	Cohesity and customer driven	Customer-driven	Customer-driven
Database backup automation	Cohesity driven	Customer-driven	Customer-driven	Customer-driven
Oracle backup types	Image copy	All RMAN supported backup types	All RMAN supported backup types	Backup-set Only
Oracle database restore	Cohesity driven	Customer-driven	Customer-driven	Customer-driven
Clone	Cohesity driven	Customer-driven	Customer-driven	Not Supported

Features	Cohesity Oracle Adapter	Remote Adapter	RMAN NFS Target	Cohesity Oracle SBT Plug-in
Dedupe	Cohesity driven (Windows only)	Source side dedupe library (Linux, AIX, and Windows)	Not Applicable	Source side dedupe library (Linux, AIX, and Windows)
Point-in-time recovery	Cohesity driven	Customer-driven	Customer-driven	Customer-driven
View expose	Cohesity driven	Customer-driven	Customer-driven	Customer-driven
RMAN channel configuration for performance	RMAN channels=MIN(2* # of Cohesity nodes, 2* # of target server CPU) The default value is 8. To configure more RMAN channels, contact Cohesity Support. You can also edit the Cohesity nodes and CPUs, if required.	Customer-driven	Customer-driven	Customer-driven

Cohesity Oracle Adapter

The Cohesity Oracle Adapter natively integrates with Oracle Recovery Manager (RMAN) to provide application-consistent backup and recovery for Oracle single-instance deployments, Oracle Exadata Databases Real Application Clusters (RAC), and Veritas Active-Passive.

Based on the policy settings, the Cohesity Oracle Adapter executes an RMAN backup command to initiate a backup of the Oracle Database. The adapter also automatically determines the optimal number of backup channels, based on the compute resources on the Oracle Database server and the number of nodes in the Cohesity cluster. The necessary Cohesity Views are automatically generated (that is, an NFS mount point on the client and an NFS View on Cohesity), where the backups are going to be written. Communication between Cohesity and the Oracle Databases is managed by the Cohesity Oracle Adapter as described in the section below.

Core components

The following lists the core components of the Cohesity Oracle adapter.

- Physical or virtual Oracle instance.
- Cohesity Data Cloud.
- Cohesity Oracle Adapter.
- Archive or tiering target at your data center or cloud provider of choice.

As illustrated in the below figure, you can integrate the Cohesity Data Cloud with a new or existing Oracle infrastructure by installing the Cohesity Oracle adapter. Once the agent is installed, you can create protection policies, jobs, and replication schedules for the Oracle database. Once the database is protected, it can be archived, sent to tape, and or replicated to another Cohesity cluster running in the cloud or on-premises as shown in the below image.

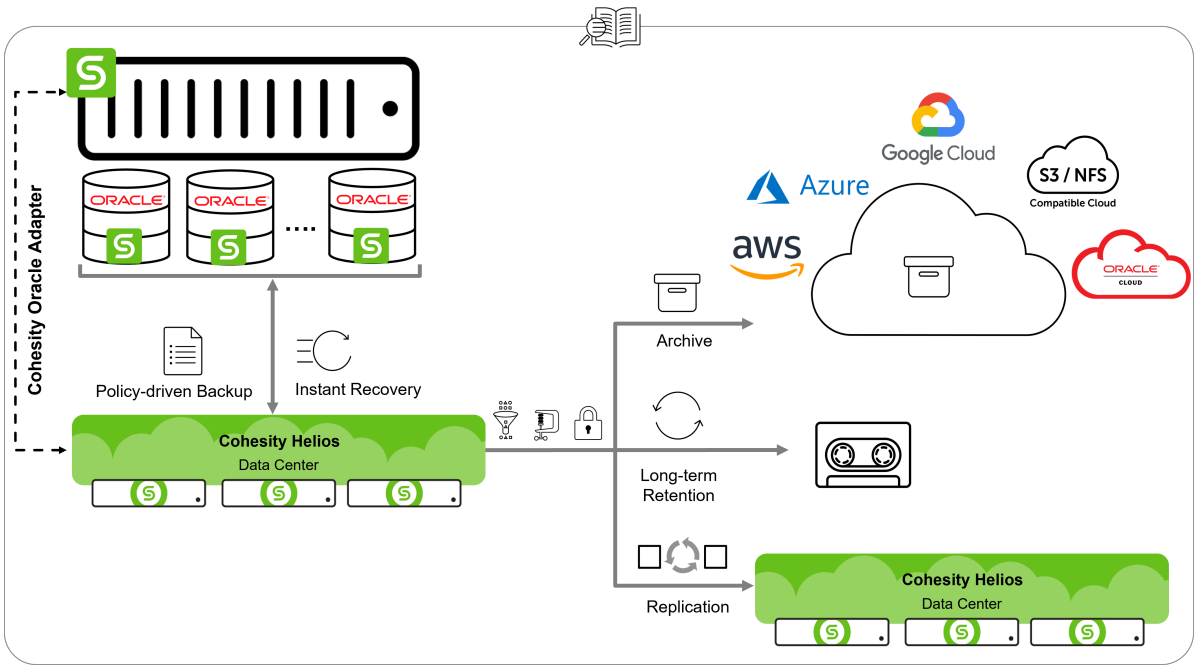

Cohesity Oracle Adapter Benefits

The following table highlights the benefits of using an Oracle Adapter for your Oracle database backup solution:

DB Backup and Storage Challenges	Solutions
Backup agent degrades database performance	The Cohesity Oracle Adapter is a lightweight agent on the target host and does not degrade the application performance.

DB Backup and Storage Challenges	Solutions
Lack of point-in-time (PIT) Support for recovery and cloning	The Cohesity Oracle Adapter offers an intuitive user interface for PIT Recovery and Clone workflows.
Expensive usage of primary storage	The Cohesity Oracle Adapter supports auto-tiering and global dedupe or compression thereby offering flexibility and cost efficiencies.
Slow data restores	The Cohesity Oracle Adapter offers faster restores as every database backup is fully hydrated using Cohesity snapshots and Oracle incremental merge capabilities. The traditional backup workflow requires repeated time consuming full backup schedules and Cohesity does not require this because customers only need to do one full backup, then the rest is 'incremental forever'.
Protection against compromised data	Cohesity performs an incremental merge after each incremental backup and validates the backup image to achieve a 'fully hydrated' image copy of the backup. The backup image is stored internally on the Cohesity cluster as an immutable snapshot. The merge will fail if the incremental backup is corrupted.
Higher Total cost of ownership	The distributed web-scale platform is designed for secondary data. Performance and capacity can be scaled linearly simply by adding nodes that eliminates the need for forklift upgrades. You can use the Cohesity Oracle adapter like a pay-as-you-go architecture.
Complex data-protection environments	The Cohesity Oracle Adapter simplifies the data protection by replacing multiple data protection silos (target storage, media servers, master servers, and cloud gateways) with a single hyper-converged solution for backup, recovery, replication, cloud tiering, archiving, and target storage.
Inability to leverage public cloud for long-term retention	The Cohesity Oracle adapter provides native cloud integration with: • Amazon Web Services (AWS) • Microsoft Azure • Oracle Cloud • Google Cloud Platform (GCP) It supports long-term retention and tiering of cold data. If the snapshots are sent to the cloud, they can be downloaded on another Cohesity cluster and restores can be performed from this cluster.
RMAN scripts maintenance overhead	No need to maintain any RMAN scripts.

DB Backup and Storage Challenges	Solutions
No single dashboard to manage Oracle Database data protection	The Cohesity Oracle adapter offers a single dashboard to manage Oracle Database data protection.
Lack of Oracle point-in-time (PIT) restore from cloud archived snapshots	Cohesity supports the PIT restore and clones from cloud archived snapshots.
Mount Point Management	The Cohesity Oracle self-healing manager periodically checks the availability of all the mount points created by the Cohesity Oracle adapter. If a mount point is not available within a specific duration, the self-healing manager performs these actions: • Treats the mount point as a non-responsive mount point. • Unmount this mount point. • Clean the RMAN sessions associated with the non-responsive mount point.
Managing Standby Databases	Cohesity supports taking backup from standby databases and restore or clone from those backups to protect Oracle data from failures, disasters, human error, and data corruption while providing high availability for mission-critical applications • Physical StandBy Support We support physical standby databases for all Oracle Configurations. • Logical StandBy Support We support Oracle Logical Standby with single instance config (non-CDB, non TDE) for Linux only. Exceptions for Logical StandBy All restrictions for Logical Standby from Oracle will apply to Cohesity. For example, If a logical standby database is not a block-for-block copy of the primary database, you cannot use a logical standby database to back up the primary database. Supported Cohesity Versions We will support this for 6.5.1x, 6.6.x, and future versions only.

Planning and Preparation for Oracle Database

The following topics provide guidance that helps you get ready to implement the Oracle Adapter:

Supported Features for Oracle Database

The following table lists the supported features for the Cohesity Oracle Adapter on different OS platforms:

Supported Oracle Features	Linux	AIX	Solaris	Windows
Standalone DB Backup (Single-instance Oracle DB)	Yes	Yes	Yes	Yes
RAC Backup (Multi-node Oracle DB)	Yes	Yes	No	No
Restore to Original Server (a.k.a. Overwrite Restore)	Yes	Yes	Yes	Yes
Restore to Alternate Server	Yes	Yes	Yes	Yes
Cloning	Yes	Yes	Yes	No
Instant Restore	Yes	No	No	No
Data Guard	Yes	Yes	No	No
CDB/PDB Support	Yes	No	No	Yes*
TDE Support	Yes	No	No	No
Archive Log Restore	Yes	No	No	No
Customizable PFile	Yes	Yes	Yes	No

* Oracle Windows adapter supports only the CDB backup and overwrite restore of the original CDB.

Prerequisites for Oracle Database

Make sure the following prerequisites are met before using Oracle Adapter:

- **DBIDs**: All the Oracle Databases that are protected using the Cohesity cluster should have a unique DBID on the Oracle source where the databases reside.
- **SCAN VIP** (*Applies only to Linux and AIX*): To register Oracle RAC, the user must use SCAN-NAME or SCAN-VIP, Otherwise, registration, backup, and recovery of the Oracle databases will fail.
- **Installation Agent Type**: You must use the same agent installation type (Script or RPM) on all the Oracle RAC nodes.
- **Authentication**: For database-level authentication, the tns-listener service name must be the same as the database name.

- **Archive Log Mode**: Archive Log mode must be enabled for the database opened in Read-Write mode, otherwise the backup task will fail.
- **Version**: The recovery source and target database must be the same version of the Oracle Database. For example, snapshots of the 11g Oracle Database cannot be recovered to the 12c Oracle Database.
- **Full Backup**: Cohesity recommends taking a full database backup after any activity that changes the control file (such as adding or removing a data file, restoring to the original location with overwrite, or changing the log file location).
- **Oracle Single Instance Deployment**:
  - For Oracle single instance database, the database must be entered into the `/etc/oratab` file for Linux and AIX and `/var/opt/oracle/oratab` file for Solaris. Otherwise, Cohesity will not be able to discover this database.
  - For Oracle single instance database running on Windows OS, `OracleService${ORACLE_SID}` service must be running.
- **VCS Deployment** (*Applies only to Linux and AIX*): For Veritas Cluster Server (VCS) deployments, the Oracle database environment should be configured in the `/etc/oratab` on all the VCS nodes.
- **Register Oracle RAC**: To register an Oracle RAC or a RAC node as a physical server, "host" command must be executed on each of the nodes of that RAC.

  The "noac" flag for NFS mounts must be set for Oracle RAC versions 11gR2 and 12cR1. If not set, Oracle RAC backup fails.
- **Oracle RAC Database** (*Applies only to Linux and AIX*)
  - Cohesity Linux Agent must be installed on all the Oracle servers that are part of an Oracle RAC.
  - The `GRID_HOME` must be mentioned in `/etc/oratab` or `/etc/oraInst.loc` file .
- **DB Authentication** If you choose DB authentication, then all the DBs on the system should have the same username and password. Special characters is not supported in the DB authentication password.
- **SBT Library** (*Applies only to Windows*): On the Windows Oracle Server, install the SBT library in the same directory where the Windows Agent is installed. By default, the agent is installed in the **c:\Program Files\Cohesity\oracle** directory. Contact Cohesity Support to obtain the SBT library installer.
- For Oracle alternate restore to the ASM disk group, check the following prerequisite before starting restores on the Cohesity cluster:
  - **ASM Compatibility (COMPATIBLE.RDBMS)**

    The disk group attributes that determine compatibility are COMPATIBLE.ASM, COMPATIBLE.RDBMS. and COMPATIBLE.ADVM. The COMPATIBLE.ASM and COMPATIBLE.RDBMS attribute settings determine the minimum Oracle

Database software version numbers that a system can use for Oracle ASM and the database instance types respectively.

The value for the COMPATIBLE.RDBMS disk group attribute determines the minimum COMPATIBLE database initialization parameter setting for any database instance that is allowed to use the disk group. Before advancing the COMPATIBLE.RDBMS attribute, ensures that the values for the COMPATIBLE initialization parameter for all of the databases that access the disk group are set to at least the value of the new setting for COMPATIBLE.RDBMS.

For example, if the COMPATIBLE initialization parameter of the databases is set to either 12.2, then COMPATIBLE.RDBMS can be set to any value between 10.1 and 12.2 inclusively. If the COMPATIBLE initialization parameter of the databases is set to 18.0, then COMPATIBLE.RDBMS can be set to any value between 10.1 and 18.0 inclusively.

The following query can be used to check existing Compatibility for ASM Disk group

select name, DATABASE_COMPATIBILITY, COMPATIBILITY from v$asm_ diskgroup;

Examples of disk group compatibility attribute settings:

COMPATIBLE.ASM	COMPATIBLE.RDBMS	Oracle ASM Instance Version	COMPATIBLE Setting for RDBMS Instance
12.1	12.1	>= 12.1	>= 12.1
12.2	12.1	>= 12.2	>= 12.1
19.0	19.0	>=19.0	>=19.0

- Oracle Software owner and ASM Management owner

  If the Oracle Grid / Oracle Restart owner is different from the Oracle RDBMS owner, make sure that the Oracle Grid OS group added as a secondary group for the Oracle RDBMS owner.

  For example,

  Oracle Grid/ Oracle Restart --- oraasm - asmdba, asmadmin (OS user and group)

  Oracle RDBMS --- oracle - oinstall, dba (OS user and group)

  In the above case add asmdba as a secondary group to oracle user.

Credentials and Privileges for Oracle Database

Once you register your physical servers with Cohesity as Oracle servers, Cohesity will discover your Oracle Database automatically. For Cohesity to successfully discover your Oracle Databases, the user running the Cohesity agent must have the appropriate credentials and privileges.

The Cohesity Linux Agent can be installed as a ROOT user or with a separate OS user (also known as the 'OS Service Account user'). The AIX Java Agent and Solaris can be installed only as a ROOT user.

When connecting to the Oracle databases, Cohesity can use either the Oracle OS Authentication or Oracle DB Authentication method. These two types of Oracle authentication are possible for both cases where the Agent is run with the ROOT user or a separate OS Service Account user. Most Oracle Adapter operations can be done using either OS or DB authentication. Some specific operations would require either OS or DB authentication. For more information, see Oracle Authentication Method Option table.

Running Agent with ROOT

You can install Linux Agent or the AIX Java Agent, or Solaris Java Agent to run with the ROOT user:

1. The agent runs every command using the ROOT user, except for Oracle commands and utilities like RMAN or SQLPLUS. To run Oracle commands and utilities, the agent will 'su' to the user who is the owner of the Oracle binary in the current Oracle Home. If an Oracle operation is run against a source database that has DB Authentication configured (where the user has previously configured DB credential for this Oracle source database in the Cohesity Oracle Adapter), DB Authentication will be used to run Oracle commands and utilities. Otherwise, OS Authentication via the Oracle binary owner will be used.
2. To use Oracle RAC Multinode Multichannel Backup, you must configure the corresponding source RAC database with DB credentials to allow the Agent to use the required DB Authentication method to run this feature.
3. When you install the Cohesity agent to run with the ROOT user, there is no need to configure additional SUDOERS privileges.

Running Agent with OS Service Account User

The Linux Agent can be installed to run with a specific OS Service Account user that meets the following requirements:

1. The following environmental variables should be set to match the Oracle user when the Cohesity agent is running as non-root and not oracle:
    - LD_LIBRARY_PATH
    - ORACLE_HOME

2. This Service Account user must be granted additional Sudoers privileges to allow the Agent to run operations like NFS mounting Cohesity views. For the complete list of the required Sudoer commands, see Sudoers Permissions for Service Account User .
3. This Service Account user must be part of the OS group that enables this user to either to have the Oracle SYSDBA or SYSBACKUP privileges. For example, on a typical Linux operating system, the OS group 'DBA' (dba) often enables the corresponding Oracle SYSDBA privilege. Another customer-defined OS Group, example: 'BACKUPDBA' (backupdba), can enable the corresponding SYSBACKUP privilege. This gives the Service Account user the ability to use the OS Authentication method to connect to Oracle databases with the intended privilege level.
4. If the Agent is running on an Oracle host that is part of a RAC cluster that uses ASM, or a non-RAC Oracle host that uses ASM that is available to that Oracle instance, the Service Account user should also be part of the OSDBA group of your ASM environment. This gives the Service Account user the ability to use the OS Authentication method to indirectly connect to the ASM instance to do things like reading ASM files during backup or restoring files into ASM during database recovery.
5. To use Oracle RAC Multinode Multichannel Backup, you must configure the corresponding source database with DB credentials to allow the Agent to use the required DB Authentication method to run this feature.

Oracle Authentication Method Option

Whether you are installing the Cohesity agent to run with the ROOT user or a separate OS Service Account user, the following table lists the Oracle Authentication method requirements for all Oracle Adapter operations.

Operation	Authentication Required	Notes
Non-RAC Backup (Single-instance Oracle DB)	OS Authentication or DB Authentication	Non-RAC single node require OS Authentication or DB authentication.
RAC Backup (Multi-node Oracle DB)	OS Authentication or DB Authentication	RAC multinode and multichannel require OS / DB authentication.
Restore to Original Server (a.k.a. Overwrite Restore)	OS Authentication or DB Authentication	Restoring data to the same server overwrites the original database.

Operation	Authentication Required	Notes
Restore to Alternate Server	OS Authentication	DB Recovery or Restore onto a different server with Oracle binaries that has the same version as the source database backup it is restoring. target oracle server must have enough free space to store the newly created database files.
Cloning	OS Authentication	DB clones onto a target server with Oracle binaries that have the same version as the source database backup it is cloning from.

Windows Agent Authentication

Windows Adapter Operation	Authentication Required
Non-RAC Backup (Single-instance Oracle DB)	OS Authentication or DB Authentication
Restore to Alternate Server	OS Authentication
Restore to Original Server (a.k.a. Overwrite Restore)	OS Authentication or DB Authentication

Sudoers Permissions for Service Account User

The following tables list the sudoers permissions for the Oracle Linux agent for OS Service Account User.

Operating System	Oracle Deployment Mode		
	Stand Alone	**RAC**	**Active Passive**
Linux	Commands for both Oracle Adapter & Linux physical agent • mount • umount • cp • chown • chmod • mkdir • rm • tee • hostname • stat • timeout • ls • rsync • touch • cat • mv Additional commands only for Linux physical agent • blkid • lsof • losetup • dmsetup • lvs • vgs • lvcreate	<COHESITY_AGENT_INSTALLATION_BASE_PATH> /cohesityagent/software/crux/bin/oracle_scripts/srvctl_from_oracle_home.sh where, COHESITY_AGENT_INSTALLATION_PATH = indicates where the cohesityagent directory got created after the Cohesity agent installation. For each unique ORACLE_HOME present in the **/etc/oratab** there would be a following entry in the /etc/sudoers. ORACLE_HOME/bin/srvctl	• <PATH TO> lltconfig • <PATH TO> hagrp • <PATH TO> hares

Operating System	Oracle Deployment Mode		
	Stand Alone	RAC	Active Passive
	• lvremove • lvchange • kill		

You can also allowlist a variant of command in sudoers. For instance, for the timeout command, you may only want to allowlist certain options such as:

- `/usr/bin/timeout --kill-after 5 --signal 1 300s`
- `/usr/bin/timeout --kill-after 5 --signal 1 300s /sbin/lvcreate *`
- `/usr/bin/timeout --kill-after 5 --signal 1 300s /sbin/lvextend *`
- `/usr/bin/timeout --kill-after 5 --signal 1 300s /sbin/lvremove *`
- The location of these commands binaries is OS vendor (RHEL, CentOS) and OS version dependent. For instance, with OEL 7.6, `/usr/bin/mount` should be allowlisted for the mount command.
- The specific options for some of these commands can also be OS vendor, OS version and shell dependent. For instance, in the case of BASH shell, the timeout option can be `/usr/bin/timeout --kill-after 5 --signal 1 300s`.
- In the case of RHEL 6.10, the timeout option `preserve-status` is not supported.

Communication Ports for Oracle Database

The following tables list the ports used for external communication when using the Cohesity Oracle adapter.

Ports used for Oracle adapter for Linux and AIX

Port	Source	Target	Direction (from node)	Network Protocol	Usage Notes	Type of Traffic
50051	Cohesity cluster	Oracle Host	Bidirectional	TCP	Linux physical server backups. Used for Cohesity Agent and Agentless data transfer to the Cluster. Needs to be open on all Nodes in the Cluster and any backup target using the Cohesity Agent, such as physical servers and Linux servers running the Cohesity installed Agent.	Backup and Recovery
111	Oracle Host	Cohesity cluster	Bidirectional	TCP	Oracle NFS based backups and restores.	Backup and Recovery
2049	Oracle Host	VIPs and VLAN IPs on Cohesity cluster	Bidirectional	TCP	Oracle NFS based backups and restores.	Backup and Recovery

Ports used for Oracle adapter for Windows

Port	Source	Target	Direction (from node)	Network Protocol	Usage Notes	Type of Traffic
50051	Cohesity cluster	Oracle Host	Bidirectional	TCP	Linux physical server backups. Used for Cohesity Agent and Agentless data transfer to the Cluster. Needs to be open on all Nodes in the Cluster and any backup target using the Cohesity Agent, such as physical servers and Linux servers running the Cohesity installed Agent.	Backup and Recovery
11117	Oracle Host	Cohesity cluster	Bidirectional	TCP	Oracle SBT based backups and restores.	Backup and Recovery

Oracle Precheck Utility

When you perform an Oracle adapter workflow such as a backup, alternate restore, or clone operation, the precheck utility automatically checks permissions, disk space, ArchiveLog mode, and so on, based on the parameters you provided.

The Precheck utility supports both C++ and AIX Java agents for Oracle.

If there are discrepancies with the parameters or the environment (server) where the workflow will run, the precheck utility displays the cause of the discrepancies on the Cohesity Dashboard. You can then resolve the discrepancies and try the workflow again.

Oracle adapter precheck is a utility that performs environmental checks and reporting for the Oracle Adapter. This precheck ensures that all key criteria are met for running any Oracle Adapter workflow.

> **Note:** This feature is supported only for Oracle databases with Linux agent.

How Precheck Utility Works

Before running the Oracle workflows like backup, alternate restore, or clone, the precheck utility performs sanity checks (permission, disk space, ArchiveLog mode, and so on) based on the parameters you provided.

If there is any discrepancy in the parameters you provide or with the environment (server) where the workflow is going to get executed, then the precheck operation fails, resulting in the bailout of Oracle workflow. The precheck utility then displays the cause for the discrepancy on the Cohesity Dashboard to notify the user about the restrictions and prerequisites for the environment and the Oracle adapter workflow.

Performing Prechecks

The precheck utility performs sanity checks for the following Oracle Adapter workflows:

- Discovery
- Backup
- Alternate restore
- Clone

Common Checks

- Checks if the Oracle binaries, SQL Plus, RMAN, oraadm, and srvctl have required access permissions for the Oracle agent for work.
- Checks if the Oracle shared libraries (.so) libsql.so has required access permission for the Oracle agent to work.
- Checks if ORACLE_HOME has required RD/WR permission for the agent to run the backup, restore, and clone workflows.

Discovery

The utility performs the following checks while discovering the DBs on the Oracle Host:

- Checks the password in the case of DB based authentication in Oracle Registration.
- Checks if the Cohesityagent user is part of the DBA group.
- Oracle Home and Oracle Base binaries and libraries are accessible by Agent user and DBA group.
- Checks the access for ORACLE_HOME/rman ORACLE_BASE
- Checks if the ORATAB file is present on the system (generally/etc/oratab).
- For RAC, the precheck utility checks if there is +ASM entry the /etc/oratab.
- Checks if the /etc/oratab entries are valid.

Backup

The utility performs the following checks before running the backup workflow:

- If Archivelog is enabled: If Archivelog is not enabled, you cannot schedule a backup job.
- If Block Change Tracking (BCT) is enabled: If BCT is not enabled, a warning message is displayed on the Cohesity Dashboard and the backup process will experience high latency.
- SYSDBA privilege: Check if the Cohesityagent user (user running Cohesity agent) is part of the DBA group.
- Adequate permission checks on ORACLE_HOME, and Oracle binaries and libs.
- MultiNode-MultiChannel check: In the case of MNMC backup, the precheck utility tries to connect to all the nodes selected for a backup.
- Checks if DB is in READ-WRITE mode for backup.

Alternate Restore

The precheck utility performs the following checks before running the alternate restore workflow:

- Checks if the Cohesityagent user is part of the DBA group.
- Adequate permission checks on ORACLE_HOME, ORACLE_BASE, and Oracle binaries and libs.
- Oracle DBSID: Checks the database whether the same SID exists. If the SID exists, the restore will fail.
- Checks for the datafile space for Alternate Restore.
- Checks for permission and existence of BCT file (if BCT file location is given from Cohesity Dashboard).
- Checks for permission and existence of DIAG_DEST (if DIAG_DEST is provided from Cohesity Dashboard).

- Checks for permission and existence of AUDIT_LOG_DEST (if AUDIT_LOG_DEST is provided from Cohesity Dashboard).
- Checks if the spfile<NEW_DB_SID> or init<NEW_DB_SID>.ora exists in the new ORACLE_HOME when performing the Restore to alternate DB. This check is to safeguard against the existence of DB with the same name in Shutdown state or a deleted DB which is not cleared properly.
- Oracle version: Restore will fail if the source and destination database versions are not the same.
- Oracle DBNAME: Checks whether the database with the same name exists. If the name exists, restore will fail.
- Existence and permission of the path given at the time of alternate restore.
- Disk space availability while performing the alternate restore.
- Existence of Oracle Home, Oracle Base, and the required permissions for these directories.
- Checks if the total memory is more than the SGA target.
- Checks if the Cohesity agent is installed with a different user than the Owning user of Oracle binaries (SQL Plus, RMAN, NID).
- Check the database instance name in /etc/oratab. If it exists, the alternate restore will fail.

Clone

The utility performs the following checks before running the clone workflow:

- Checks if the Cohesityagent user is part of the DBA group.
- Adequate permission checks on ORACLE_HOME, ORACLE_BASE, and Oracle binaries and libs.
- Checks if the Cohesity view is mounted with the correct option (nolock or lock).
- Checks if the spfile<NEW_DB_SID> or init<NEW_DB_SID>.ora exists in the new ORACLE_HOME when performing the Restore to alternate DB. This check is to safeguard against the existence of a DB with the same name in a Shutdown state or a deleted DB which is not cleared properly.
- Oracle compatibility versions: The compatibility is performed by checking the version of SQL Plus in the given ORACLE_HOME_DIR.
- Checks if the total memory is more than the SGA target
- Oracle DBNAME: Checks whether the Database with the same name exists.
- Checks if the Cohesity agent is installed with a different user than the Owning user of Oracle binaries (SQL Plus, RMAN, NID).

Best Practices for Oracle Database

Cohesity recommends you follow these best practices when you plan to deploy the Cohesity Oracle adapter:

- If you use choose to DB authentication, then all the DBs on the system should have the same username and password.
- Optional: Create a database user for Cohesity Oracle backup and restore workflow.
- Both the Oracle and Cohesity user agents should have permission to write to the `adump` and `diag` directories, the control file, and the database restores locations.
- Archivelog mode needs to be enabled to do a backup of an Oracle database opened in Read-Write mode.
- Block Change Tracking (BCT) should be enabled to improve the incremental backup performance of the Oracle adapter.

The following best practices apply only for Oracle adapter for Linux and AIX:

- Assign sudoers to the user running the Cohesity agent.
- Make the Cohesity agent user part of the Oracle `dba` group.
- In the case of a RAC or if you are using an Oracle ASM environment, then the Cohesity agent user should be part of the OSDBA group for ASM (also known as ASMDBA group).

> **Note:** In Oracle Adapter workflows, Oracle databases running on Linux and Windows must be protected in separate protection groups. Do not include Linux- and Windows-based Oracle databases in the same protection group.

Considerations for Oracle Database

Review the following considerations:

- If the original database is unavailable (disaster recovery scenario) and no pfile/spfile is present, you should start an instance with the same db_name as the original database in NOMOUNT mode before proceeding with an overwrite recovery of the original database from Cohesity.
- Cohesity does not support the backup of Oracle databases on a physical standby (with open_mode set to 'READ ONLY WITH APPLY') in a Data Guard deployed on Windows.
- Cohesity supports restoring a database from Data Guard Standby to standalone oracle database.
- Our incremental forever backup workflow uses Oracle imagecopy. You have to consider that, if the backup is done with SECTION SIZE breaking up the imagecopy, restore of the imagecopy cannot be done in SECTION SIZE as Oracle RMAN restore

does not support that. This results in longer restore time as one RMAN channel will have to restore a complete datafile, instead of the possibility of breaking down this datafile into different 'sections' using SECTION SIZE and have multiple RMAN channels to restore the same datafile in parallel.

- Cohesity does not currently support the backup of the Oracle RAC database and Oracle VCS on the Windows server.
- Parallel log backup is allowed only for policy-driven backup and not using the 'run now' option.
- If you are using the Oracle 21c version and the `Oracle Home` location is marked as read-only then the alternate restore and clone jobs on the Cohesity cluster may fail.
- Backup of the read-only database is only supported on the Standby role.
- A replication task of a log backup run of an Oracle Protection Group shows a high amount of "logical data transferred" as it replicates the last successful full/incremental backup data.
- Currently, a single instance is only supported with clones regardless of whether the source is Oracle single instance or Oracle RAC.
- Cohesity recommends performing a full backup before a log backup for the Oracle databases.

Windows

- For alternate recovery with Oracle on Windows adapter, You must use the "`oradim`" tool to create the target database service. You should run this service before the recovery, or the recoveries will fail.
- Before the 6.8.1_u2 release, for Oracle on Windows, if a full backup snapshot is deleted, then all subsequent incremental snapshots are not recoverable until the next full backup is available. From Cohesity release 6.8.1_u2 onwards, you can recover an incremental snapshot without the previous full backup on the Cohesity cluster.

AIX

- After cloning the Oracle database running on the AIX server, you must update `/etc/filesystem` manually to make the mount point used for clone persistent across reboot.
- Instant recovery is not supported for Oracle on AIX.

Install Agents for Oracle Database

This section lists the steps for installing the Oracle adapter agent, registering, and creating protection groups for the Cohesity Oracle adapter.

- Install the Cohesity Oracle Adapter Agent.
- Register the Oracle servers (physical Or virtual server).
- Register the Oracle Database server and authenticate (OS/DB authentication).
- Create a Cohesity Oracle Database Protection Group.
- Set the QoS Policy for the Protection Group.

The following table lists the operating systems and agent installation type when installing the agent on the Oracle server.

Operating System (64 Bit)	Agent Installation Type
RHEL OEL CentOS	Script installer RPM
SuSE Linux	SuSE RPM
AIX (IBM AIX)	BFF Installer
Windows	EXE Installer
Solaris	Solaris Package

For more information on installing the Oracle adapter agent, see:

> **Note:** You must use the same agent installation type (Script or RPM) on all the Oracle RAC nodes

- Installing the Linux Agent
- Installing the AIX Java Agent
- Install the Windows Agent for Oracle Database

Install the Linux Agent for Oracle Database

You can install the Linux agent to run with the ROOT user or a separate OS service account user. For complete details of Oracle Agent installation from different download packaging, please refer to Install and Manage the Agent on Linux Servers.

One of the popular installation packages offered is the Script Installer. Following is a summary of installing the Agent specifically using the Script Installer packaging.

> **Note:** You must use the same agent installation type (Script or RPM) on all the Oracle RAC nodes.

Download the Linux Agent

To download the Linux agent, perform the following steps:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent** near the top-right of the page.
3. Select the **Linux Agent** and choose **Script Installer** as the packaging options
4. Click the download arrow to download the Linux Agent.

Install the Linux Agent to run with the ROOT User

> **Important:** Cohesity recommends installing the Linux Agent as a ROOT user.

For more information on credentials and privileges for ROOT user, see Credentials and Privileges.

For detailed steps on installing the Linux Agent as a ROOT user, Install and Manage the Agent on Linux Servers.

> **Note:** In the case of Oracle RAC deployment, if the Cohesity agent is installed using RPM on one of the nodes, then all the nodes should be installed with an RPM-based installer, and; it is the same for the script installer. You cannot install one of the RAC nodes with RPM and the other with a script installer.

Install the Linux Agent to run with an OS Service Account User

To install the Linux agent on the Oracle server:

1. Run the executable file as a sudo user using the following command syntax:

   ```
   sudo /<path_to_installer_file> -- --install
   ```

   By running the above command, the installer creates the user group, `cohesityagent` and installs the agent.

   If you are running the executable file using the following command options, then the agent will be installed on behalf of the specified user and group.

- **-S** is the Oracle OS user.
- **-G** is the group where oracle_user is a member.
- **-c** prevents the user from being created if it already exists. In case the Customer is installing with -S <existing user> -G <existing Group> then -c should be 0.

For example,

```
./cohesity_agent_6.0_x64_installer -- --install -S oracle -G oinstall -
c 0
```

2. Grant sudo permission to the Service Account user who will run the agent.

   The Cohesity agent installer grants sudo permission for the following commands:

   ```
   cp, ls, rm, mount, umount, chown, chmod, mkdir, tee, hostname, stat,
   timeout, rsync, blkid, dmsetup, losetup, lsof, lvchange, lvcreate,
   lvremove, lvs, vgs, vgcreate, grep, getcap, getent, service, systemctl,
   initctl, systemd
   ```

   Example of the sudoers file addition (actual paths may change):

   ```
   NOPASSWD:/usr/bin/cp, /bin/cp, /usr/bin/ls, /bin/ls, /usr/bin/rm,
   /bin/rm, /usr/bin/mount, /bin/mount, /usr/bin/umount, /bin/umount,
   /usr/bin/chown, /bin/chown, /usr/bin/chmod, /bin/chmod, /usr/bin/mkdir,
   /bin/mkdir, /usr/bin/tee, /usr/bin/hostname, /bin/hostname,
   /usr/bin/stat, /usr/bin/timeout, /usr/bin/rsync, /usr/sbin/blkid,
   /sbin/blkid, /usr/sbin/dmsetup, /sbin/dmsetup, /usr/sbin/losetup,
   /sbin/losetup, /usr/sbin/lsof, /usr/sbin/lvchange, /sbin/lvchange,
   /usr/sbin/lvcreate, /sbin/lvcreate, /usr/sbin/lvremove, /sbin/lvremove,
   /usr/sbin/lvs, /sbin/lvs, /usr/sbin/vgs, /sbin/vgs, /usr/sbin/vgcreate,
   /sbin/vgcreate
   ```

3. This user must be part of the OSDBA or OSBACKUPDBA group. For more information, see Credentials and Privileges for Oracle Database. Notice that this Service Account user can be a newly created OS user or the 'oracle' user itself.
4. The Agent starts automatically after the installation, as well as on subsequent Oracle host reboot.
5. At the end of the installation, the commands used to start, stop or get agent status will be displayed for future reference.

Install the AIX Java Agent for Oracle Database

You can install the AIX Java agent to run with the ROOT user.

Before You Start

Before you install the AIX Java agent on the server, make sure the following software is installed on the server:

- The latest Bash package. For more information, see AIX Toolbox for Linux Applications.
- Java version, JRE 1.8.0 AIX ppc-64 bit for the agent.

> **Note:** Before installing the AIX Java agent, Cohesity recommends logging in to SQL Plus and RMAN using appropriate credentials (OS Auth or User Based). If the login is successful, proceed with the installation. If not, obtain the required user credentials and verify that you can log in.

Download the AIX Java Agent

To download the AIX Java Agent, perform the following steps:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent** near the top-right of the page.
3. Select the **AIX(java agent)**.
4. Click the download arrow to download the AIX Java Agent.

Install the Cohesity Java Agent to run with the Root User

For detailed steps on installing the AIX Agent as a ROOT user, Install and Manage the Cohesity Java Agent on AIX Servers.

Install the Windows Agent for Oracle Database

> **Note:** Cohesity does not support protection of Oracle RAC Cluster and VCS Cluster running on Windows.

To download and install Windows Agent, perform the following steps:

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agents** at the top right of the page.
3. Download the Windows Agent on the Oracle server.

   > **Note:**
   > The SBT plug-in required for Oracle on Windows Adapter will be installed with the Windows agent.

4. As an administrator with local system privileges on that server, run the executable and complete the installation wizard. Ensure that you *do not* select the **Add-ons**

components, **Volume CBT (Changed Block Tracker)** and **File System CBT (Changed Block Tracker)** to be installed.

By default, the agent is installed in the **c:\Program Files\Cohesity** directory. Ensure to install the SBT library in the <Cohesity agent installation location>/oracle directory where the Windows Agent is installed.

Install the Solaris Agent for Oracle Database

You can install the Solaris agent to run with the ROOT user.

Before You Start

Before you install the Solaris agent on the server, make sure the following software is installed on the server:

- Oracle's latest bash package running on Oracle Solaris 11 or greater on SPARC v9 (64-bit).
- Java version, JRE 8 for Oracle Solaris SPARC (64 bit) for the agent.

> **Note:** Before installing the Solaris agent, Cohesity recommends logging in to SQL Plus and RMAN using appropriate credentials (OS Auth or User Based). If the login is successful, proceed with the installation. If not, obtain the required user credentials and verify that you can log in.

Download the Solaris Agent

To download the Solaris Agent, perform the following steps:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent** near the top-right of the page.
3. Select the **Solaris**.
4. Click the download arrow to download the Solaris Agent.

Install the Cohesity Solaris to run with the Root User

To install the Solaris agent on the Oracle server:

1. Log in to the Solaris server as a root user and run the following command to install the package:

   ```
   sudo /usr/sbin/pkgadd -d <cohesity_agent_file_path>
   ```

   For example:

   ```
   sudo /usr/sbin/pkgadd -d cohesity_agent_solaris_6.8.1
   ```

2. The Cohesity Solaris Agent starts automatically after the installation. Now you can register the server with the Cohesity cluster. Reboot is not required. The Cohesity Solaris Agent starts automatically after subsequent Oracle host reboots.

Uninstall the Cohesity Solaris to run with the Root User

To uninstall the Solaris agent on the Oracle server:

Log in to the Solaris server as a root user and run the following command to uninstall the package:

```
sudo /usr/sbin/pkgrm cohesity-agent
```

Register Oracle Server

Registering an Oracle server and its databases makes them eligible to be backed up on the Cohesity cluster.

Register Oracle Server

Registering an Oracle server and its databases makes them eligible for backup on the Cohesity cluster.

- Register Oracle Standalone Source
- Register Oracle Cluster

Register Oracle Standalone Source

> **Note:** To register the Oracle Data Guard standalone source, follow the same steps listed below.

To register an Oracle standalone source, follow the steps below:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** > **Databases** > **Oracle Source**.

   > **Note:** A physical server refers to a bare metal server or a VM running the supported versions of the OS and Oracle Database with the Cohesity agent installed.

3. (*Applies only in multi-tenant environment with Data Source Connection enabled*.) From the **Data Source Connection** drop-down, select a Data Source Connection, and then click **Continue**.
4. In the **Select Source Type** drop-down, select **Oracle Standalone Source**.

5. **Enter Host Address**, for an Oracle Standalone server, specify the Hostname, Server IP, or FQDN of the Oracle server.
6. **Authentication Type**, specify the Authentication Type to register the Oracle source.

   When connecting to the Oracle databases, Cohesity can use either the Oracle OS Authentication or Oracle DB Authentication method. If you use DB authentication, provide the database username and password.
7. Click **Register**.
8. Select one of the following options from the **Cohesity Network Interface**:
   - **Auto Select**: If you select this option, the system automatically selects the interface group to use for the VLAN.
   - **Interface Group**: If you select this option, you must select the interface group to use for the source.

   > **Note:** If You select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for all the objects in the protection job.

If the registration is successful, the registered physical server appears in the list on the **Sources** page. Register this physical server as an Oracle Server to register its databases.

Register Oracle Cluster

Before You Start

Verify the following before registering the RAC source:

1. The system will first scan `/etc/oraInst.loc` on each node of RAC to get the ORA inventory file path to get the GRID HOME location. If the GRID HOME location is present in the inventory file path, we do not need to parse the `/etc/oratab` file to look for GRID_HOME.
2. If GRID HOME is not found in the`/etc/oraInst.loc` file, Cohesity will proceed to look for GRID HOME from `/etc/oratab` file. Entries in the ORATAB file that starts with the '+' character will be considered as ORATAB entries referring to the GRID HOME. One of such example would be the ASM instance, which would have an entry that starts with '+ASM1:' in the ORATAB file. If multiple such entries are found, Cohesity will use the first entry to retrieve the GRID HOME. The following example indicates the ASM instance and its correspond GRID HOME: +ASM3:/u01/app/12.1.0/grid_1:N.
3. The registration should be done using SCAN IP or NAME (Short Name) so that the Cohesity agent will detect it as RAC.
4. The Cohesity agent must be running on all the nodes in the RAC cluster.

5. Choose your Oracle Authentication method: OS Authentication (default) or DB Authentication. For more information, see Credentials and Privileges for Oracle Database.

To register the Oracle cluster, follow the steps listed below:

1. Navigate to **Data Protection > Sources**.
2. For an Oracle cluster (RAC - Real Applications Clusters or AP - Active-Passive), Click **Register** > **Databases** > **Oracle Source**.
3. In the **Select Source Type** drop-down, select **Oracle Cluster**.
4. **Enter SCAN/VIP Address**, for an Oracle cluster, specify the SCAN IP or NAME (Short Name or FQDN) of the Oracle cluster.
5. **Enter Connect Endpoint (optional)**, specify a reachable endpoint for your Oracle cluster, this endpoint should be reachable the first time when Cohesity discovers the Oracle cluster. This optional endpoint can be provided if the Scan endpoints are not reachable from the Cohesity cluster.
6. **Authentication Type**, specify the Authentication Type to register the Oracle source.

   When connecting to the Oracle databases, Cohesity can use either the Oracle OS Authentication or Oracle DB Authentication method. If you use DB authentication, provide the database username and password.
7. **Cohesity Network Interface**, Select one of the following options to assign the interface group for the Oracle source.
   - **Auto Select**: If you select this option, the system automatically selects the interface group to use for the VLAN.
   - **Interface Group**: If you select this option, you must select the interface group to use for the source.

   > **Note:** If You select a different interface group from what was selected during the Physical source registration, the system will use this VLAN for all the objects in the protection job.
8. Click **Register**.

If the registration is successful, the registered Oracle cluster appears in the list on the Sources page.

Register an Oracle RAC Cluster

> **Note:** Cohesity does not support protection of Oracle RAC Cluster running on Windows.

Before You Start

Verify the following before registering the RAC source:

1. The system will first scan `/etc/oraInst.loc` on each node of RAC to get the ORA inventory file path to get the GRID HOME location. If the GRID HOME location is present in the inventory file path, we do not need to parse the `/etc/oratab` file to look for GRID_HOME.
2. If GRID HOME is not found in the`/etc/oraInst.loc` file, Cohesity will proceed to look for GRID HOME from `/etc/oratab` file. Entries in the ORATAB file that starts with the '`+`' character will be considered as ORATAB entries referring to the GRID HOME. One of such example would be the ASM instance, which would have an entry that starts with '`+ASM1:`' in the ORATAB file. If multiple such entries are found, Cohesity will use the first entry to retrieve the GRID HOME. The following example indicates the ASM instance and its correspond GRID HOME:

   ```
   +ASM3:/u01/app/12.1.0/grid_1:N
   ```

3. The registration should be done using `SCAN IP or NAME` (Short Name) so that the Cohesity agent will detect it as RAC.
4. The Cohesity agent must be running on all the nodes in the RAC cluster.
5. Choose your Oracle Authentication method: OS Authentication (default) or DB Authentication. For more information, see Credentials and Privileges for Oracle Database.

Register an Oracle RAC Cluster

The steps for registering an Oracle RAC cluster are the same as registering the Oracle standalone server. For more information, see Register Oracle Server.

Register an Oracle Veritas Cluster Server

> **Note:** Cohesity does not support protection of Oracle VCS Cluster running on Windows.

1. Navigate to **Data Protection > Sources**. On the **Sources** page, click **Register** > **Physical Server**.

   > **Note:** A physical server refers to a bare metal server or a VM running the supported versions of the OS and Oracle Database with the Cohesity agent installed. If you register Oracle without selecting the Physical Server, you should populate some of the fields manually when you register the database.

2. Enter the **VCS VIP** of the VCS cluster and then click **Register**.

   > **Note:** You can enter only one IP address in the field.

   Use the individual IP addresses or the nodes participating in a VCS cluster to register them as a standalone physical machine.

3. Select one of the following options from the Cohesity network interface:
   1. **Auto Select**: If you select this option, the system automatically selects the interface group to use for the VLAN.
   2. **Interface Group**: If you select this option, you must select the interface group to use for the source.

If the registration is successful, the registered server appears in the list on the Sources page.

Backup Oracle Database

The Cohesity Oracle adapter performs backup and recovery procedures and protect your Oracle databases against data loss or corruption and reconstructs the data if a data loss occurs.

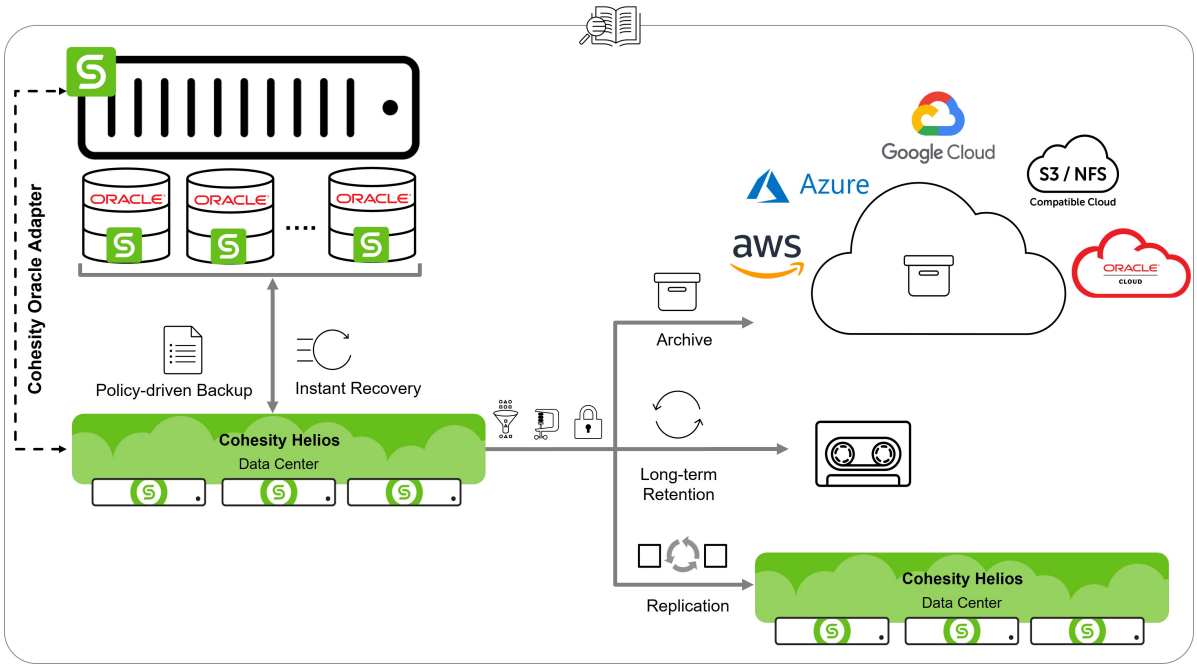

The Recovery Manager (RMAN) is an Oracle utility that can back up, restore, and recover database files. Cohesity Oracle adapter integrates with Oracle RMAN to deliver an integrated and centralized approach to manage Oracle backup jobs and recovery jobs using a single dashboard.

Cohesity uses the Incremental Merge capability of the Oracle database to create an image copy of a database and periodically update the image copy by merging incremental changes into a new backup copy. The resulting image copy can be used for restore and recovery, duplicating databases, and switching a database to the image copy as shown below.

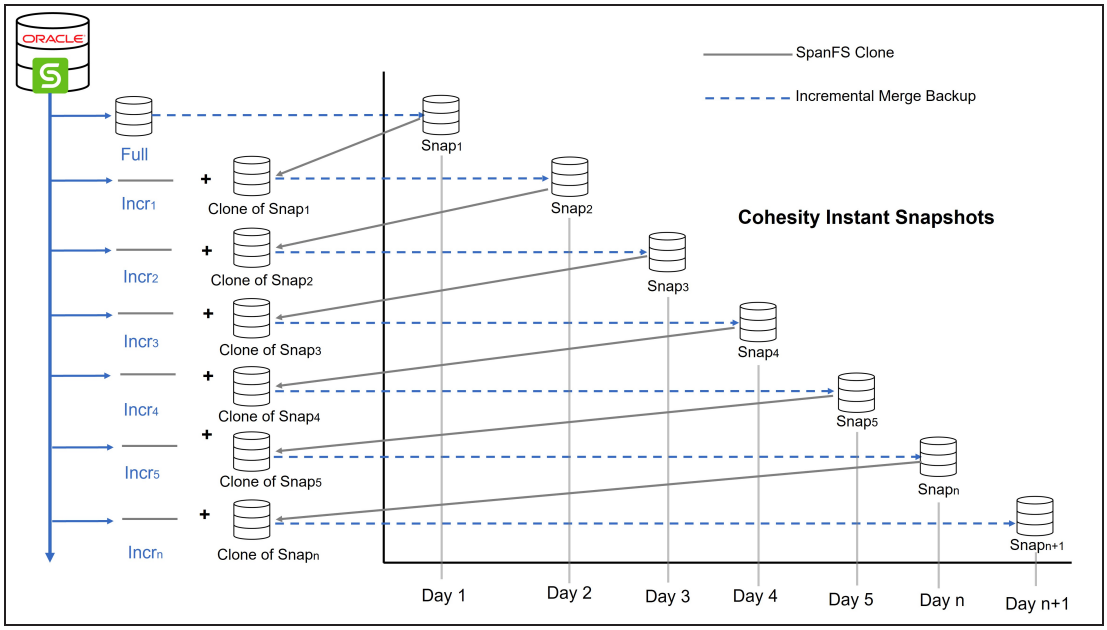

Cohesity Oracle Adapter Performance

Cohesity Oracle adapter supports multi-threading using multiple RAC nodes and multiple channels for improving backup performance and uses RMAN channels creation Logic given below:

```
RMAN channels=MIN(2* # of Cohesity nodes, 2* # of target server CPU
```

> **Note:**
> From Cohesity version 6.4, the backup channels on a single instance and RAC database nodes can be increased from default to 64. (However, the limit from Oracle is 255 channels).

Protect Oracle Database

When you create a protection group, you can select an existing source, policy and storage domain or you can also register a new oracle source, or create a new policy or create a new storage domain while creating a new oracle protection group. However, you might find it easier to create them before creating the protection group, as described in the following topics.

- Register Oracle Server
- Create or Edit a Standard Policy

- Create or Edit Storage Domains

> **Note:** You can now protect Oracle databases through an Oracle adapter and SBT plugin on Linux Oracle servers running in dual-stack mode(IPV4 and IPV6). This applies to both RAC and standalone Oracle servers.

Create a Protection Group

Follow the steps listed below to create a Protection Group for the Cohesity Oracle adapter:

1. Navigate to **Data Protection > Protection** and do one of the following:
    - To create a job, click **Protect** > **Databases** > **Oracle Database**, where *Job Type* is the type of the Protection Group.
    - Alternatively, click the action icon (⋮) next to the job and select the appropriate options.
        - **Run Now**: Start a run of this Protection Group. You have the option to back up one or more objects.
            - **Backup**: Select the objects you want to backup during the protection group run.
            - **Backup Type**: If the policy supports multiple backup types, you can select an **incremental**, **full**, or **log** backup.

              If the policy associated with the Protection Group is configured with replication and archive targets, you can select those targets and set the archive parameters.

              If you select **Logs** as the backup type, after the log backup run is completed, the logs will be archived to the target configured in the policy.

              > **Note:** When archiving Oracle archive logs, the logs are stored locally and not archive retention specified in the policy.
        - **Resume**: Resumes the paused protection job run.
        - **Edit**: Edits the selected protection group.
        - **Delete**: Deletes the selected protection group.
2. In the **New Protection** pop-up, do one of the following:

- Protect the Oracle database by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
  i. **Add Objects**: Select the registered Oracle source or click **Register Source** to register an Oracle source. Select the database instances to protect.
  ii. **Protection Group**: Specify a name for the Protection Group.
  iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
  iv. **Storage Domain**: Select the appropriate storage domain.

  > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

  v. Click **Protect**.
- To protect the Oracle database by specifying all the options, click **More Options**.

3. On the New Oracle Job page, enter the **Job Title** and select the **Sources** from the list.

   You can either specify an existing protection group or create a new one

4. On the Oracle Servers page, select the database instances to protect, and then click **Add**.

   When creating a new Protection Group, a shield next to the Object name indicates the Object is already been protected by another Protection Group.

   Click the icon to see details about the object.

   If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.

   The New Oracle Job page returns. Under Objects, verify that the database instances you selected are listed.

5. Set the policy for your Protection Group. Under **Policy**, select the appropriate Protection Policy, choose one of the defaults (Bronze, Silver, or Gold), or create a New Policy if the available policies do not meet your schedule and retention needs.

   For Oracle databases running on Windows OS, you must create or select a policy with the **Periodic Full Backup** option set.

Note:

- For Oracle RAC databases, you can now choose one, some, or all of the RAC system nodes for backups or restores.
- If the protection job includes standby databases, there's a risk that incremental or full backups may miss some logs needed for recovery. These logs will be captured in later backups, depending on when the standby databases received them. To mitigate this, log backup snapshot retention should match that of incremental/full snapshots.

  If the incremental/full snapshots lack all necessary logs for recovery, the backup process will issue a warning indicating that the "Backup may not be self-sufficient."

Tip: If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified error states, it can still be added to a Protection Group. However, during the Protection Group run, no snapshot for the object will be created. If the object becomes accessible, a snapshot is created the next time when the Protection Group runs.

6. Click the edit icon next to the selected Object and depending on the operating system of the server, set the following options for the Object, and then click **Save**.

Options	Description	
	Linux	**Windows**
System selects active node	Cohesity will auto-select an active RAC node and configure the number of RMAN channels for the database object. This option is enabled by default.	In the **SBT Library Path** field, specify the path to the Cohesity SBT Library. Cohesity SBT Library is installed on the Oracle Server along with the Cohesity Agent. This option is enabled by default.

Options	Description	
	Linux	Windows
Select specific nodes	For the RAC database, If you select this option, you can choose the number of RAC nodes and the number of RMAN channels and port to be used for each RAC node of the database object. For a standalone database, if you select this option, you can choose the number of RMAN channels and port to be used for the node of the database object.	You can also select this option and to specify the path to the Cohesity SBT Library and to configure the number of RMAN channels for the database object.
Update Database credentials	Enable this option to provide the updated database credentials of the user with SYSDBA privilege.	NA
Delete Archive Log	Toggle on and specify the days after which the archived logs on the source database should be deleted. If you specify the value as "0" days, source archived logs will be deleted immediately after each successful Cohesity backup. If you choose this option, Cohesity issues `RMAN DELETE ARCHIVELOG` command to delete archived logs after each successful Cohesity backup. Logs are deleted if they are backed up at least once by RMAN and they are older than the days specified in this option. In addition, if you have archived logs not backed up previously by Cohesity and the time satisfies the log delete eligibility, then these logs will be deleted. However, If you do not enable this option, Cohesity will not delete the archived logs after each Cohesity backup. If you are using the “Oracle RMAN archived log delete policy” to manage archived logs deletion on this database, depending on the policy, it may prevent Cohesity from deleting the archived logs as Oracle manages the policy. If you have upgraded Cohesity Data Cloud (both the Cohesity cluster and Cohesity agent) from a previous version, then the archived logs for the existing Protection Group will not be deleted. Starting with release 6.4.1, the archived logs are not deleted by default.	

Options	Description	
	Linux	Windows
Data Guard role-based backup	Enable the Data Guard role-based backup option to back up Oracle databases according to their role in the Data Guard configuration. Options include: • **Backup when in Primary role**: Choose this option to back up the database if it is in the Primary role only. • **Backup when in Standby role**: Choose this option to back up the database if it is in the Standby role only. Allow backup of archive logs in the Standby/Primary role: If your Protection Policy includes log backup configuration, this enables exclusive log backup based on the database role defined in the Data Guard configuration. If you choose the Backup when in the Primary role option, this allows backing up archive logs when the database is in the Standby role. Conversely, if you choose the Backup when in Standby role option, this allows backing up archive logs when the database is in the Primary role. This ensures the continuity of Cohesity's log backups during Data Guard role transitions and supports Oracle's archived log deletion, provided the feature is enabled on the Cohesity cluster. > **Note:** Enabling this option may increase backup storage requirements due to the additional log backups.	

7. Specify the number of channels to use for each node and a single instance DB.
8. Select the appropriate **Storage Domain**.

   A Storage Domain defines the level of storage efficiency, security, and resilience for the data being protected in Cohesity Data Cloud. In a Storage Domain, you can enable or disable deduplication, compression, and/or encryption for the domain, and select the necessary level of resilience (that is, the number of faults tolerated).

   If a Storage Domain with the right settings exists, select that one. Otherwise, create a new one with Add Storage Domain. See Create or Edit Storage Domains.

   If you need to create a new Storage Domain:

   - You can enable or disable deduplication and compression at any time.
   - You can enable or disable encryption only when you create the Storage Domain for the first time.

- There are no benefits if you perform deduplication and compression on the encrypted data.

9. After selecting a Storage Domain, if you need to change any of the Advanced settings on the New Oracle Job page, scroll down and click **Edit** on the right.

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the job should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
End Time	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. This option is selected by default. > **Note:** In the 6.6.0a release, The Backup Auto option is available only on the C6000 Series. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Advance Settings	Description
Pre & Post Scripts	Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If the Protection Group is protecting Oracle databases from different hosts, then the pre and post scripts are executed for each Oracle server. For more details, see Configure Pre & Post Scripts.
NFS Protocol	Select the NFS protocol version for the backup job. Options include NFSv3 and NFSv4.1. **Note:** • NFS protocol is not supported when performing an SBT backup. • NFS protocol is not supported for Windows and AIX databases.
Persist Mountpoints	If you enable this option, NFS mounts created during the Oracle database backup job are retained after the job is completed. Disabling this would remove the NFS mounts after the backup job is completed. By default, this option is enabled. You can disable this if needed.
Alerts and Priority	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: 1. High-priority jobs 2. Medium-priority jobs 3. Low-priority jobs.
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the job.

Advance Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Cluster Interface	Select one of the following options to assign the interface group for the protection group. • **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration. • **Interface Group**: Select the interface group to use for the backup. **Note:** If You select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for all the objects in the protection job.
Description	Optional. Click the plus sign and enter a description for the Protection Group.

10. **Run Now**: Start a run of this Protection Group. You have the option to back up one or more objects.

    If the policy associated with the Protection Group is configured with replication and archive targets, you can select those targets. If the policy supports multiple backup types, you can select an **incremental**, **full**, or **log** backup.

    If you select **Logs** as the backup type, after the log backup run is completed, the logs will be archived to the target configured in the policy.
11. Click **Protect**.

Your Protection Group runs appear in the list on the Protection Groups page.

Download Debug Logs for Oracle Backup Jobs

You can now download the debug logs (`rman_shell.INFO`, `precheck_logs.INFO`, `alert_log.INFO`, `tnsnames_ora.INFO`, `sparse-db-config.txt`) as a compressed file for both successful and failed backup runs.

To download the debug logs:

- **For a job run**: To download the debug logs for a job run, click the **Options** (⋮) icon next to a completed or a failed job run and select **Download Debug Logs** to download debug logs for all the databases in the job.
- **For a Specific database within the job run**: Optionally, you can download logs for a specific database within the job run. Click the **Options** (⋮) icon next to a database and select **Download Debug Logs**.

> **Note:** You can download the debug logs only if the local snapshot is available on the Cohesity cluster.

Once your Protection Group completes its first run, the database will be ready for recovery, replication, and archival.

Recover Oracle Database

The Cohesity Oracle adapter performs recovery procedures to restore a physical backup of a data file or control file and make it available to the Oracle database server using the Recovery Manager (RMAN).

Cohesity integrates with the Oracle Recovery Manager (RMAN) to restore data files from backup sets or image copies on disk, either to their current location or to a new location. RMAN automatically restores the archived logs that are needed for recovery and deletes them after the recovery is finished as shown in the illustration below.

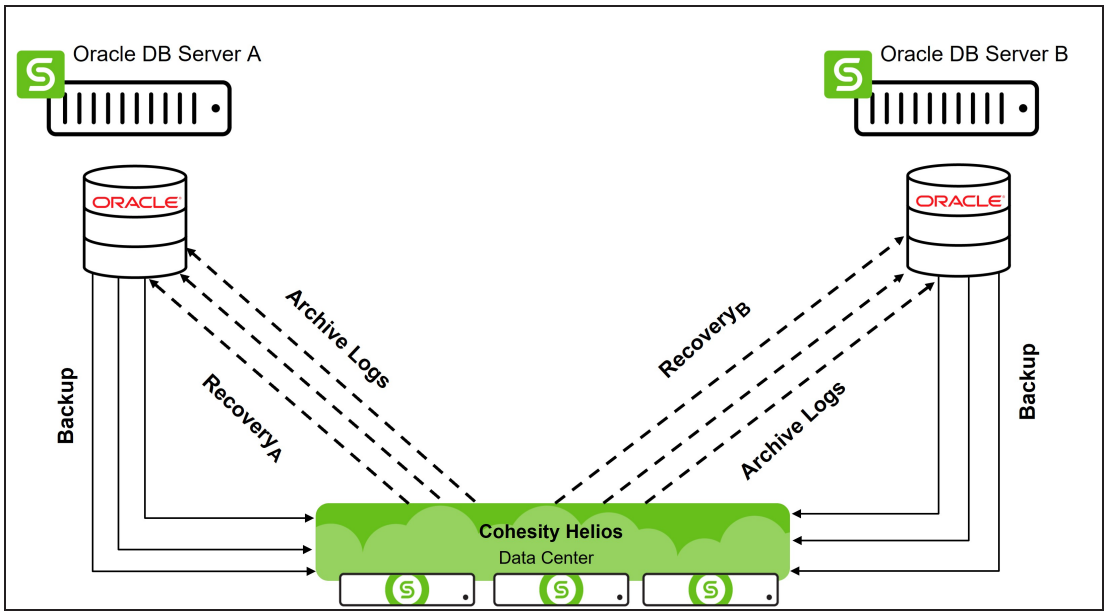

The Cohesity Oracle adapter uses the Incremental Merge capability of the Oracle database to perform a backup of database files. The Incremental Merge process involves IMAGE COPY backups that are typically indistinct copies needing redo to be applied to make them consistent. The process makes use of a single RMAN script for backup and restore. The

merge process creates a single backup dump that you can use for instantaneous recovery of the backed-up data.

Oracle Database Restore Types

The following options are available for an Oracle database recovery:

Recovery Type	Single Instance	RAC	VCS
Overwrite Existing Database	Original Oracle Server	Original Oracle RAC Server	Original Oracle VCS Server
As a new database	Original Oracle Server*	Standalone DB on ASM	Original VCS Oracle Server
	Alternate Oracle Server	Standalone DB on ASM	Alternate VCS Oracle Server

*If you are restoring the Oracle database to the original server as a new database, then you must manually enter the new database name in the **tnsnames.ora** file.

Considerations for Oracle Restore

Review the following considerations before you restore from the Oracle RMAN tool and is applicable for the Cohesity Oracle Adapter.

- Cross-platform Oracle database restore is not supported. (for example, Linux vs Other : AIX, SLES, HP UNIX, etc.)
- Restore across versions is not supported (for example, restoring 11G database to 12c and vice versa)
- Restore from the incremental snapshot is not supported when there is no preceding full backup (as it might have expired) for Oracle on Windows backups.

Recover Oracle Database

Enterprise customers may choose to use the Oracle container Database (CDB) model to leverage the benefits in consolidation and data management. This configuration is also known as Oracle’s multitenant (MT) architecture.

A CDB includes zero, one, or many customer-created pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and non-schema objects that appear to an Oracle Net client as a non-CDB.

Cohesity now supports fully automated data protection and management for Oracle CDBs and PDBs Cohesity provides flexible data protection capabilities at both the CDB and PDB level.

To recover a recent or previous backup of an MT and non-MT Oracle Database, perform the steps listed below.

1. Navigate to **Data Protection > Recoveries** > **Recover** > **Databases** > **Oracle**.
2. In the Recover page, search for the Oracle instance that needs to be restored or use a wildcard name to search for a specific database instance.
    a. **CDB**: Search the Oracle CDB that you want to restore.
        i. Select the entire CDB (with all the PDBs) or a CDB and a subset of PDBs associated with that CDB for restore.
        ii. Select the CDB and click **Next**.
    b. **PDB**: Search for the Oracle PDB that you want to restore.
        i. You can either search for the parent CDB associated with the PDB or directly search for the PDB.
        ii. Select the PDB and click **Next**.
    c. **Oracle Database**: Search for the Oracle Database you want to restore.
    d. **RAC**: Search for the Oracle RAC you want to restore.
3. Accept the latest **Recovery Point** or click the change icon ( ) and select a different recovery point.
    1. **Snapshots**: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
    2. **Point in Time**: Use the slide control to adjust the date and time and click **Save**.
4. After you select an instance for recovery, configure these fields in the Recover page:

Alternate CDB/PDB/RAC

Provide the following information to **Alternate CDB/PDB/Database**:

> **Note:**
> - If restoring to an alternate Oracle server, the Oracle server must be registered with Cohesity before you perform the restore operations. Ensure that the Cohesity Oracle agent is installed on the Oracle server.
> - PDB alternate restore jobs will fail if the target CDB is registered with DB Authentication.

Field Name	Description

Alternate CDB/PDB	• **Oracle Host and RACs**: The host to which to restore the database.
	• **Configure Channels**: Click these options to configure the following:
	• **System selects active node**: If you select this option, Cohesity will select the channels for the recovery job.
	• **Select specific nodes**:
	  • **Hostname**: Select the hostname to which you want to perform the restore.
	  • **Node**: Select the node to which you want to perform the restore. This option is valid only in the case of a multi-node cluster.
	  > **Note:** In the case of PDB alternate restore it is mandatory to select the node as we have to provide the target CDB instance name.
	• **Channels**: Specify the number of channels to perform the restore job. The system will run these channels in parallel to improve the restore performance.

Field Name	Description

Recovery Options	• **Restore Database Files to**: Specify the path to an existing empty directory. The newly created database files will reside in this path.
	
	> **Note:** You can restore to an ASM path using this option. Specify just the ASM volume name instead of the entire path. Example: The restore path for ASM volume data1 is +data1. If you specify the entire path, +data1/dbname the restore job will fail.
	
	• **Oracle Home**: The ORACLE_HOME value for the host where the database is restored.
	• **Base Directory**: The directory for the restored database.
	• **Target Database Name**: The name for the alternate CDB database.
	
	If you are restoring a PDB, you must provide the target CDB on to which the PDB is restored.
	
	> **Note:** When recovering to an alternate location, make sure a database with the same name does not exist on the Oracle server instance. Otherwise, the recovery task will fail.
	
	If you are restoring **RAC** database, then you can either select standalone or RAC database.
	
	> **Note:** RAC to Standalone restore is supported. RAC-specific options are displayed only when Target Database Type is set to RAC.
	
	• **Target RAC Configuration** (This option is displayed when Target Database Type is set to RAC)
	  • **Hostname**: Displays the hostname of the RAC node that can participate in the target RAC database.
	  • **Node IP**: Shows the IP address associated with the target RAC node. This option is valid only in the case of a multi-node cluster.
	  • **Source Instance Name**: Displays the source RAC instance selected by the system for this node. You can manually select the source instance from the drop-down list.
	  • **Target Instance Name**: Enter the name of the Oracle instance to be created or used on the target RAC node. Each target instance name must be unique within the cluster.

Field Name	Description
	• **Configure Chanels**: Specify the number of Migration Streams to perform the restore job. For details, refer to the descriptions above. • **SPfile location for Target RAC**: Depending on the Recovery options, the SPfile location will be automatically populated. • **Pfile Parameters**: The Pfile parameters option allows you to see the target database parameters before the restore job is initiated. The Pfile inherits some source database parameters to the target database and allows you to edit the target database parameters based on your requirement by specifying the additional parameters as a `key=value` pair. The Pfile lists the following parameters: • **Default Parameters**: Default parameters of the target database and you can edit the default • **Inherited Parameters**: Parameters from the source database and you can edit the inherited parameters, if required. • **Restricted Parameters**: You cannot edit the restricted parameters in the Pfile. The following are the restricted parameters: • INSTANCE_NAME • DB_NAME • DB_UNIQUE_NAME • DB_BLOCK_SIZE • MAX_STRING_SIZE • COMPATIBLE If you edit the following values in the **Recover Oracle** GUI, You must click the **Reset** button on the top-right corner to reset the Pfile and get the updated parameter values in the Pfile box. • Restore database files to • Oracle Home • Oracle Base • Database Name Cohesity Platform will create a target database with the specified parameters in the Pfile. For example, if `SGA_TARGET` is specified as a part of the Pfile parameters, the restore operation will use the value specified in the Pfile during the restore.

Field Name	Description
	> **Note:** Customizable pfile functionality is not supported for archival backups. • **NFS Protocol** - Select the NFS protocol version for the recovery job. Options include NFSv3 and NFSv4.1. > **Note:** NFS protocol is not supported for Windows and AIX databases. • **Rename Database**: Rename the database, if required. • **Enable Archive Log mode for the Database**: Select the checkbox if you want to enable the redo log archiving on the recovered database. • **BCT File Path**: The BCT file path specifies the location where the block change tracking file will be created. If not provided, BCT is not enabled for the restored database. • **Number of Temporary Files to Use**: Number of temp files to rename from the source database. • **Redo Log Group Settings**: The Redo log group configuration stores all changes made to the database as they occur. Every instance of an Oracle Database has an associated redo log to protect the database in case of an instance failure. You can provide the number, size, and prefix of redo log groups along with the redo log group member path for the recovery database. • **Leave Database in Recovery Mode**: Select the checkbox if you do not want the recovered database to be opened. • **Shell Environment**: Toggle on the Shell Environment to configure your shell environment that executes the Cohesity restore workflow using this setting. For example, when restoring the Oracle EBS database, the restore environment often needs to define a different TNS_ADMIN value to use a different sqlnet.ora file in the same Oracle Home shared with other databases. You can use Shell Environment to define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your restored target database. For TDE support, the wallet location for your restore target database may depend on a shell variable. This is the case when your restored database wallet location is specified either by the ENCRYPTION_WALLET_LOCATION or WALLET_LOCATION keyword in your sqlnet.ora file. A typical value for these parameters will reference a shell variable. For Example, `ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)`

Field Name	Description
	`(METHOD_DATA=(DIRECTORY=/etc/ORACLE/WALLETS/$ORACLE_UNIQUE_NAME/)))` In this example, you will need to set ORACLE_UNIQUE_NAME in the SHELL ENVIRONMENT setting in the GUI to set up the Cohesity restore shell environment for the Oracle restore to find the Keystore. • **Source CDB Keystore Password**: The source CDB keystore password is the same as that of the source CDB, duplicated on the target CDB host. This duplicated keystore is necessary for the PDB alternate restore workflow, which temporarily clones the source CDB onto the target CDB host. You can configure the location of the duplicated keystore using the WALLET_ROOT pfile parameter or the ENCRYPTION_WALLET_LOCATION parameter in the sqlnet.ora file. Consider the following guidance to determine the best location for the duplicated keystore. Consider the following guidance to decide the location of the duplicated keystore. • Configuring the keystore location using WALLET_ROOT is recommended if the database version is 18c or above. • If WALLET_ROOT is configured in the pfile, TDE_CONFIGURATION must be set to 'KEYSTORE_CONFIGURATION=FILE'. • If ENCRYPTION_WALLET_LOCATION of the sqlnet.ora file is used to determine the keystore location, don't include variables such as ORACLE_SID in the path because the temp cloned DB instance name/SID are dynamically generated. • If the sqlnet.ora file is not in the default directory $ORACLE_HOME/network/admin, add its directory to the shell environment variable in the Cohesity UI of the PDB alternate restore. • The duplicated keystore can be deleted after the PDB alternate restore is finished. **Use External Store** Enable this option to utilize the external store for the password of the duplicated keystore from the source CDB, rather than entering the password manually in the UI. The external store must be the 'file tde_seps/cwallet.so' at the path specified by the WALLET_ROOT pfile parameter or cwallet.so at the location indicated by the EXTERNAL_KEYSTORE_CREDENTIAL_LOCATION pfile parameter. • **Target CDB Keystore Password**: The Target CDB keystore password protects the Keystore for the Target CDB and must be unlocked before the PDB alternate restore.

Field Name	Description

	Use External Store
	Enable this option to utilize the external store for the password of the duplicated keystore from the target CDB, rather than entering the password manually in the UI.
	The external store must be the 'file tde_seps/cwallet.so' at the path specified by the WALLET_ROOT pfile parameter or cwallet.so at the location indicated by the EXTERNAL_KEYSTORE_CREDENTIAL_LOCATION pfile parameter.
	• **Cluster Interface**: Select one of the following options to assign the interface group for the recovery job.
	  • **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration.
	  • **Interface Group**: Select the interface group to use for the recovery.
	If You select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for the recovery job.
	• **Use NOFILENAMECHECK**: Using NOFILENAMECHECK forces RMAN to skip checking whether target datafile names are the same as source datafile names.
	Warning: Use NOFILELNAMECHECK with caution, as this restore can inadvertently clobber the source datafiles if restored back to the source host.
	Note: Some restore use case requires restoring the datafiles and/or online logfiles using the same names and file paths as the source database, but to a different host (i.e. not restoring back to the source database host).
	For example, using this restore to create a standby database of the source database with exactly the same directory structure on a different host. In this case, it is required to set the 'Use NOFILENAMECHECK' option for this Cohesity Alternate Restore.
	• **Customize Destination Paths with SET NEWNAME**: Use the SET NEWNAME command to define a new location for individual datafiles, tempfiles, tablespaces, or the whole database.
	If not specified, Cohesity uses DB_FILE_NAME_CONVERT, if not empty. Otherwise, it uses your input to restore the destination location. Use SET NEWNAME specifications to determine the location of your restored datafiles and logfiles
	You can specify DB_FILE_NAME_CONVERT together with this option to augment paths that are not affected by DB_FILE_NAME_CONVERT.

Field Name	Description
	Syntax Example: `SET NEWNAME FOR DATAFILE 1 TO 'u01/home/test_1.dbf';` • **Task Name**: (Optional) Provide the name for the recovery task.

Overwrite Original CDB/PDB/Database

> **Note:** For PDB overwrite restore, PDB can be rolled back only within the current CDBs set.

Provide the following information to **Overwrite Original CDB/PDB**:

Field Name	Description
Overwrite Original CDB/PDB	• **Oracle Host and RACs**: The host to which to restore the database.

Field Name	Description

Recovery Options	• **View Name Identifier**: A view will be created with this name as identity.
	• **Leave Database in Recovery Mode**: Select the checkbox if you do not want the recovered database to be opened.
	• **NFS Protocol** - Select the NFS protocol version for the recovery job. Options include NFSv3 and NFSv4.1.
	> **Note:** NFS protocol is not supported for Windows and AIX databases.
	• **Shell Environment**: Toggle on the Shell Environment to configure your shell environment that executes the Cohesity restore workflow using this setting. For example, when restoring the Oracle EBS database, the restore environment often needs to define a different TNS_ADMIN value to use a different sqlnet.ora file in the same Oracle Home shared with other databases. You can use Shell Environment to define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your restored target database. For TDE support, the wallet location for your restore target database may depend on a shell variable. This is the case when your restored database wallet location is specified either by the ENCRYPTION_WALLET_LOCATION or WALLET_LOCATION keyword in your sqlnet.ora file. A typical value for these parameters will reference a shell variable.
	For Example,
	`ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)(METHOD_DATA=(DIRECTORY=/etc/ORACLE/WALLETS/$ORACLE_UNIQUE_NAME/)))`
	In this example, you will need to set ORACLE_UNIQUE_NAME in the SHELL ENVIRONMENT setting in the GUI to set up the Cohesity restore shell environment for the Oracle restore to find the Keystore.
	• **Cluster Interface**: Select one of the following options to assign the interface group for the recovery job.
	  • **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration.
	  • **Interface Group**: Select the interface group to use for the recovery.
	If You select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for the recovery job.
	• **Perform point-in-time incomplete recovery** - Select this option to recover the Oracle database from media failures, tablespace corruption, or tablespace being dropped

Field Name	Description
	accidentally. The database can be restored till the backup is available with Cohesity. Optionally the database can also be rolled forward, if archive logs are still available on the Oracle host . Using this option database will open with the RESETLOGS option. By default, this option is selected. • **Attempt complete recovery** - Select this option to roll forward the databases from the latest backup available with Cohesity. For a complete recovery to succeed, you must have the complete set of archived logs (after the last Cohesity backup), current control file, and online redo logs available on your source database host. If complete recovery is successful database will open with "ALTER DATABASE OPEN" and if the complete recovery attempt failed database will be open with the RESETLOGS option. • **Task Name**: (Optional) Provide the name for the recovery task.
Roll Forward	You can use the Roll Forward recovery option to restore the Oracle databases to the state it was in just before the failure occurred. We will auto-select the latest backup available and apply the provided archive logs to roll-forward the database. The following inputs will be needed: • **Archive log location**: Path at which archive logs are present on the source Oracle host. In the case of a clustered Oracle host, this location can be an ASM path or a path that is present (shared) in all the cluster nodes. • **Restore point-in-time**: Optional field to specify the date-time to which the database must be rolled forward. This date-time must be in the future of the latest Cohesity backup. If nothing is specified, then the database will be restored at the latest time possible. If the Roll Forward option is disabled the database will be restored to the selected point-in-time. > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Create Cohesity View with DB Files - CDB/PDB/Database

Provide the following information to **Create Cohesity View with DB Files**:

Field Name	Description
Create Cohesity View with DB Files - CDB/PDB	• **Oracle Host and RACs**: The host to which to restore the database.

Recovery Options	• **View Name Identifier**: A view will be created with this name as identity.
	• **Cluster Interface**: Select one of the following options to assign the interface group for the recovery job.
	  • **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration.
	  • **Interface Group**: Select the interface group to use for the recovery.
	If You select a different interface group from what was selected during the Oracle source registration, the system will use this VLAN for the recovery job.
	• **NFS Protocol** - Select the NFS protocol version for the recovery job. Options include NFSv3 and NFSv4.1.
	Note: NFS protocol is not supported for Windows and AIX databases.
	• **Task Name**: (Optional) Provide the name for the recovery task.

Instant Recovery - CDB/Oracle Database (Standalone and RAC)

Note:

- The Instant Recovery workflow is applicable only for CDB and Oracle Databases.
- Instant recovery will be slower as compared to running a DB on primary till the migration is completed.

Instant Recovery provides near-instant RTO for databases. This is particularly helpful in disaster recovery situations and in cases where complete database failure or downtime can have a significant business impact on the users.

With the Oracle Instant Restore functionality, you can quickly bring up the Oracle databases similar to the Clone operations. Once the database is successfully brought

up, you can choose to migrate the DATAFILES and LOGFILES from the Cohesity storage to the target Oracle host while the newly restored Oracle database is active and open for read-write operation.

The migration operation in the Instant Restore is seamless, and there is no downtime of the Oracle database. Hence using this functionality, you can achieve a near-instant RTO for the Oracle databases.

> **Note:** The Instant Recovery functionality is supported from Oracle 12c and later.

Provide the following information to **Instant Recovery**:

Select one of the following options to proceed with the instant recovery of the source database:

Instant Migration

If you select this option, the database will be brought up, but the DATAFILES and LOGFILES are present in the Cohesity storage. This is similar to the Clone procedure of the Cohesity recovery workflow for Oracle.

> **Note:** The Instant Recovery functionality is only applicable to Linux platform.

Manual Migration

In manual migration, you will need to manually migrate the DATAFILES from Cohesity storage to your source database. An option for migration will be available on the Restore page, after successful restore.

To start the database migration, provide the following details:

Field Name	Description
Instant Recovery	• **Oracle Host and RACs**: The host to which to restore the database. • **Disaster Recovery** - If this option is enabled, it creates a clone database using the same database name, database ID, and database parameters as the original database. After migration, all database files will be restored to the original location of the source database. > **Note:** If the disaster recovery option is selected, you cannot change the host to which the database is restored to. **Configure Channels** - If **Disaster Recovery** is enabled, the **Hostname** and **Node IP** of the host are selected by default. You can specify the number of Migration Streams to perform the restore job. If **Disaster Recovery** is disabled, then share the following details: • **System selects active node**: If you select this option, Cohesity will select the channels for the recovery job. • **Select specific nodes**: • **Hostname**: Select the hostname to which you want to perform the restore. • **Node IP**: Select the node to which you want to perform the restore. This option is valid only in the case of a multi-node cluster. > **Note:** In the case of PDB alternate restore it is mandatory to select the node as we have to provide the target CDB instance name. • **Channels**: Specify the number of channels to perform the restore job. The system will run these channels in parallel to improve the restore performance.

Field Name	Description
Recovery Options	• **Shell Environment**: Toggle on the Shell Environment to configure your shell environment that executes the Cohesity restore workflow using this setting. For example, when restoring the Oracle EBS database, the restore environment often needs to define a different TNS_ADMIN value to use a different sqlnet.ora file in the same Oracle Home shared with other databases. You can use Shell Environment to define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your restored target database. For TDE support, the wallet location for your restore target database may depend on a shell variable. This is the case when your restored database wallet location is specified either by the ENCRYPTION_WALLET_LOCATION or WALLET_LOCATION keyword in your sqlnet.ora file. A typical value for these parameters will reference a shell variable. For Example, `ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE) (METHOD_DATA=(DIRECTORY=/etc/ORACLE/WALLETS/$ORACLE_UNIQUE_ NAME/)))` In this example, you will need to set ORACLE_UNIQUE_NAME in the SHELL ENVIRONMENT setting in the GUI to set up the Cohesity restore shell environment for the Oracle restore to find the Keystore. • **NFS Protocol** - Select the NFS protocol version for the recovery job. Options include NFSv3 and NFSv4.1. **Note:** NFS protocol is not supported for Windows and AIX databases.

5. Toggle on the **Shell Environment** to configure your shell environment that executes the Cohesity restore workflow using this setting. For example, when restoring the Oracle EBS database, the restore environment often needs to define a different TNS_ADMIN value to use a different sqlnet.ora file in the same Oracle Home shared with other databases. You can use Shell Environment to define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your restored target database. For TDE support, the wallet location for your restore target database may depend on a shell variable. This is the case when your restored database wallet location is specified either by the ENCRYPTION_WALLET_LOCATION or WALLET_LOCATION keyword in your sqlnet.ora file. A typical value for these parameters will reference a shell variable.

   For Example,

   ```
   ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)
   (METHOD_DATA=(DIRECTORY=/etc/ORACLE/WALLETS/$ORACLE_UNIQUE_NAME/)))
   ```

In this example, you will need to set ORACLE_UNIQUE_NAME in the SHELL ENVIRONMENT setting in the GUI to set up the Cohesity restore shell environment for the Oracle restore to find the Keystore.

6. Click **Recover** to initiate the restore process.

Your recovery job runs appears in the list on the **Recoveries** page.

Download Debug Logs for Oracle Recover Jobs

Debug logs are system-generated logs that are used in troubleshooting Oracle database backup, recovery, and performance issues. You can now download the debug logs (backup logs and restore logs) as a compressed file for both successful and failed recovery job runs.

To download the debug logs:

Click the **Options** (⋮) icon next to a completed or a failed recovery job run and select **Download Debug Logs**.

Recover Oracle TDE Database

Introduction

Transparent Data Encryption (TDE) enables you to encrypt sensitive data that you store in tables and tablespaces. After the data is encrypted, this data is transparently decrypted for authorized users or applications when they access this data. TDE helps protect data stored on media (also called data at rest) in the event that the storage media or data file is stolen.

You can encrypt sensitive data at the column level or the tablespace level. At the column level, you can encrypt data using selected table columns. TDE tablespace encryption enables you to encrypt all of the data that is stored in a tablespace.

Supported Oracle TDE Wallet types

Cohesity does not impose restriction on wallet types used by the source database for the backup workflow. Your source database can use hardware or software keystore, as well as Oracle Key Vault to centralize your wallet keys management.

For restore and clone workflow, Cohesity requires you to set up the appropriate duplicated keystore on your target hosts before using Cohesity restore or clone workflow. In these workflows, only software keystores are supported and keystores must be put into the auto log-in mode before issuing the Cohesity restore or clone workflow.

Recover TDE Database

When you recover a Transparent Data Encryption (TDE) database backup, Oracle performs media recovery on files restored from Cohesity. To perform media recovery on encrypted blocks, Oracle requires access to the respective TDE keystore.

TDE keystore considerations

- Oracle mandates each database has its own keystore.
- A cloned Oracle database must have its TDE keystore with keys that were used during the previous restore operation.
- Cohesity does not store, manage, or solicit any TDE keystore passwords from users.
- Cohesity requires users to manage the TDE keystore for your source and restore target databases.

This design allows you to control and manage the operation and placement of your TDE keystores to ensure it meets your organization's security requirements.

To use Cohesity to restore from a backup taken from a TDE database, you must set up a duplicate keystore on the restore target host with a chosen keystore configuration mechanism that is supported by Oracle.

To duplicate a keystore from your source database, you can use one of the following methods as recommended by Oracle:

- Backup and close your source keystore and physically copy the keystore to a new location.
- Use the `ADMINISTER KEY MANAGEMENT CREATE KEYSTORE` command to first create an empty keystore, then use the `ADMINISTER KEY MANAGEMENT MERGE KEYSTORE` command to populate your new keystore with the keys from the source keystore.

Oracle supports keystore configuration in the following ways:

- For Oracle versions 18c and later, you can use DB parameters `WALLET_ROOT` and `TDE_CONFIGURATION` to direct your database on how to locate your keystore for your database.
- Use the `ENCRYPTION_WALLET_LOCATION` parameter in `sqlnet.ora` to configure your keystore location. The file path of this `sqlnet` parameter can be made dependent on a shell variable that can uniquely identify your database. For example, the value of the shell variable can be your `DB UNIQUE NAME`.
- Place your keystore in one of the following default locations:

```
$ORACLE_BASE/admin/<db_unique_name>/wallet
```

```
$ORACLE_HOME/admin/<db_unique_name>/wallet
```

> **Note:** Cohesity requires the restore target keystore to be in auto log-in mode. You can create an auto log-in keystore from your duplicated password-protected keystore. For more information, see Creating an Auto-Login or a Local Auto-Login Software Keystore.

If your Oracle database version is 18c or above, and you would like to use the DB parameters `WALLET_ROOT` and `TDE_CONFIGURATION` to configure your restore target

database keystore, you can use the Cohesity Version 6.5.1 and later to add the configuration of these DB parameters as part of the restore workflow.

If you choose to configure your restore target database auto log in keystore using `sqlnet.ora` and you are using a shell variable as part of the file path of the wallet, you have to provide the variable name and its values to be used to the Cohesity Restore workflow. For more information, see Recover Oracle Database.

Source database key rotations considerations

- Cohesity does not store, manage, or solicit any TDE keystore passwords from users. Ensure that you manage the TDE keystore for your source and restore target databases. Refresh your duplicated keystore on the restore target host from time to time if master keys are changed in your source database.
- You can change your source database master keys because a new master key is created, or the keys are rotated.
- A successful Cohesity clone operation from a chosen point-in-time is only possible for a TDE database if your duplicate clone target host keystore has the correct set of master keys from that point-in-time. You do not need the exact set of keys for that point-in-time. It is sufficient if you have those keys as well as more keys than you need in your duplicated keystores. Therefore, the easiest way is to always refresh your duplicate keystore from the latest version of your source keystore.
- Oracle does not provide commands or APIs to delete old master keys from the keystore. Keys can only be added or re-used by key rotation.

Granular Recovery for Oracle Database

To perform Granular recovery of Oracle databases, mount the corresponding NFS Views from Cohesity. After the Oracle server is protected by Cohesity, you can leverage the simplicity of Cohesity to easily mount the NFS View in which the RMAN backup files are stored as an image copy.

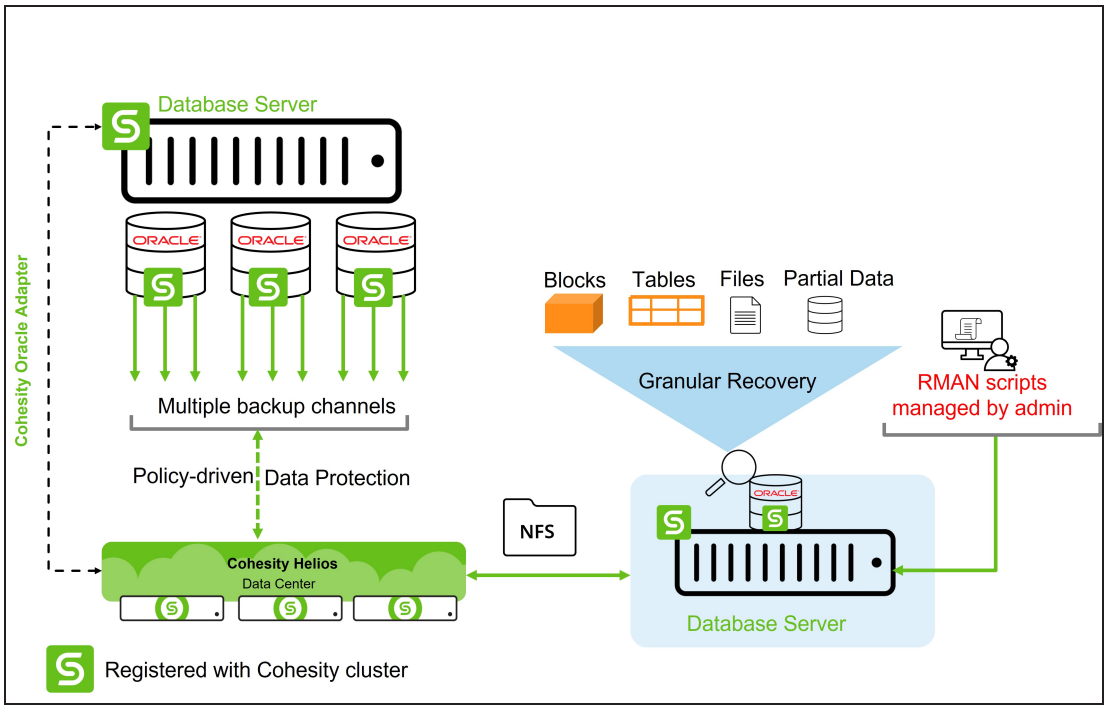

You can perform these recovery actions using NFS mount:

- Block-level recovery
- Table and tablespace recovery
- DB files (controlfile, datafile, SPfile) recovery
- Partial data recovery
- Cohesity backup view projecting as NFS
- Restore Oracle Database in recovery mode.

> **Note:**
>
> - To mount NFS to an alternate Oracle server, register the Oracle server with Cohesity Data Cloud. Ensure that the Oracle software is installed on the Oracle server and Recovery Admin should be self-sufficient with RMAN scripts needed for granular recovery. This feature requires the Cohesity Linux agent with version 6.4 and above.
> - While granular recovery is not performed directly through the Cohesity UI, Oracle DBAs can recover individual tables or tablespaces using standard Oracle scripts once the backup is exposed as a view.

To recover a recent or previous backup of an Oracle database, navigate to select **Data Protection > Recoveries** > **Recover** > **Databases** > **Oracle**. In the Recover Oracle

page, search for the protected Oracle instance and select the database planned for granular recovery or use a wildcard name to search for a specific database instance.

After you select an instance for recovery, configure these fields in the Recover page:

Field	Description
Objects Being Recovered	Under **Objects Being Recovered**, verify that the database instances you selected are listed.
Recover Point	If there are multiple backups, select the point of time that corresponds to the desired restore point from the drop-down.
Settings	Click **Create Cohesity View with DB Files** and enter the server details to which Recovery View NFS mount is planned.
Oracle Host	From the drop-down list, browse for the host which is registered as Oracle Server. View expose works only with registered Oracle servers.
View Name Identifier	Provide a name to identify the View at the target server when mounted successfully.

Once all the fields are configured, click **Recover** to restore your Oracle Server.

Recovery Validation for Oracle Database

> **Note:** The recovery validation task will affect the CPU, memory, and other resources on the staging target you select.

The Recovery Validation option allows you to check for physical and logical corruptions in the database before you can initiate a recovery job for an Oracle database.

To perform a recovery validation of an MT and non-MT Oracle Database, perform the steps listed below.

1. Select **Data Protection** > **Recoveries** > **Recover** > **Databases** > **Oracle**.
2. On the Recover page, select **Validation**.
3. Accept the latest **Recovery Point** or click the change icon () and select a different recovery point.
4. After you select an instance for recovery validation, configure these fields in the Recover page:

1. **Temporary Database**: Cohesity will create a temporary database on the selected host for this recovery validation task.
2. **Oracle Home**: ORACLE_HOME specifies the directory containing the Oracle database software binary files. By default, Oracle Universal Installer (OUI) installs the Oracle Database software binary files by version and Oracle Home Name in a subdirectory of the Oracle base directory.
3. **Base Directory**: ORACLE_BASE specifies the directory at the top of the Oracle software and administrative file structure. The Oracle base directory is the location where Oracle software and configuration files are stored.
4. **Task Name**: (Optional) Provide the name for the recovery task.

5. Click **Recover** to start the validation.

Recover Archive Logs for Oracle Database

You can also choose to recover the archive logs for a specific range of the target database. You can restore Archive logs using the LSN (Log Sequence Number), time range, and SCN (System Change Number).

To recover archive logs of an MT and non-MT Oracle Database, perform the steps listed below.

1. Select **Data Protection** > **Recoveries** > **Recover** > **Databases** > **Oracle**.
2. Select **Archive Logs** as the Recover Type.
3. Search for the Oracle instance that needs to be restored or use a wildcard name to search for a specific database instance.
4. Click the change icon (ICON) and select one of the following options:
    1. **Log Sequence Number**: Use this option to restore the Archive logs by the Log Sequence Number (LSN). Specify the LSN range to restore the archive logs.
    2. **Time**: Use this option to restore the Archive logs by time frame. Specify a time range between which you need the archive logs.
    3. **System Change Number**: Use this option to restore Archive logs by System Change Number (SCN). Specify the LSN range to restore the archive logs.
5. Click **Next: Recovery Options** and specify the recovery options.
    1. **Restore Archive Logs to**: Specify the path to an existing empty directory. The restored Archive logs will reside in this path.
    2. **Oracle Home**: The ORACLE_HOME value for the host where Archive logs are restored. This is required when recovering the archive logs to an alternate Oracle host.
    3. **Base Directory**: The directory for the restored database. This is required when recovering the archive logs to an alternate Oracle host.

4. **Oracle SID(optional)**: The system identifier (SID) of the Oracle database instance on a specific host. This is required when recovering the archive logs to an alternate Oracle host.

   > **Note:** Restoring archive logs to a database on an alternate host is supported only if that database has the same DB ID as the source DB ID.

5. **Cluster Interface**: Cluster Interface: Select one of the following options to assign the interface group for the recovery job.
   1. **Auto Select**: Toggle on this option if you want to use the same interface group that was selected during the Oracle source registration.
   2. **Interface Group**: Select the interface group to use for the recovery.
6. Click **Recover**.

Test Data Management for Oracle Database

Data and application retention is critical to organizations that seek to prevent data loss and meet security, legal, compliance, disaster-recovery, and business-continuity requirements. Cohesity provides the following solutions as part of the data management with Oracle Instances:

- Clone Oracle Database
- Refresh Clone
- Tear Down Clone

Clone Oracle Database

Cohesity provides the ability to clone Oracle databases (Standalone, RAC, Multi-container Database (CDB), and VCS). The Oracle database is cloned at Cohesity Storage Domains and requires no additional storage from the Oracle Host. Multiple clones from a single source of Oracle Database backup may require little to no additional storage space

in the Cohesity cluster.

Supported Oracle Release for TDM Cloning

Cohesity supports TDM cloning in the following Oracle releases:

Oracle Release	Clone Support for Oracle Database
11gR2	Supported
12cR1	Supported
12cR2	Supported
18c	Supported
19c	Supported

Prerequisites

- Install the latest Cohesity Agent on the system running the Oracle Database.
- Capture Snapshots of the Oracle database to be cloned. To create Snapshots, see create and run a Protection Group.

Clone an Oracle Database

1. Navigate to **Test & Dev** > **Clone** > **Databases**.
2. Search and select the Oracle database. (You can optionally specify the wildcard character *in the search.) You can also optionally narrow the search results by specifying filter criteria.
3. Provide the following information listed in the table, and then click **Clone**.

CLI Option	Description
Task Name	(Optional) Change the default name for the Clone task.
Objects Being Cloned	The Oracle database is being cloned. You can click the change icon to select a different database.
Clone Point	The latest clone point. You click the change icon and select a different clone point. Select the Snapshot to clone. By default, the latest Snapshot is cloned and the selected Restore Point is listed. You can optionally click **Edit** and choose the previous Snapshot to clone.
Oracle Host	The host to which you clone the database.
Oracle Home	The home directory for the cloned database.
Base Directory	The base directory for the cloned database.
Database Name	The name for the cloned database. The maximum number of characters for the database name is eight. The database name should be unique for the given Oracle Host.

CLI Option	Description

Pfile Parameters	The Pfile parameters option allows you to see the target database parameters before the clone job is initiated. The Pfile inherits some source database parameters to the target database and allows you to edit the target database parameters based on your requirement by specifying the additional parameters as a `key=value` pair.
	
	The Pfile lists the following parameters:
	
	• **Default Parameters**: Default parameters of the target database and you can edit the default
	• **Inherited Parameters**: Parameters from the source database and you can edit the inherited parameters, if required.
	• **Restricted Parameters**: You cannot edit the restricted parameters in the Pfile. The following are the restricted parameters:
	  • INSTANCE_NAME
	  • DB_NAME
	  • DB_UNIQUE_NAME
	  • DB_BLOCK_SIZE
	  • MAX_STRING_SIZE
	  • COMPATIBLE
	
	If you edit the following values in the Clone Oracle GUI, You must click the Reset button on the top-right corner to reset the Pfile and get the updated parameter values in the Pfile box.
	  • Oracle Home
	  • Oracle Base
	  • Database Name
	
	Target database will be created with the specified parameters in the Pfile. For example, if `SGA_TARGET` is specified as a part of the pfile parameters, the clone operation will use the value specified in the pfile during the cloning of the source database.

SGA (System Global Area)	The SGA is an area of memory (RAM) allocated to an Oracle Instance when it starts up. Ideally, SGA's value varies from 1024M to 102400M.

CLI Option	Description
Shell Environment	User can configure their shell environment that executes the Cohesity clone workflow using this setting. An example would be, define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your clone target database. For TDE support, the wallet location for your clone target database may depend on a shell variable. Use this to specify your wallet location for cloning a backup taken from a TDE database. For more information, see Clone Oracle TDE Database.
NFS Protocol	Select the NFS protocol version for the clone job. Options include NFSv3 and NFSv4.1. **Note:** NFS protocol is not supported for Windows and AIX databases.
Change DBID	If this option is disabled, the cloned database UUID will be the same as the source database and it will not be displayed on the Source page. Later, If the cloned databases are required to be discovered by Cohesity, then follow the steps below: 1. Start the database in MOUNT mode. 2. Run the following query to get the DBID: "select DBID from GV$database". 3. Change the database DBID using the NID command. 4. Run the following query "select DBID from GV$database" and confirm that the DBID is different from the one we obtained in step 2. 5. Once confirmed, start the database in OPEN mode. 6. Add the database entry in the oratab file. 7. Do a source refresh. 8. The database should be listed there now. If none of the other databases on the Oracle host have DBID the same as that of the cloned database then, skip steps 1 to 5.

You will be redirected to the status details page. Once the clone task that you initiated completes, the status will be displayed as "Success".

You can cancel a clone task for an individual database. Canceling the clone task leaves the database in an unusable state. You can delete this database manually at the Oracle host and then restart the clone task as needed.

Cancel an Oracle Database Clone Task

You can cancel a clone task for an individual database. Canceling the clone task leaves the database in an unusable state. You can delete this database manually at the Oracle host and

then restart the clone task as needed.

1. Navigate to **Test & Dev**.
2. Click a Task Name link in the table to display the details page.
3. Click the **Cancel** button. (The Cancel button is available only if the task is running.)
4. Respond to the confirmation prompt.

Refresh Oracle Database Clone

After you clone the Oracle database on the Cohesity cluster, you can refresh the cloned Oracle database with the latest snapshot, with a selected snapshot, or with a point in time snapshot.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Steps to Refresh a Clone

1. Navigate to **Test & Dev**.
2. Click the cloned task you want to refresh and then click **Refresh** located at the top right of the page. Alternatively, you can hover your mouse over the actions menu located next to the cloned task.
3. To refresh the cloned task with the latest snapshot, select **Refresh to Latest Snapshot**, and then click **Confirm**.

   To refresh the cloned task with an available snapshot, click Refresh to Snapshot, and then perform the following steps:

   a. From the **Recovery Point** column, select the snapshot to which the cloned task must be refreshed.
   b. To refresh the cloned task with a specific time of the selected snapshot, click **Restore to a New Time**, and then use the slide control to adjust the date and time.
   c. Click **Save**.

Tear Down Oracle Database Clone

Deletes the cloned Oracle database from the Oracle Host and remove the database files from the Cohesity Storage Domain.

To delete an Oracle database clone from the Cohesity Storage Domain, perform the following steps:

1. Navigate to **Test & Dev**.
2. Click the cloned task or hover your mouse over the actions menu located next to the

cloned task you want to delete.

3. Click **Tear Down Clone** > **Yes, Tear Down**.

The Oracle database clone will be deleted from the Cohesity cluster Storage Domain.

Clone Oracle TDE Database

Introduction

Transparent Data Encryption (TDE) enables you to encrypt sensitive data that you store in tables and tablespaces. After the data is encrypted, this data is transparently decrypted for authorized users or applications when they access this data. TDE helps protect data stored on media (also called data at rest) in the event that the storage media or data file is stolen.

You can encrypt sensitive data at the column level or the tablespace level. At the column level, you can encrypt data using selected table columns. TDE tablespace encryption enables you to encrypt all of the data that is stored in a tablespace.

Supported Oracle TDE Wallet types

Cohesity does not impose restriction on wallet types used by the source database for the clone workflow. Your source database can use hardware or software keystore, as well as Oracle Key Vault to centralize your wallet keys management.

For restore and clone workflow, Cohesity requires you to set up the appropriate duplicated keystore on your target hosts before using Cohesity restore or clone workflow. In these workflows, only software keystores are supported and keystores must be put into the auto log-in mode before issuing the Cohesity restore or clone workflow.

Clone TDE Database

When you clone a Transparent Data Encryption (TDE) database backup, Oracle performs media recovery on files cloned from Cohesity. To perform media recovery on encrypted blocks, Oracle requires access to the respective TDE keystore.

TDE keystore considerations

- Oracle mandates each database has its own keystore.
- A cloned Oracle database must have its TDE keystore with keys that were used during the previous clone operation.
- Cohesity does not store, manage, or solicit any TDE keystore passwords from users.
- Cohesity requires users to manage the TDE keystore for your source and clone target databases.

This design allows you to control and manage the operation and placement of your TDE keystores to ensure it meets your organization’s security requirements.

To use Cohesity to clone from a backup taken from a TDE database, you must set up a duplicate keystore on the clone target host with a chosen keystore configuration mechanism that is supported by Oracle.

To duplicate a keystore from your source database, you can use one of the following methods as recommended by Oracle:

- Backup and close your source keystore and physically copy the keystore to a new location.
- Use the `ADMINISTER KEY MANAGEMENT CREATE KEYSTORE` command to create an empty keystore, then use the `ADMINISTER KEY MANAGEMENT MERGE KEYSTORE` command to populate your new keystore with the keys from the source keystore.

Oracle supports keystore configuration in the following ways:

- For Oracle versions 18c and later, you can use DB parameters `WALLET_ROOT` and `TDE_CONFIGURATION` to direct your database on how to locate your keystore for your database.
- Use the `ENCRYPTION_WALLET_LOCATION` parameter in `sqlnet.ora` to configure your keystore location. The file path of this `sqlnet` parameter can be made dependent on a shell variable that can uniquely identify your database. For example, the value of the shell variable can be your `DB UNIQUE NAME`.
- Place your keystore in one of the following default locations:

```
$ORACLE_BASE/admin/<db_unique_name>/wallet
```

```
$ORACLE_HOME/admin/<db_unique_name>/wallet
```

> **Note:** Cohesity requires the clone target keystore to be in auto log-in mode. You can create an auto log-in keystore from your duplicated password-protected keystore. For more information, see Creating an Auto-Login or a Local Auto-Login Software Keystore.

If your Oracle database version is 18c or above, and you would like to use the DB parameters `WALLET_ROOT` and `TDE_CONFIGURATION` to configure your cloned target database keystore, you can use the Cohesity version 6.5.1 and later to add the configuration of these DB parameters as part of the cloning workflow.

If you choose to configure your clone target database auto log in keystore using `sqlnet.ora` and you are using a shell variable as part of the file path of the wallet, you have to provide the variable name and its values to be used to the Cohesity clone workflow. For more information, see Clone Oracle Database.

Source database key rotations considerations

- Cohesity does not store, manage, or solicit any TDE keystore passwords from users. Ensure that you manage the TDE keystore for your source and clone target databases. Refresh your duplicated keystore on the cloned target host from time to time if master keys are changed in your source database.

- You can change your source database master keys because a new master key is created, or the keys are rotated.
- A successful Cohesity clone operation from a chosen point-in-time is only possible for a TDE database if your duplicate clone target host keystore has the correct set of master keys from that point-in-time. You do not need the exact set of keys for that point-in-time. It is sufficient if you have those keys as well as more keys than you need in your duplicated keystores. Therefore, the easiest way is to always refresh your duplicate keystore from the latest version of your source keystore.
- Oracle does not provide commands or APIs to delete old master keys from the keystore. Keys can only be added or re-used by key rotation.

Unregister an Oracle Server

You can unregister an Oracle server from the Cohesity cluster. Before you unregister it, you must delete the Protection Groups associated with the Oracle database. Cohesity provides the following solutions as part of unregistering Oracle Instances:

- Delete a Protection Group
- Unregister Oracle Server
- Uninstall Cohesity Agent

Delete a Protection Group

Perform the following steps to delete a Protection Group:

> **Note:** If the protection group is deleted with snapshots, then no restores can be done from that job unless the snapshots are archived.

1. Select **Data Protection > Protection**.
2. Find the Job in the list and in the **Actions** column select **Delete**.

   Alternatively, drill-down on a Job name in the **Job** column of the list. In the Job Details page, click the (**Delete Job**) icon.
3. Choose one of the following:
   - **Delete Job Only:** Select this option to delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Job are not deleted.
   - **Delete Job and Snapshots:** If you select the option to delete the Job and all Snapshots, any local Snapshots captured by this Job and stored on this Cluster are deleted. In addition, any archived Snapshots stored on external targets that were captured by this Job are also deleted. (To selectively delete the Snapshots for specific protection runs without deleting the Job, see Deleting Snapshots or Objects.)

If the current job is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should probably cancel these replication or archive tasks before deleting the job. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Group Run.

If the Job is replicating to a remote cluster, the inactive job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the inactive job.

Unregister Oracle Server

Follow the steps listed below to unregister the Oracle server:

1. Select **Data Protection > Sources**.
2. Navigate the source hierarchy to the Oracle server, mouse over its actions menu, and click **Unregister**.
3. Click **Unregister**.

Uninstall Cohesity Agent

For information about uninstalling the Cohesity agent from an Oracle server (physical or VM), see Install and Manage the Agent on Linux Servers.

Oracle Database Remote Adapter

> **Note:** Cohesity Remote Adapter functionality is not supported if SmartFiles is not licensed on your Cohesity cluster.

The Cohesity remote adapter is used for cases where you need to use custom RMAN scripts (not the one provided by the Cohesity adapter agent). You can schedule agentless Oracle backups with the remote adapter using your own scripts. However, the storage management and backup scheduling is from Cohesity.

Cohesity provides built-in support for Oracle databases. You can either use the Cohesity UI or Remote Adapter Jobs, depending on your preference and requirements for backing up your Oracle database.

The workflow for executing a Protection Group using the Oracle RMAN remote adapter is as follows:

- Create a Cohesity data Protection Group using the Cohesity Remote Adapter. You should add the SSH public key on the Oracle Database server so that the Cohesity cluster can connect to the remote adapter and trigger remote adapter jobs. For the complete set of instructions on how to add the SSH public on the Oracle database, see Add or Edit a Remote Adapter Protection Group.

- Mount the Cohesity Views 'NFS Mount Path' on the Oracle Database Server.
- Set a QoS Policy for the RMAN job.
- Run the RMAN Protection Group.

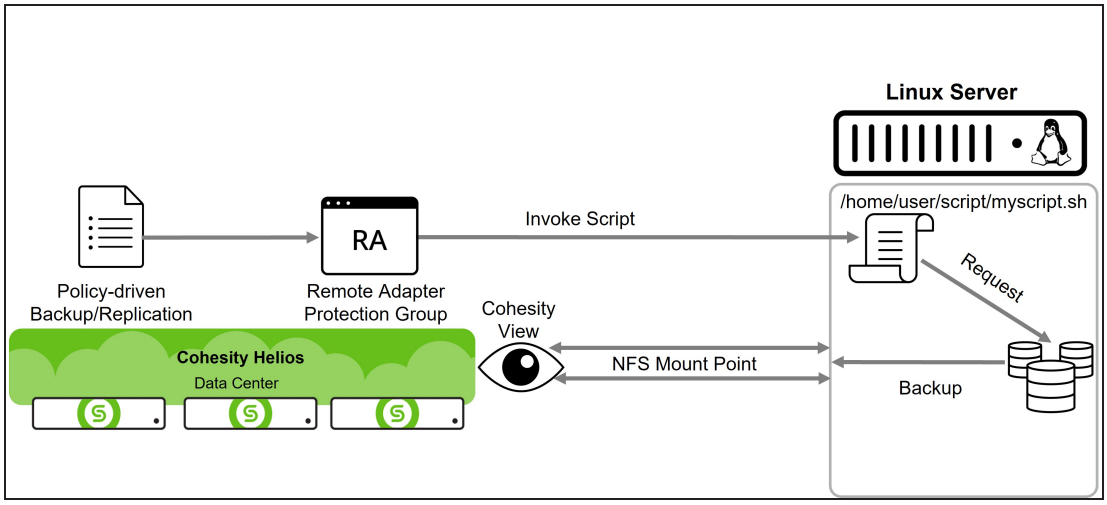

Ports Used for Communication

The following table lists the ports used for external communication when using the Cohesity Remote adapter.

Port	Source	Target	Direction (from node)	Network Protocol	Usage Notes	Type of Traffic
11113	Oracle Host	Cohesity cluster	Incoming	TCP		Backup and Recovery
22	Cohesity cluster	Oracle Host	Outgoing	TCP	SSH client to remotely trigger the scripts on the Oracle Host.	Backup and Recovery

Protecting Oracle Databases using Remote Adapter

When you create a Remote Adapter Protection Group, you will create a View or use an existing View to store the result of the scripts. You must then mount the View on your system and configure your script to write the results to a directory in the mounted View. At the end of the Remote Adapter protection run, the Cohesity cluster captures a Snapshot of the View.

When you create a Remote Adapter Protection Group, you can select an existing Policy, storage domain, and View. You can also create them while creating the job. However, it is easy to create them prior to creating the job, as described in the following topics.

- Create or Edit a Standard Policy
- Create or Edit Storage Domains
- Create a View

Creating Remote Adapter Job

1. You will create an RMAN data Protection Group leveraging a pre-existing RMAN script. Start by navigating to the Protection Group landing page through the Protection menu option and then click **Remote Adapter**. Then choose Remote Adapter as the job type. This will start the Protection Group setup workflow with each step of the workflow identified at the top of the screen. For the complete set of instructions on how to create a remote adapter Protection Group, see Add or Edit a Remote Adapter Protection Group.
2. From the New Remote Adapter Protection Group page, either select **New Policy** to create a new policy or select an existing "Protection Policy". With the existing policy, you can ignore the default "Protection Schedule and Options" as this schedule does not apply for Oracle backups. For more information, see Create or Edit a Standard Policy.
3. In the next **Specify Script Settings** page, specify the Oracle database host information, SSH keys, and script locations. Ensure the SSH key is verified to allow remote login and execution of scripts. This can be verified by running a test script that echoes an output. For more details, see Creating a NFS Export on the Cohesity Storage.
4. Specify a descriptive job name, an appropriate Cohesity storage domain, and a Cohesity View. For more details, see Create or Edit Views.

A summary of the settings appears. Review the summary to ensure that all the settings are correct.

Remote Adapter Considerations for Oracle Database

- **Limited Support**: Remote Adapter is not supported for Windows.
- **Register Recovery Scripts**: There is no single location to register recovery scripts when using the remote adapter.
- **Tracking**: You cannot track and report all the recovery tasks and alerts from Cohesity UI when using the remote adapter.
- **Data Read Column**: The Data Read column in the Protection Group Run details page shows the total size of the Oracle Databases protected by the job. This is because of an RMAN limitation.

RMAN NFS Target for Oracle Database

Cohesity provides an efficient integration option with Oracle's Recovery Manager (RMAN) using the NFS mount points. You can present Cohesity as an NFS target for native RMAN backups using the NFS mounts.

To store your RMAN backups on the Cohesity cluster, you need to create a Cohesity View. The View allows you to set the appropriate options for deduplication, compression, and QoS Policies and lets you use NFS mounts to assign channels to every node in a Cohesity cluster.

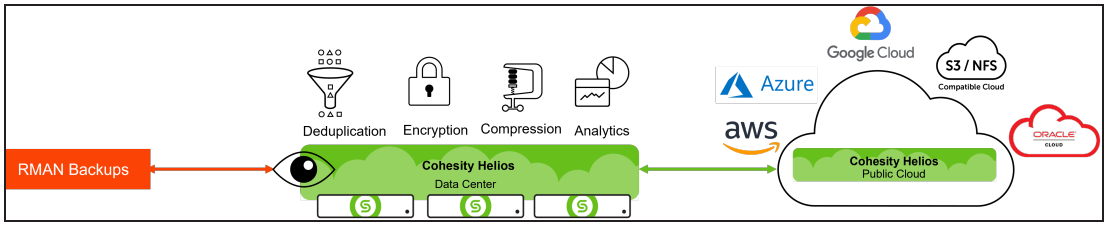

Each node in a Cohesity cluster has a Virtual IP (VIP). You can use these VIPs in your RMAN scripts when you allocate channels to every node in a Cohesity cluster. By assigning separate channels to every node, you are assigning parallel resources of each node to share the workload and thus making your backups faster.

You can use multiple backup RMAN channels in the backup operation and map several channels to the same VIP to increase the parallel processing. A combination of multiple channels and multiple mount points (defined using VIPs) will yield faster backup performance.

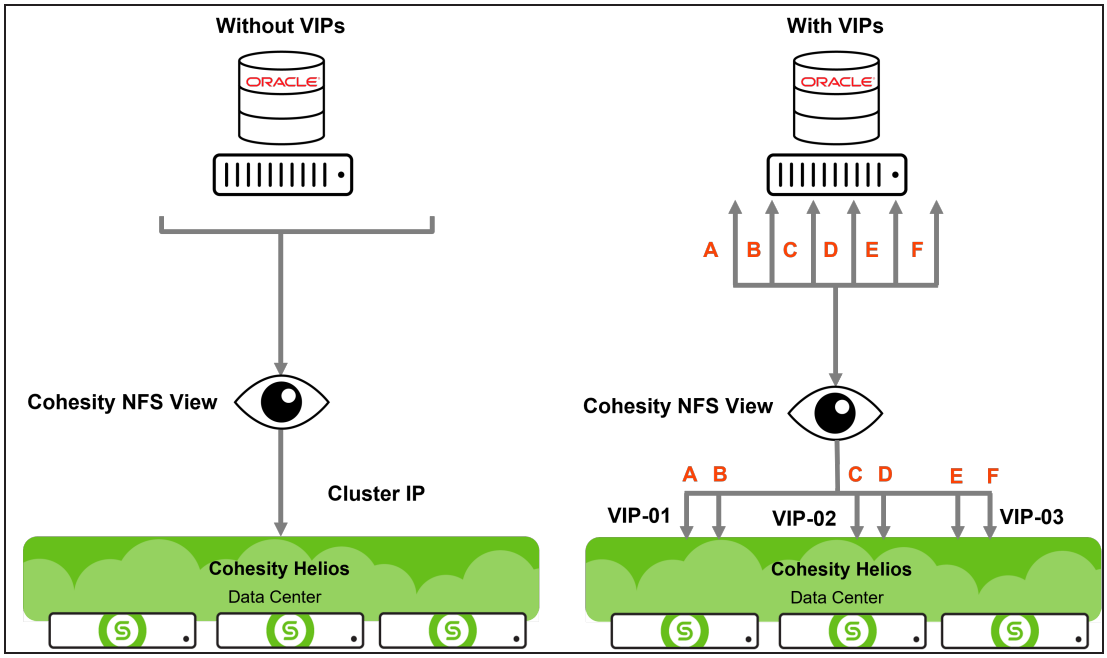

> **Tip:** For best performance, create a mount point for each node in your Cohesity cluster using that node's unique VIP.

Allocating too many or too few channels can degrade the backup performance. Cohesity recommends that you determine the optimal combination of channels and mount points based on the backup operation on the Oracle side to give you the best performance. The number of channels should be in multiples based on your cluster node count.

For example, if your Cohesity cluster has 3 nodes, then use channels in multiples of 3 (3, 6, 9, 12).

Create NFS Export on Cohesity Storage for Oracle Database

In order to create an NFS export from the Cohesity Data Cloud, follow the steps listed below:

1. Create a storage domain to define the policy settings on the exported shares. For more information, see Create or Edit Storage Domains.
2. Set the data management policies within the storage domain, such as enabling data reduction functionality, the type of data deduplication, and whether inline or post process.
3. Create a View and enter the View Name. The View name is also used to create the exported share name. In the Advanced Settings, select a specific protocol such as NFS. In this example, NFS is the protocol choice for Linux and Unix. For more details, Create a View.
4. Associate this view with the Storage Domain you created. This View inherits the policies associated with the Storage Domain.
5. Allowlist the IP subnet of the Oracle database server. Add the IP address or IP subnet of the Oracle database servers that are allowed to access exported NFS share in the Cohesity View. For more information, see Add Subnets to Allowlists.

Backup and restore for NFS are completely driven at the user or customer end. If you choose Cohesity Data Cloud as NFS target, it just acts as storage.

RMAN NFS Best Practices for Oracle Database

For Oracle RMAN backup configured with NFS Export, the capacity savings are extremely dependent on how RMAN is used. Oracle datafiles contain data blocks, which have block headers at the start of the block. These contain numbers, which are unique for each datafile, making deduplication impossible at the database block size. Also, the end of each block contains a tail check section that features a number generated using data such as the SCN. Even if the block is divided into two, the second half would offer a limited opportunity for deduplication while the first half would offer none. To maximize deduplication efficiency when using RMAN, understand these settings and best practices:

- **FILESPERSET=1**

  This setting suppresses Oracle multiplexing of more than one data file into a backup set. When you suppress Oracle multiplexing, Oracle creates the backup set identically each time the backup runs and maximizes the dedupe rate.

```
BACKUP
FILESPERSET=1
%BACKUP_TYPE%
FORMAT 'bk_u%%u_s%%s_p%%p_t%%t'
DATABASE;
```

- **Compression/Encryption**

  Both compression and encryption must be disabled in RMAN to see effective deduplication. However, compression should be enabled when backing up transaction logs and control files because they don't see good deduplication rates.

```
RMAN> CONFIGURE BACKUP OPTIMIZATION OFF;
RMAN> CONFIGURE ENCRYPTION FOR DATABASE OFF;
```

- **Mount options for NFS client**

  Depending on the Oracle Server OS, different NFS client mount options are required for successful backup and recovery of Oracle database. The following table lists the NFS mount options for the supported Oracle Server OS:

<table>
<tr><th>OS</th><th>Deployment Type</th><th>Version</th><th>Default Mount Options</th><th>Full/Incremental Backup</th><th>Log Backup</th><th>Overwrite Restore</th><th>Alternate Restore</th><th>Clone</th></tr>
<tr><td>AIX</td><td rowspan="3">RAC</td><td>12c</td><td rowspan="4">bg,hard,intr,rsize=32768,wsize=32768,vers=3,proto=tcp,sec=sys,rw,acdirmax=30</td><td></td><td></td><td rowspan="3">noac</td><td></td><td>noac</td></tr>
<tr><td></td><td>11g</td><td>noac</td><td></td><td></td><td>llock</td></tr>
<tr><td></td><td rowspan="2">Standalone</td><td>12c</td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td></td><td>11g</td><td></td><td></td><td></td><td></td><td></td></tr>
</table>

OS	Deployment Type	Version	Default Mount Options	Full/Incremental Backup	Log Backup	Overwrite Restore	Alternate Restore	Clone
Linux	RAC	12c	rw,relatime,vers=3,rsize=1048576,wsize=1048576,namlen=255,hard,proto=tcp,timeo=600,retrans=2,sec=sys,mountvers=3,mountproto=tcp,local_lock=none					noac llock
		11g		noac				
	Standalone	12c						llock
		11g						

- **QoS settings for NFS view**

  NFS view should be created with TestAndDev High QoS settings. This will provide higher priority when Cohesity runs other tasks such as VM backup jobs.

  The following table lists the QoS policies in the increasing order of QoS and helps in I/O tuning.

QoS Policy	Description
Backup Target Low	Use this policy for non-latency sensitive workloads such as backups, which keep many I/Os outstanding. Data is written to HDD and has higher latency compared to other QoS policies writing to a SSD.
Backup Target High	The same as **Backup Target Low**, except that the priority for processing workloads with this policy is higher than workloads with Backup Target Low when there is contention.
Backup Target SSD	Use this policy for workloads such as backups, which keep many I/Os outstanding. However, in this case, Cohesity sends both sequential and random I/Os to SSD. The latency is lower than other Backup Target policies.
Backup Target CommVault	Use this policy to intelligently detect and exclude application-specific markers to achieve better deduplication when CommVault backup application is writing to a Cohesity View. Data is written to SSD and has the same priority and latency as **TestAndDev High**.

QoS Policy	Description
TestAndDev Low	Use this policy for workloads that require lower I/O latency or do not keep many I/Os outstanding, as the I/Os are given higher priority compared to other QoS policies. Data is written to SSD and has lower latency compared to Backup Target SSD.
TestAndDev High	The same as **TestAndDev Low**, except that the I/Os with this QoS policy are given higher priority compared to I/Os with TestAndDev Low when there is contention.

- **Load balancing RMAN jobs**

  To get the best performance with Oracle RMAN, the load from backup and restore jobs should be distributed between all the Cohesity nodes in the cluster. To load balance work between the nodes, create multiple mount points with different VIP of the Cohesity Data Cloud on the Oracle Server as RMAN target.

  For example, assume a view named `nfs_view` created (using Step 4) by selecting the suitable QoS at Cohesity Domain. The same view is used to create multiple mount points with different VIP of the Cohesity Data Cloud.

  This is an example how the mount is performed:

```
create directory structure

                                    mkdir -p /orcl1
mkdir -p /orcl2
mkdir -p /orcl3
mkdir -p /orcl4

The following Linux NFS settings are the recommended options for the
Cohesity as they provide the best performance:

noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,hard,intr,nolock
```

  Example: Mount VIPs using the options above:

```
sudo mount -t nfs -o
noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,hard,intr,nolock
192.10.2.10:/nfs_view /orcl1
sudo mount -t nfs -o
noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,hard,intr,nolock
192.10.2.11:/nfs_view /orcl2
sudo mount -t nfs -o
noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,hard,intr,nolock
192.10.2.12:/nfs_view /orcl3
```

```
sudo mount -t nfs -o
noatime,vers=3,proto=tcp,rsize=1048576,wsize=1048576,hard,intr,nolock
192.10.2.13:/nfs_view /orcl4
```

> **Note:** Ensure that you update the `/etc/fstab` file with above mount options to make the changes persistent across server reboots.

RMAN Full backup script looks similar to this example:

```
CONFIGURE DEFAULT DEVICE TYPE TO DISK;
CONFIGURE DEVICE TYPE DISK PARALLELISM 4 BACKUP TYPE TO BACKUPSET;

run {
allocate channel t1 type disk FORMAT '/orcl1/%d_%T_%U';
allocate channel t2 type disk FORMAT '/orcl2/%d_%T_%U';
allocate channel t3 type disk FORMAT '/orcl3/%d_%T_%U';
allocate channel t4 type disk FORMAT '/orcl4/%d_%T_%U';

sql 'alter system archive log current';
backup spfile tag 'spfile';
CONFIGURE CONTROLFILE AUTOBACKUP OFF;
backup current controlfile tag 'controlfile_1';
backup incremental level 1 for recover of copy with tag 'Full_1'
database plus archivelog delete input;
RECOVER COPY OF DATABASE WITH TAG 'Full_1';
}
```

- **Archive Logs and Control Files**

  By default, RMAN interleaves archive (transaction) logs and control files with database content. This defeats deduplication because it changes the alignment of the data. When used with deduplication, these additional files should be written to separate streams.

- The Cohesity cluster will work with Oracle compression and encryption functionality. If encryption and compression are turned on, make sure to disable data deduplication on the Cohesity View.
- When the native deduplication built-in to the Cohesity Data Cloud is not used, the alternate storage and time-saving option is native Oracle compression with the multi-channel target. For storage, RMAN compression shows around 50% reduction and also shows a good backup time reduction. However, these are dependent highly on Oracle RMAN performance and tuning. This may cause extra load to the Oracle DB server.

For more information, see Cohesity and Oracle RMAN Guide.

Cohesity Oracle SBT Plug-in

The Cohesity Oracle SBT Plug-in allows you to use Cohesity as target storage for backing up Oracle databases.

Cohesity Oracle SBT Plug-in protects the Oracle databases using the Oracle Recovery Manager (RMAN) script. Cohesity natively integrates with Oracle Recovery Manager (RMAN) to provide application-consistent backup and recovery for Oracle databases including single instance or Real Application Clusters (RAC).

Check the **Databases** section in the Supported Software topic for a complete list of supported Oracle versions for Cohesity SBT Plug-in and the supported Cohesity cluster versions.

Backup and Restore using Cohesity Oracle SBT Plug-In

> **Note:**
> - Cohesity recommends you schedule periodic full backups for the Oracle database protection using Oracle SBT Plug-In.
> - For Oracle 19c, the Oracle SBT plug-in supports backing up of containerized databases (CDB/PDB).

To back up a database:

1. Create views as the target where the databases need to be backed up.
2. Add the Oracle host(s) to the allowlist of the view or global allowlist.
3. Allocate SBT channels in the RMAN scripts to include the Cohesity-specific parameters.

   For example:

   ```
   'ALLOCATE CHANNEL c1 DEVICE TYPE SBT_TAPE PARMS 'SBT_
   LIBRARY=/opt/cohesity_sbt/lib/libsbt_linux_x86_64.so, SBT_PARMS=(data_
   view=/SbtView1/db1,vips=192.0.2.1)'
   ```

4. Schedule the backup with Cohesity Remote Adapter or external scheduler by running the script from the previous step.

To restore a database:

1. Select the view or snapshot where the backup data exists for a specific Database for restore.
2. Expose the snapshot to a view if necessary and optionally mount to a target source if the customer wants to see the backup pieces on the view.
3. Write RMAN restore scripts to include the Cohesity-specific parameters same as the

backup.

4. Run the script manually to perform the restore.

The following table lists the ports used for communication between the Cohesity cluster and the Oracle server. Please open the port based on your choice of protocol. The default protocol is gRPC and the TCP port is 11113.

Port	Source	Target	Direction (from node)	Network Protocol	Usage Notes	Type of Traffic
11113	Oracle Server	Cohesity cluster	Incoming	TCP	gRPC used by Oracle SBT	Backup and Recovery
11117	Oracle Server	Cohesity cluster	Incoming	TCP	Secure gRPC used by Oracle SBT	Backup and Recovery
11111	Oracle Server	Cohesity cluster	Incoming	TCP	SunRPC used by Oracle SBT (Linux only)	Backup and Recovery

Installing Cohesity Oracle SBT Plug-in

The Cohesity Oracle SBT Plug-in can be installed on the Linux, AIX, and Windows platforms.

- Install Cohesity Oracle SBT Plug-in (For Linux)
- Install Cohesity Oracle SBT Plug-in (For AIX)
- Install Cohesity Oracle SBT Plug-in (For Windows)

> **Note:**
> - Cohesity recommends using the latest download package available on the downloads portal.
> - The latest version of the SBT plugin is compatible with all the previous versions. You can replace the plugin to upgrade to the latest version without changing the RMAN scripts.
> - If Packet Inspection is enabled on the firewall between the Oracle host and the Cohesity cluster, the backup performance with gRPC is affected. To resolve this issue, either turn off the Packet Inspection on the firewall or use SunRPC or secure gRPC instead of gRPC.

Install Cohesity Oracle SBT Plug-in (For Linux)

1. Download the Cohesity Oracle SBT Plug-in for Linux from the Cohesity Package Download page.
2. Place the `libsbt_linux_x86_64.so` file in any directory of the Oracle server. For the Oracle RAC setup, the file should be put in the same directory in all the Oracle RAC nodes.

Install Cohesity Oracle SBT Plug-in (For AIX)

There are two ways to download and install Cohesity Oracle SBT Plug-in for AIX.

Install from the single library file

1. Download the Cohesity Oracle SBT Plug-in for AIX from the Cohesity Package Download page.

   > **Note:** Perform this step for each Oracle Database that is protected by Cohesity.

2. Place the `libsbt_aix_powerpc.so` file in any directory of the Oracle server. For the Oracle RAC setup, the file should be put in the same directory in all the Oracle RAC nodes.

Install from the installer

1. Download the Cohesity Oracle SBT Plug-in installer for AIX.
   1. Download the installer `cohesity_plugin_sbt_release-<yyyymmdd>_<commit_id>_aix_powerpc_installer` from the Cohesity Package Download page.
2. Add the executable permission to the downloaded installer file then execute it.
   1. Install to default location `/opt/cohesity/plugins/sbt`.

      Run the following command with root privilege.

      ```
      sh <installer_name>
      ```

   2. Optionally, You can install to a customized location by defining the installation directory to which you have the write permission.

      ```
      sh <installer_name> -- -d <installation_dir>
      ```

3. For the Oracle RAC setup, repeat the steps above for all the nodes.

> **Note:** To use Cohesity SBT Plugin on AIX 7.3, the AIX largesend attribute of the network interface should be disabled.

To disable the `largesend`, run the following command. This configuration will be lost after rebooting.

```
ifconfig <interface> -largesend
```

To persist the configure after rebooting, run the following command.

```
chdev -l <interface> -a mtu_bypass=off
```

To check if the `largsend` is enabled, run the following command.

```
ifconfig <interface> |grep LARGESEND
```

Install Cohesity Oracle SBT Plug-in (For Windows)

1. Download the Cohesity Oracle SBT Plug-in library for Windows Server 2012 and later from the Cohesity Package Download page. The RPC library shared on the Cohesity Downloads page is applicable for Cohesity versions 6.5.1 and later. Alternatively, you can install the Cohesity Oracle SBT Plug-in library by selecting the **Install SBT** option checkbox while installing the Cohesity agent.

   > **Note:** There are two files for the library. `Cohesity_sbt.dll` and `Cohesity_sbt.exe`.

2. Place the two files in the same directory of the Oracle server. Ensure the directory is accessible by Oracle services and the directory path does not contain space.
3. For Oracle RAC setup, repeat the steps above for all the nodes.

Creating Cohesity View and Configuring RMAN script

1. Create views on the Cohesity cluster and add the following to the allowlist of the view:
   a. The Oracle host IP for a DB server. If there are multiple IPs on the Oracle server, suggest adding all the IPs to the allowlist.

      The read-only option does not apply when creating a view for use with sbt, as the sbt plugin can modify files in the view regardless of this setting.
2. Edit the RMAN script to include the following information:
   a. Parameters that indicate backups will be sent to a tape library and the library's location:
      i. Device Type is SBT_TAPE
      ii. **For Linux**

      ```
      SBT_LIBRARY=<library_directory>/libsbt_linux-x86_64.so
      ```

      iii. **For AIX**

      ```
      SBT_LIBRARY=<library_directory>/libsbt_aix_powerpc.so
      ```

iv. **For Windows**

```
SBT_LIBRARY=<library_directory>\cohesity_sbt.dll
```

b. Provide the required Cohesity specific details:

i. **mount_path**: <any of the Cohesity cluster VIPs>:/<view name>/<path>, for example, `192.0.2.0:/SbtView1/db1`.

ii. **vips**: One of the VIPs configured for the Cohesity cluster, for example, `192.0.2.1`.

iii. For IPv6 only environment, add `sbt_use_ipv6=true` in the `SBT_PARMS` list.

iv. To enable enable security certificates, perform the following steps.

a. Make sure the TCP port 11117 is reachable from the Oracle host.

b. log in to the Cohesity cluster backend and run the following commands:

```
/home/cohesity/bin/client_tls_cert_generator_exec -
validity_days <days> -ca_type <cohesity_ca|self_signed>
```

By default, Cohesity will generate a certificate with 7300 days validity of type `cohesity_ca`.

If you want to fall back to generate a self-signed certificate, specify the option `-ca_type self_signed` when running the command.

> **Note:**
> - SBT library will now try to renew the certificate automatically if the cert type is `cohesity_ca` and it will expire in 60 days. To disable the auto renewal, specify `sbt_certificate_auto_renew=false` in the `SBT_PARMS` when configure or allocate RMAN channels.
> - Run this command from the **cluster host shell**, not from the default Secure Shell session (the prompt that shows restricted). If you only need iris_cli commands, remain in the Secure Shell.

c. Copy the generated certificate file `/tmp/cert.cfg` to Oracle host.

d. Add the `sbt_certificate_file=<path of the certificate file>` to the SBT_PARMS.

> **Note:** The IPv6 and TLS features are supported on the cluster version to be 6.5.1 and later and only support Linux platforms.

3. Run the RMAN script either by Cohesity Remote Adapter jobs or customer external schedulers. If an external scheduler is used to invoke the RMAN script, make sure to take snapshots of the target view after each successful backup run. This can be configured by calling the Cohesity REST API in the script or by scheduling the view protection job in the policy.
4. Cohesity SBT plug-in supports single node single channel, single node multi-channel, and multi-node multi-channel based on the VIPS settings in the RMAN scripts. You must use the same `mount_path` for all the channels.

Linux Tools

For information on debug tools, see SBT Linux Tools.

SBT Channel Parameters

The following table lists the Oracle SBT Channel Parameters.

Name	Value	Mandatory/Optional	Notes	Example
vips	string	Mandatory	List of host names or VIPs of the Cohesity cluster. Separated by","for IPv4 or"|"for IPv6 if there are multiple vips are provided.	vips=sac01-broadwell28-p1-vip.eng.cohesity.com vips=10.14.53.42 Note: vips can be specified with FQDN
mount_path or data_view	string	Mandatory	mount_path or path of the view.	mount_path=sac01-broadwell28-p1-vip.eng.cohesity.com:/sbt_view/db1 data_view=sbt_view/db1

Name	Value	Mandatory/Optional	Notes	Example
sbt_ certificate_file	string	Optional	The full path on the Oracle host of the certificate file copied from the Cohesity cluster Note: Specifying this option will imply using secure gRPC	sbt_certificate_ file=/opt/cohesity_sbt/cert.cfg
sbt_use_grpc	bool	Optional	Set this value to false to use SunRPC, this option only applies for Linux platform.	sbt_use_grpc=false
disable_ source_side_ dedup	bool	Optional	Set this value to true to disable source side dedup for full/incremental backups. Archivelog backups will not do source side dedup regardless of this value. You can separately enable or disable source side dedup for your full and incremental backup, if you use different script or RMAN run block to run different types of backup.	disable_source_side_ dedup=true
sbt_use_ipv6	bool	Optional	Enable ipv6 for SBT	sbt_use_ipv6=true

Source-Side Deduplication for Oracle Database

The source-side deduplication feature is built into the Cohesity SBT plugin. As long as you use the SBT plugin, the source-side deduplication will be enabled automatically by default.

The Cohesity source-side deduplication provides storage-efficient deduplication that enhances your Oracle environment with faster backups and reduced network traffic between Oracle servers and the Cohesity.

This functionality is enabled by default on the Cohesity cluster. It shifts the deduplication processes to the Oracle server and reduces CPU consumption on the Cohesity allowing further scaling to support more workloads. Additionally, the source-side deduplication

plugin Cohesity integrates with the Platform's built-in deduplication and minimizes data traffic by moving only the new data blocks.

Advantages of the Source-Side Deduplication Plugin

- Faster backups
- Reduce network traffic between Oracle servers and Cohesity
- Enables more frequent backups resulting in higher RPOs
- Offloading deduplication reduces the CPU consumption and enables the Cohesity to run more backup workload sessions.

The following diagram illustrates how source-side deduplication reduces CPU consumption by offloading the deduplication to the source client.

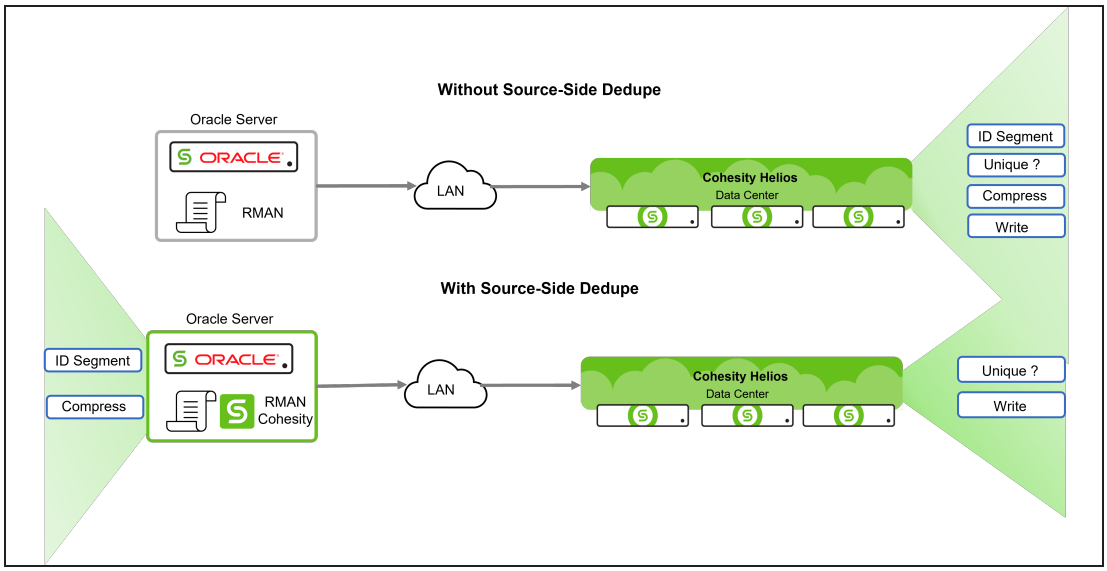

Enable Source-Side Deduplication for Oracle

For a complete set of instructions on how to enable source-side deduplication for Oracle when using the Remote Adapter RMAN script, see Install Cohesity Oracle SBT Plug-in (For Linux).

Cohesity Oracle SBT Plug-In for ZDLRA

The Cohesity SBT plug-in for ZDLRA will be used in the customer data center to facilitate ZDLRA copy jobs to the Cohesity. These copy jobs are queued and executed by ZDLRA according to the configured copy job policies.

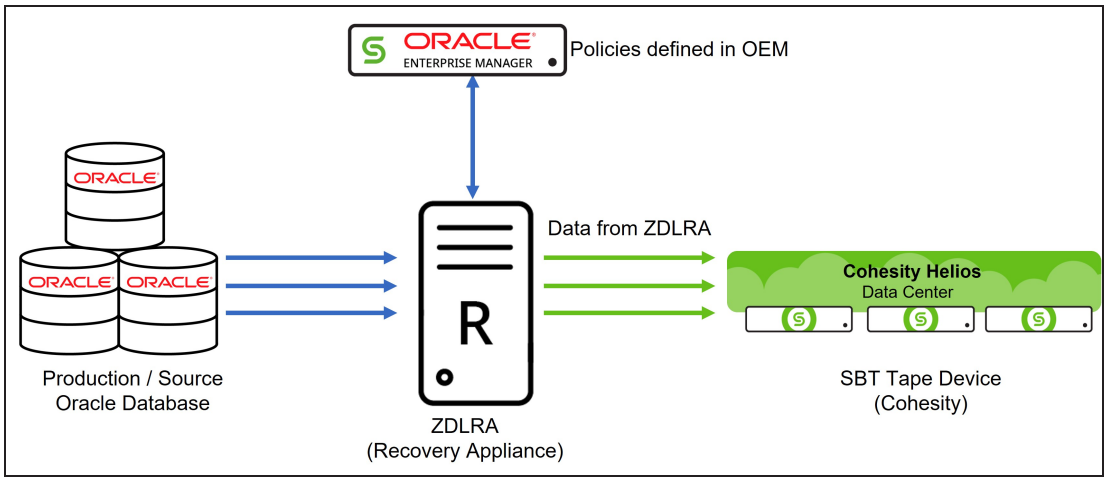

Cohesity acts as a tape device for the ZDLRA appliance and provides cost-efficiency by offloading backups to Cohesity for long-term storage. Backups are stored on Cohesity cluster and snapshotted into immutable views to guarantee data protection. Cohesity does its cataloging of backups by collecting metadata of backups and stores this metadata in the Catalog View.

Once the backup is on the Cohesity cluster, protection jobs can be configured to replicate **backups** to other Cohesity clusters, or archived to Cloud storage. Cohesity provides flexible restore options through ZDLRA or directly from Cohesity without ZDLRA. Refer to the Restore section for more details.

Roles

There are 3 roles in a ZDLRA environment with a Cohesity cluster. Cohesity administrator, ZDLRA administrator, and Oracle DBA.

- **Cohesity Administrator**: Cohesity administrator manages the views and protection jobs for ZDLRA to copy the backup data to and restore from the Cohesity cluster.
- **ZDLRA Administrator**: ZDLRA administrator configures ZDLRA copy jobs to copy the backup data to the Cohesity views created by the Cohesity administrator.
- **Oracle DBA**: Oracle DBA performs the restore from ZDLRA or directly from Cohesity with the view information provided by the Cohesity administrator.

Cohesity Views for ZDLRA

A view provides a storage location in a storage domain on the Cohesity cluster. These views can only be accessed by machines specified in the allowlisted subnets. A view that is created and used specifically as a storage location for the ZDLRA copy job is only accessible via the Cohesity SBT driver that is installed on the ZDLRA.

Two distinct types of views are required for ZDLRA backups.

Data View

a. A Data View is used to store the backup data that is copied to the Cohesity cluster through ZDLRA.
b. Multiple Data Views can be configured for one ZDLRA. Each Data View can be used for more than one database.
c. The Cohesity administrator will provide the ZDLRA administrator the information about the view that is required to configure the ZDLRA copy job attributes.
d. The data view name needs to be configured as the value of the parameter `data_view` in the ZDLRA copy job Attribute Set.

Catalog View

a. A Catalog View is used for storing the metadata of the backup data that is copied to the Cohesity Data Views through ZDLRA.
b. This metadata is used by the Cohesity SBT code when ZDLRA issues a restore or crosscheck operation to the Cohesity SBT.
c. The Cohesity administrator will provide the ZDLRA administrator with information about the Catalog View. The ZDLRA administrator will need to know this information to correctly install the Cohesity ZDLRA Package on ZDLRA.
d. The Catalog View name needs to be configured as the value of the parameter a part of `catalog_view` in the ZDLRA copy job Attribute Set and in the configuration file `/opt/cohesity_sbt/config/sbt_config.cfg` on ZDLRA hosts.

Data Flow

When the copy job runs on the ZDLRA, via the SBT library, it will write data to two separate locations on the Cohesity. Each backup piece received from ZDLRA will be written to the Data View and its corresponding metadata will be written to the Catalog View with a filename that is the same name as the backup piece filename.

> **Note:** The file names on both views are the same; however, the contents of the files are different. Do not configure the Data View and Catalog View as the same view, or else your copy job execution will fail.

Prerequisites

This section lists the prerequisites for ZDLRA and Oracle host.

ZDLRA Requirements

The following files are required on the ZDLRA:

1. **SBT Library File**: The SBT library file is a single .so file extracted from the ZDLRA Package.

2. **Configuration File**: The configuration file is a text file created by the ZDLRA administrator on ZDLRA hosts to set the catalog view for ZDLRA copy jobs.
3. **ZDLRA Utility File**: The ZDLRA utility file is an executable that can be run on ZDLRA to list the total number of files and the total size of the data copied to Cohesity for each DB in a particular time range. This file is an optional tool. The ZDLRA backup and restore are not affected without this tool.

Oracle Host Requirements

The following files are required on the Oracle host when restoring directly from Cohesity:

1. **RMAN Script**: Oracle DBA needs to write a restore script based on the sample RMAN script provided by the Cohesity administrator.
2. **SBT Library File**: The SBT library file is a single .so file extracted from the ZDLRA Package.

Installation and Configuration

This section lists the configuration steps for both the Cohesity cluster and the ZDLRA appliance.

- Configuring Cohesity Cluster
- Configuring ZDLRA

Configuring Cohesity Cluster

The steps listed in this section are performed by the Cohesity Administrator.

1. Verify the Cohesity Cluster Version.

   The cluster must be running with the locked view and apiKey features. Cluster versions 6.5.1 or above support these features. If your cluster runs on a special version, contact Cohesity support to confirm if it supports ZDLRA. If your cluster is running on another version, contact Cohesity Support to confirm if it supports ZDLRA.
2. Configure View Indexing.
3. Create a Cohesity Protection Policy.

   Create Cohesity Protection Policies to be used for the view protection jobs for both the catalog view and all the data views. The retention period of the data views should be equal or greater than the retention period of the corresponding SBT copy job. For more information, see Protect Oracle Database.

> **Note:** The schedule of the Cohesity protection policy for data views will be ignored because the protection jobs are triggered by ZDLRA instead of Cohesity. Cohesity will expire the data based on the retention of the Cohesity protection policy, it won't affect the data and metadata on the ZDLRA end. If the SBT retention on the ZDLRA side is less than the one on the Cohesity side, the catalog entries in the Cohesity catalog view will be deleted when ZDLRA expires the data but the data in the Cohesity snapshots of the data view will NOT be deleted.

4. Install the ZDLRA Package on a Linux host which will be used to manage the Cohesity views and jobs for ZDLRA/SBT backups.
   a. Download the ZDLRA package cohesity_zdlra.tar.gz from the Cohesity Downloads Site to the management Linux host.
   b. Run the following command to extract the file:

      ```
      tar xvf cohesity_zdlra.tar.gz
      ```

   c. Run the following command to install the installer in the extracted directory cohesity_zdlra:

      `sudo ./cohesity_zdlra_sbt_linux_x86_64_installer` for default installation which will install the package to `/opt/cohesity_sbt` or `./cohesity_zdlra_sbt_linux_x86_64_installer -- -d <install_dir>` for customized installation.

   d. Refer to the "Cohesity Installer Details" section in the Appendix for more information about what the package installs.
5. Create a Cohesity Data Security user.

   Log in to Cohesity UI as an admin. Create a user with a Data Security role. The user will be used to manage the views for the ZDLRA copy and restore jobs. For more information, see Manage Roles.

   Optionally, you can generate an API key by running the following CLI command in the Cohesity CLI:

   ```
   <install_dir>/bin/view_cli generate --host <cluster_name_or_ip> --
   username <user_with_admin_role> --password <password> --apikeyusername
   <data_security_username> --apikeyname <choose_your_own_keyname>
   ```

   For more information on arguments description, see ZDLRA CLI Commands.
6. Create Cohesity Views and Protection Jobs.

There are two ways to create the Cohesity views and view protection jobs for ZDLRA copy jobs. One is to run the script from the Cohesity admin's Linux host. The other is to create them from Cohesity UI.

a. Create Cohesity Views and Protection Jobs with Script

   The ZDLRA package installer contains a convenient script that will create the Cohesity View and Protection Job at the same time. The script is called `cohesity_create_view_for_zdlra.sh`. It is located in `/opt/cohesity_sbt/bin` for default installation or `<install_dir>/bin` for customized installation.

   The parameters that can be used with this script are explained in detail in the Appendix section of this document. The `-u` and `-p` options are not needed if the API key (generated in step 5) is supplied. Each ZDLRA can have one Catalog View and one or more Data Views.

   i. As a user with the Data Security Role, create a **Catalog View** by running the following command:

   ```
   ./cohesity_create_view_for_zdlra.sh -V <zdlra_catalog_view_name> -I <cluster_ip> -P <policy_name> -W <allowlist/subnet_mask> -s <storage_domain> [-u <user_with_data_security_role> -p <password_for_user> -J <job_name>]
   ```

   ii. As a user with the Data Security Role, create a **Data View** by running the following command:

   ```
   ./cohesity_create_view_for_zdlra.sh -V <zdlra_data_view_name> -I <cluster_ip> -P <policy_name> -W <allowlist/subnet_mask> -s <storage_domain> -C <zdlra_catalog_view_name> [-u <user_with_data_security_role> -p <password_for_user> -J <job_name>]
   ```

   > **Note:** A data view can only be created after the catalog view is created and must set the catalog view name with the `-C` option.

   The views created using the script are locked views that cannot be deleted or updated from Cohesity UI.

   The allowlist should be a valid IPv4 address and the subnet_mask should be an integer between 0 and 32. If the value of `subnet_mask` is 32, it is for the single IP, in this case, the "/32" can be omitted. If multiple IPs need to be allowlisted, separate them with ",".

   For example,

```
./cohesity_create_view_for_zdlra.sh -V zdlra_catalog -I
cluster01 -P zdlra_365 -W 169.22.61.10,169.22.61.12/24 -s
DefaultStorageDomain
```

b. Create Cohesity Views and Protection Jobs from Cohesity UI.

   This section lists steps the to create the views and jobs for ZDLRA from Cohesity UI. Please note the views created from Cohesity UI are not locked views so that they can be updated from Cohesity UI.

   i. Create the catalog view as a regular view

      I. From the Cohesity UI, navigate to **File Services** > **Views**, then click the **+** button in the upper-right corner and click **New View** from the dropdown list.

      II. Provide the view attributes and click the **Create** button.

   ii. Create a view protection job for the catalog view

      I. From the Cohesity UI, navigate to **Data Protection** > **Protection**, then click the **+** button in the upper-right corner and click the **View** menu from the dropdown list.

      II. Select the catalog view name and click the **Save Selection** button.

      III. Provide the job attributes and click the **Protect** button.

   iii. Create the data view from the ZDLRA view template

      I. From the Cohesity UI, navigate to **File Services** > **Views**, then click the **+** button in the upper-right corner and click the **New View from Template** menu from the dropdown list.

      II. Select **ZDLRA** from the dropdown list.

      III. Provide the view attributes and click the **Create** button.

   iv. Create a view protection job for the data view

      I. From the Cohesity UI, navigate to **Data Protection** > **Protection**, then click the **+** button in the upper-right corner and click the **View** menu from the dropdown list.

      II. Select the data view name and click the **Save Selection** button.

      III. In addition to the job name and policy, configure the catalog view by clicking the **Additional Settings** dropdown and select **SBT** as the client type of **Externally Triggered Settings** and set the catalog name to the **Catalog View** option.

      IV. Click the **Protect** button to create the job.

7. Share the following Information for the ZDLRA administrator:

a. **Data View name and VIPs**

   Data View name and VIPs are used in the configuration of the environment variables for the ZDLRA copy job Attribute Set on the ZDLRA.

b. **Catalog View name and VIPs**

   The ZDLRA administrator can set the catalog view name and VIPs in the configuration file `/opt/cohesity_sbt/config/sbt_config.cfg` on ZDLRA hosts.

Configuring ZDLRA

The steps listed in this section are performed by the ZDLRA Administrator.

1. Install the Cohesity ZDLRA Plug-in.

   Optionally, obtain the file `libst_linux_x86_64.so` from the Cohesity administrator and place it in a directory on saved on the Oracle databse server.

   For example:

   ```
   /opt/cohesity_sbt/lib.
   ```

2. Create the configuration file.

   Manually create a configuration file called `/opt/cohesity_sbt/config/sbt_config.cfg` and add the following entry to this file for each Cohesity cluster:

   ```
   Cohesity_cluster : {
   catalog_view: "<catalog_view_name>"
   catalog_vips: "<cluster_VIPs>"
   }
   ```

   - The `catalog_vips` could be one or more IP addresses or hostnames of the Cohesity cluster VIPs. If more than one VIPs are specified, separate them with `:`. For example:

     ```
     catalog_vips:"10.2.100.178:10.2.100.180:10.2.100.182"
     ```

   - By default, the SBT log will be written to the `sbtio.log` of ZDLRA, you may add the optional parameter `log_dir` to the configuration file to redirect the logs to any directory of your choice. For example:

     ```
     log_dir: "/opt/cohesit_sbt/logs"
     ```

   - By default, the `SBT log_level` is `1` which will generate normal information such as the file names, speed, parameters, etc. This value can be set to `0` in the configuration file to allow only error messages to be generated. For example:

     ```
     log_level: 0    #log_level: 0 for error, 1 for info, 2 for debug
     ```

The `log_level` can also be configured in the copy job attribute by adding `log_level=n` in the ENV of PARMS. In this case, the `log_level` only applies to the specific copy job.

A sample `sbt_config.cfg` file is given below:

```
[oracle@rack110dbadm01 ~]$ cat /opt/cohesity_sbt/config/sbt_
config.cfg
log_level: 1

cohesity_cluster: {
  catalog_view: "catalog"
  catalog_vips: "10.14.25.96:10.14.25.97:10.14.25.98"
}

cohesity_cluster: {
  catalog_view: "catalog"
  catalog_vips: "10.14.53.50:10.14.53.103:10.14.53.111"
}

cohesity_cluster: {
  catalog_view: "catalog_400d"
  catalog_vips: "10.2.200.64"
}
```

This file should only be created on the ZDLRA, not on the Oracle hosts.

3. Create the ZDLRA SBT library using `DBMS_RA.CREATE_SBT_LIBRARY`.

   For example:

```
DBMS_RA.CREATE_SBT_LIBRARY(
            LIB_NAM            => 'COHESITY_SBT',
            DRIVE                      => 24,
            RESTORE_DRIVES     => 8,
PARMS                  => 'SBT_LIBRARY=/opt/cohesity_sbt/lib/libsbt_linux_x86_64.so'
         );
```

4. Create the ZDLRA copy job Attribute Set. Coordinate with your Cohesity administrator on the data views and catalog views to be used for copy jobs.

   For example:

```
'ENV=(data_view=<data_view_name>,vips=<cluster_vips>,catalog_
view=<catalog_view_name>,is_zdlra_copy_job=true,<log_level=<log_
level>>)'
    );
```

The VIPs could be one or more IP addresses of the Cohesity cluster VIPs. If more than one VIPs are specified, separate them with ":".

> **Note:** All the VIPs configured in the ZDLRA copy job attributes must exist in the corresponding cluster entry in the configuration file `/opt/cohesity_sbt/config/sbt_config.cfg` on the ZDLRA hosts.

By default, the `log_level` is `1`, which will generate normal information such as the file names, speed, parameters, etc. This value can be set to `0` to allow only errors to be generated or set to `2` for debug logs.

Example 1: Single VIP

```
 DBMS_RA.CREATE_SBT_ATTRIBUTE_SET (
     LIB_NAME          => 'COHESITY_SBT',
     ATTRIBUTE_SET_NAME        => 'ZDLRA_TEST',
     STREAMS                   => 16,
            PARAMS                      => 'ENV=(data_view=zdlra_
35,vips=10.2.32.123,catalog_view=zdlra_400d,is_zdlra_copy_job=true)'

    );
```

Example 2: Multiple VIPs

```
DBMS_RA.CREATE_SBT_ATTRIBUTE_SET (
     LIB_NAME          => 'COHESITY_SBT',
     ATTRIBUTE_SET_NAME        => 'ZDLRA_TEST',
     STREAMS                   => 16,
            PARAMS                      => 'ENV=(data_view=zdlra_
100,vips=10.2.32.123:10.2.32.125:10.2.32.127:10.2.32.129,catalog_
view=zdlra_400d,is_zdlra_copy_job=true)'
```

Example 3: Multiple VIPs and disable regular logs

```
DBMS_RA.CREATE_SBT_ATTRIBUTE_SET (
     LIB_NAME          => 'COHESITY_SBT',
     ATTRIBUTE_SET_NAME        => 'ZDLRA_TEST',
     STREAMS                   => 16,
            PARAMS                      => 'ENV=(data_view=zdlra_
100,vips=10.2.32.123:10.2.32.125:10.2.32.127:10.2.32.129,catalog_
view=zdlra_400d,is_zdlra_copy_job=true,log_level=0)'
```

5. Create a ZDLRA copy job template for each copy job.

   For example:

```
DBMS_RA.CREATE_SBT_JOB_TEMPLATE (
    TEMPLATE_NAME                    => 'DB123_FULL',
    PROTECTION_POLICY_NAME     => '90_DAYS',
    ATTRIBUTE_SET_NAME               => 'ZDLRA_TEST',
    BACKUP_TTYPE                     => 'FULL, ARCH');
```

6. Queue the copy job using a `Cron job` or the `DBMS_SCHEDULER` package. Use the following `DBMS_RA` call to queue the SBT copy job to ZDLRA:

```
DBMS_RA.QUEUE_SBT_BACKUP_TAST('DB123_FULL');
```

Restore

This section lists the steps for restoring data from ZDLRA and direct restore without ZDLRA, using the Cohesity cluster.

- Restore through ZDLRA from the Local Cohesity Cluster
- Restore through ZDLRA from Replicated Cohesity Cluster
- Direct Restore without ZDLRA
- Direct Restore from the Remote Cluster without ZDLRA
- Restore from the Remote Cluster through ZDLRA
- Direct Restore from Archival without ZDLRA
- Restore from Archival through ZDLRA

Restore through ZDLRA from the Local Cohesity Cluster

Connect to the ZDLRA catalog to run the restore. ZDLRA will automatically read the backup data from the Cohesity cluster if they were copied to Cohesity and the local copy has been deleted.

Restore through ZDLRA from Replicated Cohesity Cluster

Cohesity administrator updates the whitelists of both replicated data view and catalog view to the IP addresses of the ZDLRA nodes.

ZDLRA administrator adds the replicated Cohesity cluster VIPs and replicated catalog view name in the configuration file /opt/cohesity_sbt/config/sbt_config.cfg on the ZDLRA hosts.

Connect to the ZDLRA catalog to run the restore. ZDLRA will automatically read the backup data from the replicated Cohesity cluster if they were copied to Cohesity and the local copy has been deleted.

Note: The files already expired on the local Cohesity cluster can not be restored by this way. To restore the files older than the retention of the policy on the local cluster, please follow the direct restore from Cohesity instructions.

Direct Restore without ZDLRA

The steps listed in this section are performed by the Oracle DBA and Cohesity Administrator.

1. Oracle DBA requests the `cohesity_zdlra.tar.gz` tar file from the Cohesity administrator and extract the Cohesity SBT plug-in `libsbt_linux_x86_64.so` file.

   The `libsbt_linux_x86_64.so` file is then placed in the directory of their choice. For example:

   ```
   /opt/cohesity_sbt/lib.
   ```

2. Oracle DBA provides the following information to the Cohesity administrator:

   i. The path of the Cohesity SBT plugin. For example:

   ```
   /opt/cohesity_sbt/lib/libsbt_linux_x86_64.so.
   ```

   ii. The DB name to be restored.

   iii. The time range of the files to be used for restore. The time is determined by the time point when the files were copied to the Cohesity cluster.

   iv. The IP address of the database server to which the files need to be restored. If the database server has more than one IP address, provide the IP address used to connect the Cohesity cluster. If you are not sure which one is used for connecting to the Cohesity cluster, all the IP addresses can be provided.

3. The Cohesity administrator generates an RMAN re-catalog script by running the `cohesity_prepare_restore.sh` script and with the information provided by the Oracle DBA.

   The following example generates the re-catalog script for database `testdb` for backup pieces within the time range `2020/04/14-23:53:01` to `2020/07/16-19:48:43`.

   ```
   ./cohesity_prepare_restore.sh -I 10.2.148.135 -S /opt/cohesity_
   sbt/lib/libsbt_linux_x86_64.so -D testdb -W 0.0.0.0/0 -s 2020/04/14-
   23:53:01 -e 2020/07/16-19:48:43 -J jobname
   ```

4. The Cohesity administrator provides the sample script generated by the `cohesity_prepare_restore.sh` to the Oracle DBA.

   The sample script can be found from /`tmp/cohesity_sbt_<timestamp>/cohesity_catalog_sbt.<dbname>.rman`.

5. Oracle DBA performs the re-catalog and restores the database based on the sample script provided by the Cohesity administrator by following steps below:

   i. Oracle DBA opens the RMAN script to read the comments, which contains the list of available backup pieces, each with an associated time stamp and file size, that matches your input to the script.

   ii. Identify the control file name that needs to be restored and write an RMAN script to restore the control file according to the SBT parameters in the sample script.

iii. Use RMAN to directly restore from individual backup pieces, or run the sample script in RMAN to re-catalog all the backup pieces if the target Oracle database is already mounted or open.

iv. Write an RMAN script to restore the database according to the SBT parameters in the sample script.

6. After the restore/recovery is finished, the Cohesity administrator cleans the cloned views and temp files with the script `cohesity_end_restore.sh`.

   For example:

   ```
   ./cohesity_end_restore.sh -I 10.2.148.135 -d /tmp/cohesity_sbt_
   20200611_100424
   ```

Direct Restore from the Remote Cluster without ZDLRA

To restore the files replicated to the remote cluster, follow the same procedure of Restore through ZDLRA from the Local Cohesity Cluster from the local cluster but set the cluster IP to the remote cluster.

Restore from the Remote Cluster through ZDLRA

To restore the files replicated to the remote cluster, follow the following steps.

1. The Cohesity administrator runs the script `cohesity_prepare_restore.sh` as described in the section Direct restore without ZDLRA and adds the option -Z.

   For example:

   ```
   ./cohesity_prepare_restore.sh -Z -I 10.2.148.135 -S /opt/cohesity_
   sbt/lib/libsbt_linux_x86_64.so -D testdb -W 0.0.0.0/0 -s 2020/04/14-
   23:53:01 -e 2020/07/16-19:48:43 -J jobname
   ```

2. ZDLRA administrator adds the remote cluster vip and the catalog view created in the previous step to the `/opt/cohesity_sbt/config/sbt_config.cfg` files on all the ZDLRA nodes.
3. Oracle DBA runs the restore on the target Oracle host against ZDLRA.
4. After the restore, the Cohesity administrator runs the script `cohesity_end_restore.sh` to clean up the cloned views and the ZDLRA administrator deletes the catalog view entries from the files `/opt/cohesity_sbt/config/sbt_config.cfg` on all the ZDLRA nodes.

Direct Restore from Archival without ZDLRA

To restore the files archived to the cloud, follow the same procedure of Direct restore without ZDLRA from the local cluster and add the `-A <archival_target>` option to the `cohesity_prepare_restore.sh` command.

Restore from Archival through ZDLRA

To restore the files archived to the cloud through ZDLRA, follow the same procedure of Direct Restore from Archival without ZDLRA from the local cluster and add the `-z` option to the `cohesity_prepare_restore.sh` command.

ZDLRA and SBT restore visibility

The ZDLRA and SBT restore visibility feature allows you to check the reading activities against externally triggered views through the Cohesity SBT plugin. Typically the reading happens when restoring Oracle databases from the data view snapshots or clones of them.

To enable the ZDLRA and SBT restore visibility feature, you have to configure the File Service Audit Logs on the cluster and upgrade both the SBT plugin and the ZDLRA package to the latest version.

Enable or Disable File Services Audit Logs

By default, the File Services Audit setting is disabled. You should enable the File Services Audit option for Cohesity to start logging the Cohesity SBT read events.

You can disable or enable it back from the Settings tab of the Audit Logs page. When you disable the option, the old logs are retained until the retention period ends.

To enable or disable File Services Audit, complete the following:

1. Navigate to **System** > **Audit Logs** > **Log Settings**.
2. In the **Log Settings** tab, turn on/off the **File Services Audit** toggle to enable or disable.

After enabling the **File Services Audit**, the **Log Retention Period** audit information option appears. Use this option to set the audit logs retention period.

Set Log Retention Period for File Services Audit Logs

You can set the retention period for file services audit logs on the Cohesity cluster.

> **Note:** The default retention period is 90 days. You can increase or decrease the renewal period. Once the retention period expires, the audit logs are deleted automatically.

To set a retention period for File Services Audit Logs, complete the following:

1. Navigate to **System** > **Audit Logs** > **Log Settings**.
2. In the **File Services Audit** section of the **Log Settings** tab, click the icon for the **Log Retention Period** option, type the desired number, and select the retention period (Days, Weeks, Months, or Years).

The Log Retention Period is displayed in days as Cohesity converts weeks, months, or years into days.

3. Click **Save**.

Add IP Allowlist for File Services Audit Logs

The host that installed the ZDLRA package has to be included in the IP allowlist of the File Service Audit Logs to print the SBT read report. By default, every client that is part of the Global Allowlist set at the cluster level can access the File Services audit logs. You can change access permission by the Override Allowlist option to allow or deny access to a subnet. To override the global IP allowlist set at the cluster level and add a subnet to the allowlist for file service audit:

1. Navigate to **System** > **Audit Logs** > **Log Settings**.
2. In the **IP Allowlist** section of the **Log Settings** page, turn on the **Override Global IP Allowlist** toggle.
3. Click **Add**.
4. In the **Add Allowlist** dialog, complete the following:
   a. Type the **Subnet IP** address in CIDR format.

      Example: Subnet in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).

   b. The **NFS Permissions**, **SMB Permissions**, or **S3 Permissions** can be any value. (Read/Write, Read Only, or Disabled).

   c. Optionally, type a description.

   d. Click **Add**.
5. The added allowlist IP subnet details appear in the File Services Audit section of the page.

Print the SBT Read Report

The utility sbt_report delivered in the ZDLRA package can be used to print the SBT reading report. For example:

```
/opt/cohesity_sbt/bin/sbt_report -vips=CohesityCluster

START_TIME            END_TIME              DB_NAME               CLIENT_IP
          STATUS                FILE_SIZE            READ_SIZE
 FILE_NAME
2024/01/21-07:28:01   2024/01/21-07:28:09   CDB
10.14.53.99           Success               2301362176           2301362176
           132h4flq_1_1
2024/01/21-07:28:16   2024/01/21-07:28:18   CDB
10.14.53.99           Success               617873408            617873408
           142h4fmj_1_1
```

```
2024/01/21-07:28:19   2024/01/21-07:28:21   CDB
10.14.53.99           Success               563871744             563871744
            152h4fn3_1_1
```

Refer to Appendix for the detailed usage of the sbt_report.

Enabling SSL/TLS for ZDLRA

Cohesity SBT supports SSL/TLS by transferring the encrypted data through secure gRPC.

- Enable SSL/TLS for ZDLRA Copy Job
- Enable SSL/TLS for Restore from Cohesity

Enable SSL/TLS for ZDLRA Copy Job

To enable the SSL/TLS for ZDLRA copy job, follow the steps below:

1. Log in to the Cohesity cluster and run the following command to generate the certificate file for ZDLRA.

   ```
   ./client_tls_cert_generator_exec -cert_name <certificate_name> -
   validity_days <expire_days> -output_file <certificate_file_name>
   ```

   For Example,

   ```
   [cohesity@zdlra-cluster-bqkp82500972-node-2 bin]$ client_tls_cert_
   generator_exec -cert_name zdlra1 -validity_days 3650 -output_file
   /tmp/zdlra1.cert
   I0811 09:10:16.374521 15601 cluster_config_service.go:246] Starting
   cluster config service...
   I0811 09:10:16.374548 15601 cluster_config_service.go:255] Cluster
   config service successfully started
   I0811 09:10:16.374678 15601 cluster_config_service.go:298] Starting
   cluster config watcher.......
   I0811 09:10:16.374726 15601 config_watcher.go:73] Loading cluster
   config...
   I0811 09:10:16.382387 15601 config_watcher.go:207] Config change
   notification sent
   I0811 09:10:16.382427 15601 config_watcher.go:84] Load cluster config
   done
   I0811 09:10:16.382426 15601 cluster_config_service.go:313] Config
   change notification received
   I0811 09:10:16.382477 15601 config_watcher.go:155] Watch for cluster
   config change, version=583
   I0811 09:10:16.382551 15601 nexus_proxy_service.go:669] Request to
   watch for cluster config change version=583
   I0811 09:10:16.557071 15601 certificate_utils.go:80] Generated new
   client TLS certificate configuration.
   ```

```
I0811 09:10:16.558456 15601 cluster_config_service.go:402] Received
request to stop cluster config service
I0811 09:10:16.558804 15601 cluster_config_service.go:283] Received
notification to stop the config service
I0811 09:10:16.558815 15601 config_watcher.go:114] Received
notification to stop the config watch
I0811 09:10:16.558824 15601 cluster_config_service.go:414] Cluster
config service successfully stopped
I0811 09:10:16.558828 15601 client_tls_cert_generator.go:145] Saved
3650 days validity client TLS certification to /tmp/zdlra1.cert
successfully.
```

2. Copy the generated certificate file to all the ZDLRA nodes and add the full path of the generated certificate file in the `field sbt_certificate_file` attribute of the `/opt/cohesity_sbt/config/sbt_config.cfg` on all the ZDLRA nodes.

   For Example,

```
$ cat /opt/cohesity_sbt/config/sbt_config.cfg
cohesity_cluster: {
catalog_view: "catalog"
catalog_vips: "10.14.25.96:10.14.25.97:10.14.25.98"

sbt_certificate_file: "/opt/cohesity_sbt/config/zdlra1.cert"
}
```

Enable SSL/TLS for Restore from Cohesity

To enable the SSL/TLS for restore from Cohesity, follow the steps below:

1. Log in to the Cohesity cluster and run the following command to generate the certificate file for the Oracle host which the DB to be restored to.

```
./client_tls_cert_generator_exec -cert_name <certificate_name> -
validity_days <expire_days> -output_file <certificate_file_name>
```

   For Example,

```
[cohesity@zdlra-cluster-bqkp82500972-node-2 ~]$ client_tls_cert_
generator_exec -cert_name zdlra1 -validity_days 3650 -output_file
/tmp/oracle1.cert
I0811 11:58:44.593153 23873 cluster_config_service.go:246] Starting
cluster config service...
I0811 11:58:44.593205 23873 cluster_config_service.go:255] Cluster
config service successfully started
I0811 11:58:44.593275 23873 cluster_config_service.go:298] Starting
cluster config watcher.......
I0811 11:58:44.593316 23873 config_watcher.go:73] Loading cluster
```

```
config...
I0811 11:58:44.600283 23873 config_watcher.go:207] Config change
notification sent
I0811 11:58:44.600299 23873 config_watcher.go:84] Load cluster config
done
I0811 11:58:44.600314 23873 cluster_config_service.go:313] Config
change notification received
I0811 11:58:44.600354 23873 config_watcher.go:155] Watch for cluster
config change, version=583
I0811 11:58:44.600439 23873 nexus_proxy_service.go:669] Request to
watch for cluster config change version=583
I0811 11:58:44.866526 23873 certificate_utils.go:80] Generated new
client TLS certificate configuration.
I0811 11:58:44.867805 23873 cluster_config_service.go:402] Received
request to stop cluster config service
I0811 11:58:44.867827 23873 cluster_config_service.go:283] Received
notification to stop the config service
I0811 11:58:44.867845 23873 config_watcher.go:114] Received
notification to stop the config watch
I0811 11:58:44.867862 23873 cluster_config_service.go:414] Cluster
config service successfully stopped
I0811 11:58:44.867868 23873 client_tls_cert_generator.go:145] Saved
3650 days validity client TLS certification to /tmp/oracle1.cert
successfully.
```

2. Copy the generated certificate file to the Oracle hosts and add the `-c <full path of the certificate file>` when running the script `cohesity_prepare_restore.sh`.

   For Example,

```
$ ./cohesity_prepare_restore.sh -I konacluster -S /opt/cohesity_
sbt/lib/libsbt_linux_x86_64.so -D OL7312C -W 10.0.0.0/1 -c
/opt/cohesity_sbt/config/oracle1.cert
/tmp/cohesity_sbt_20210811_150508/cohesity_catalog_sbt.OL7312C.rman is
the sample RMAN script to catalog your requested backup for the
database OL7312C
```

Non-ZDLRA Backup and Restore

Backup Oracle Database without ZDLRA

The following steps can be used to back up Oracle databases to immutable views on a Cohesity cluster.

1. Oracle DBA provides the IP addresses of the Oracle servers which are used to connect to the Cohesity cluster to the Cohesity administrator.

2. The Cohesity administrator creates the data view and catalog view for the database to be backed up and adds the IP addresses to both the data view and catalog view for the database and provides the vips, data view name and catalog view name to the Oracle DBA.
3. Oracle DBA puts the SBT library to any directory and writes RMAN scripts with vips, data_view and catalog_view in the parameters of the RMAN channels in the following format.

```
SBT_PARMS=(vips=<vips of the cluster>, data_view=<data view name>,
catalog_view=<catalog view name>)
```

or

```
ENV=(vips=<vips of the cluster>, data_view=<data view name>, catalog_
view=<catalog view name>)
```

For example

```
run {
ALLOCATE CHANNEL c1 DEVICE TYPE sbt PARMS   'SBT_LIBRARY=/opt/cohesity_
sbt/lib/libsbt_linux_x86_64.so, SBT_PARMS=(data_
view=data,vips=10.14.53.50,catalog_view=catalog)';
ALLOCATE CHANNEL c2 DEVICE TYPE sbt PARMS  'SBT_LIBRARY=/opt/cohesity_
sbt/lib/libsbt_linux_x86_64.so, SBT_PARMS=(data_
view=data,vips=10.14.53.103,catalog_view=catalog)';
backup database;
}
```

> **Note:**
> If restoring the database to a host other than the one used for backup, the Cohesity administrator needs to whitelist the IPs of the destination host to both the catalog and data views.
>
> For example:
>
> ```
> ./view_cli update -host test_cluster -whitelist
> 10.14.56.80,10.14.58.241 -view catalog
> ```
>
> ```
> ./view_cli update -host test_cluster -whitelist
> 10.14.56.80,10.14.58.241 -view data
> ```

Restore Oracle Database from Non-ZDLRA Backups

There are two scenarios to restore an Oracle database from the Non-ZDLRA backups. One is that the RMAN catalog is available. The other is that the RMAN catalog is not available.

Restore Database with RMAN Catalog from Local Cluster to Alternate Host

If restoring the database to a host other than the one used for backup, the Cohesity administrator needs to add the IPs of the destination hosts to the catalog view whitelist either with the command `view_cli` or through Cohesity REST API.

After the restore target hosts are whitelisted for the catalog view, add the parameter `sbt_allow_copy_catalog_whitelist=true` to the SBT channel of the RMAN restore script.

For example:

```
run { ALLOCATE CHANNEL c1 DEVICE TYPE sbt PARMS 'SBT_LIBRARY=/opt/cohesity_sbt/lib/libsbt_linux_x86_64.so, SBT_PARMS=(data_view=data,catalog_view=catalog,vips=vmdp,sbt_allow_copy_catalog_whitelist=true)';  restore database; }
```

The format of the view_cli command is `` `view_cli update -host <cluster_name> -view <view_name> -whitelist <new_whitelists>` ``.

For example:

```
./view_cli update -host test_cluster -whitelist 10.14.56.80,10.14.58.241 -view catalog
```

The `view_cli` update command will overwrite the existing whitelist of the views. To keep the existing whitelist, first run the `view_cli listwhitelist` command to get the existing whitelist and then append the new IPs to the existing one for the update.

For example:

```
./view_cli listWhitelist -host test_cluster -view catalog 10.0.0.1/32./view_cli update -host test_cluster -whitelist 10.0.0.1/32,10.14.56.80,10.14.58.241 -view catalog
```

Restore Database Without RMAN Catalog or from Replicated Cluster or Archival

If the RMAN catalog is not available, a re-catalog is needed for the restore. The procedure is the same as the one of direct restore from Cohesity for the ZDLRA backups.

SBT Alerts on Backup and Restore Failures

Starting with Cohesity cluster version 7.3 and SBT release 20250626_0548b3bd, the cluster generates alerts for failed SBT backup and restore operations if:

- The catalog view is used, and
- The SBT requests can reach the Backup Manager service through the I/O Operation service when the error occurs.

If the SBT plug-in successfully connects to the Backup Manager service at the start of a backup or restore but later loses the connection due to a network or I/O Operation error, the Backup Manager will raise a timeout alert after the configured threshold (30 minutes by default).

> **Note:** The timeout alert may not be generated if the Backup Manager service is restarted during the timeout period.

Appendix

This section lists the Cohesity ZDLRA installer details, ZDLRA CLI commands, and ZDLRA scripts used for creating ZDLRA views and ZDLRA restore.

- Cohesity ZDLRA Installer Details
- ZDLRA CLI Commands
- ZDLRA Scripts and Supported Arguments

Cohesity ZDLRA Installer Details

The name of the installer package is `cohesity_zdlra_sbt_linux_x86_64_installer`.

To run the installer, use the following command:

```
cohesity_zdlra_sbt_linux_x86_64_installer -- -d <install dir> -u <user> -g
<group>
```

The installer creates the following directory structure:

<install dir>	/opt/cohesity_sbt if no -d option is provided		
<install dir>	/bin	cohesity_create_view_for_zdlra.sh	Used for creating the Cohesity Views and Cohesity Protection Jobs by Cohesity admin (see details below)
		cohesity_end_restore.sh	Used for cleaning up the cloned views and temp files by the Oracle DBA after performing the direct restore (see details below)
		cohesity_prepare_restore.sh	Used for cloning the data views and generating the template file with the file list to be cataloged for restore (see details below)
		sbt_list	Used for listing the files that exist in a given View (see details below)
		view_cli	Used for manipulating Views (see details below)
		sbt_report	Used for reporting the restore activities (see details below)
	/config		Used for setting the catalog view name and VIPs in the configuration file `sbt_config.cfg` on ZDLRA hosts.
	/lib	libsbt_linux_x86_64.so	SBT library
	/logs		Used for writing logs to the `sbtio.log` of ZDLRA. You may add the optional parameter `log_dir` to the configuration file to redirect the logs to any directory of your choice.

ZDLRA CLI Commands

- **sbt_list**

  This CLI command is used for listing the files that exist in a given View. It is called internally by `cohesity_prepare_restore.sh`. For example:

  ```
  sbt_list --view=<data view name> --vips=<vips> --db_name=<db name> [--start_time=<yyyy/mm/dd-HH24:MI:SS>] [--end_time=<yyyy/mm/dd-HH24:MI:SS>]
  ```

- **sbt_report**

  This utility is used for printing the report of the SBT read activities. The syntax is:

  ```
  sbt_report --vips=<vips> [--start_time=<yyyy/mm/dd-HH24:MI:SS>] [--end_time=<yyyy/mm/dd-HH24:MI:SS>][--db_names=<db_names>]
  ```

```
[--sbt_certificate_file=<certificate_file_name>]
[--client_ips=<client_ips>]
[--status=[InProgress|Success|Failure|TimeOut]]
[--sbt_use_ipv6=[true|false]] [--log_dir=<dir>]
```

The arguments of the command sbt_report are listed as below. Please note the vips option is mandatory, all other options are optional.

Command	Description	Arguments	Argument Description
sbt_report	Print the report of the SBT reading events on a Cohesity cluster.	--vips <string>	The Cohesity cluster name or VIP
		--start_time <string>	The start time of the reading events to be reported.
		--end_time <string>	The end time of the reading events to be reported.
		--db_names <string>	The comma separated database names of the reading events to be reported.
		--sbt_ certificate_file <string>	The full path of the SBT certificate file if secure gRPC is enabled.
		--client_ips <string>	The comma separated client IPs of the reading events to be reported.
		--status <string>	The comma separated status of the reading events to be reported.
		--sort_by <string>	The column name used to sort the results.
		--show_units <bool>	If displaying the size units such as KB, MB, GB etc.
		--sbt_use_ipv6 <bool>	If using IPv6 to connect to Cohesity cluster.
		--log_dir <string>	The full path of the sbt_report logs.

- **view_cli**

  This CLI command is used for manipulating Views. The currently supported commands are:

Command	Description	Arguments	Argument Description
create	Create a View	--host <string>	The Cohesity cluster name or VIP
		--username <string>	*The username for the Cohesity UI
		--password <string>	**The password for the Cohesity UI
		--view <string>	The name of the view to create
		--viewbox <string>	The name of the Storage Domain. The JPMC viewbox is “DefaultStorageDomain”
		--whitelist <string>	The allowlist in the format of x.x.x.x/x,x.x.x.x/x
update	Update a View allowlist	--host <string>	The Cohesity cluster name or VIP
		--username <string>	**The username for the Cohesity UI
		--password <string>	**The password for the Cohesity UI
		--view <string>	The name of the view to update
		--whitelist <string>	The allowlist in the format of x.x.x.x/x,x.x.x.x/x
		--isDataview <true/false>	True to change the view to a data view; False to change the view to a regular view.
		--isLocked <string>	True to lock the view; False to unlock the view.

Command	Description	Arguments	Argument Description
clone	Clone a View	--host <string>	The Cohesity cluster name or VIP
		--username <string>	*The username for the Cohesity UI
		--password <string>	**The password for the Cohesity UI
		--source <string>	The name of the view to be cloned
		--target <string>	The name of the cloned view
		--whitelist <string>	The allowlist in the format of x.x.x.x/x,x.x.x.x/x
protect	Create a Protection Job	--host <string>	The Cohesity cluster name or VIP
		--username <string>	*The username for the Cohesity UI
		--password <string>	*The password for the Cohesity UI
		--view <string>	The name of the view to be protected
		--policy <string>	The name of the Protection Policy
delete	Delete a View	--host <string>	The Cohesity cluster name or VIP
		--username <string>	*The username for the Cohesity UI
		--password <string>	*The password for the Cohesity UI
		--view <string>	The name of the View to delete
search	Search Views based on DB name	--host <string>	The Cohesity cluster name or VIP
		--username <string>	*The username for the Cohesity UI
		--password <string>	*The password for the Cohesity UI
		--dbname <string>	The DB name

Command	Description	Arguments	Argument Description
generate	Generate API key	--host <string>	The Cohesity cluster name or VIP
		--username <string>	*The username for the Cohesity UI
		--password <string>	*The password for the Cohesity UI
		--apikeyname	The name of the API key to create
		--apikeyusername	**The username for which to create the key
		--apikeydir	***The location in which to create the API Key file

*Username and password are optional if the API Key file is present.

**If not specified, the value from `--username <string>` will be used.

***If not specified, the file will be created in `${HOME}`.

Examples

- **create**

```
./view_cli create --host cluster01.xxx.com --username dsuser --password
abcd1234 --view zdlra_data --viewbox DefaultStorageDomain --whitelist
10.5.68.178,10.5.68.179
```

- **update**

```
./view_cli update -host cluster01.xxx.com -username dsuser -password
abcd1234 -view zdlra_data -whitelist 10.5.68.179,10.5.68.180
```

- **clone**

```
./view_cli clone --host cluster01.xxx.com --source zdlra_data --
whitelist 10.5.68.118,10.5.68.119 --target zdlra_data_clone --username
db_clone --password abcd1234
```

- **protect**

```
./view_cli protect --host cluster01.xxx.com --policy "90_day_policy" --
username cohesity-admin --password admin --view zdlra_data
```

- **generate**

```
./view_cli generate --host cluster01.xxx.com --username cohesity-admin
--password admin --apikeyusername oracle_zdlra --apikeydir /tmp --
apikeyname oracle_zdlra_20200703
```

ZDLRA Scripts and Supported Arguments

This section lists the scripts and arguments used for creating ZDLRA views and preparing to restore directly from the Cohesity cluster.

`cohesity_create_view_for_zdlra.sh`

Arguments that can be used with the `cohesity_create_view_for_zdlra.sh script`:

Argument	Description
-V	View name to create
-I	Cohesity Cluster IP
-P	Cohesity Cluster Policy name
-W	Allowed IP address/netmask bits separate by commas if multiple
-u	Cohesity username (optional - if not provided will use API key)
-p	Cohesity password (optional - if not provided will use API key)
-s	Cohesity Storage Domain name (optional - if not provided will use DefaultStorageDomain)

`cohesity_prepare_restore.sh`

Arguments that can be used with the `cohesity_prepare_restore.sh script`:

Argument	Description
-D	DB name to be restored
-I	Cohesity Cluster IP
-S	Cohesity SBT library path
-W	Allowed IP address/netmask bits separate by commas if multiple
-u	Cohesity username (optional - if not provided will use API key)
-p	Cohesity password (optional - if not provided will use API key)

Argument	Description
-V	The Cohesity data view used for backup the DB (optional - if not provided, the script will search all the views that have a backup for this DB name)
-s	start_time with format yyyy/mm/dd-HI:MM:SS (optional - if not provided, no limit on start time)
-e	end_time with format yyyy/mm/dd-HI:MM:SS (optional - if not provided, no limit on end time)
-A	The archival target if the restore needs to get data from the archival (optional - if not provided, only local snapshots will be searched)
-J	The job name or job UID of the job to be used for the restore (optional - if not provided, will use Yoda search to get the snapshots)
-c	The sbt certificate file on the Oracle host if secure gRPC will be used between the Cohesity cluster and the Oracle host (optional - if not provided, regular gRPC will be used)
-C	The sbt certificate file on the Linux management host if secure gRPC will be used between the Cohesity cluster and the Linux host (optional - if not provided, regular gRPC will be used)
-M	The file contains the list of views to be merged (optional - if provided, the script will merge the cloned views listed in the file instead of searching the snapshots internally)
-K	Keep the views undeleted (optional - if not provided, will delete the temp views generated by the script)
-Z	Indicates the restore will be done through ZDLRA. A catalog view will be created after running the script with this option. (optional - if not provided, only a data view will be created.)
-O	Overwrite the duplicate entries in the catalog view. (optional - if not provided, the script will fail if duplicate files are found.)

`cohesity_end_restore.sh`

Arguments that can be used with the `cohesity_end_restore.sh` script:

Argument	Description
-d	The temp directory generated by cohesity_prepare_restore.sh
-l	Cohesity Cluster IP
-u	Cohesity username (optional - if not provided will use API key)
-p	Cohesity password (optional - if not provided will use API key)

Replicate Oracle Database

Oracle application administrators can use Cohesity replication for cost-effective disaster recovery (DR). Cohesity provides a policy-based data replication solution from the core to the cloud or to the edge, from one cluster to another cluster. As part of replication, Cohesity always performs source-side compression and deduplication first, and sends only the changed data over the network.

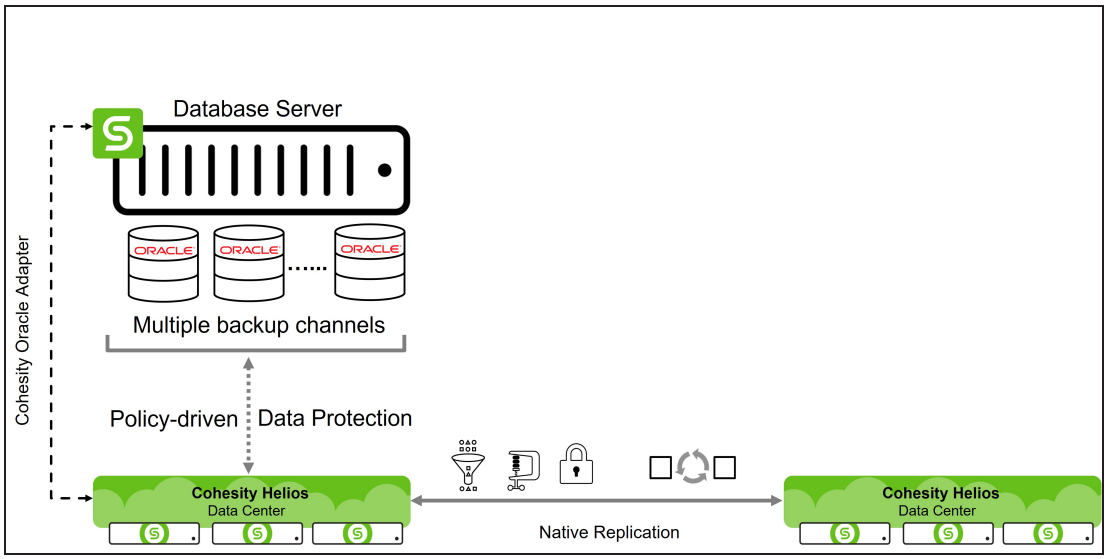

Replicate Oracle Database

For more information about database replication, see Creating a Connection to the Remote Cluster.

Considerations

1. Overwrite restore is not supported for replication sites.
2. Failover is not supported for the Cohesity Oracle adapter. You have to perform a force re-register of the same sources on the replication cluster and start the backup process all over again.
3. In case of a failover, after the re-registration, a full backup is required before you can perform an incremental or a log backup.

CloudArchive and CloudRetrieve for Oracle Database

Cohesity offers cost-effective, reliable data storage and protection solutions for rapidly growing environments that deploy Oracle Databases. The data in these environments can grow to terabytes and even petabytes in size.

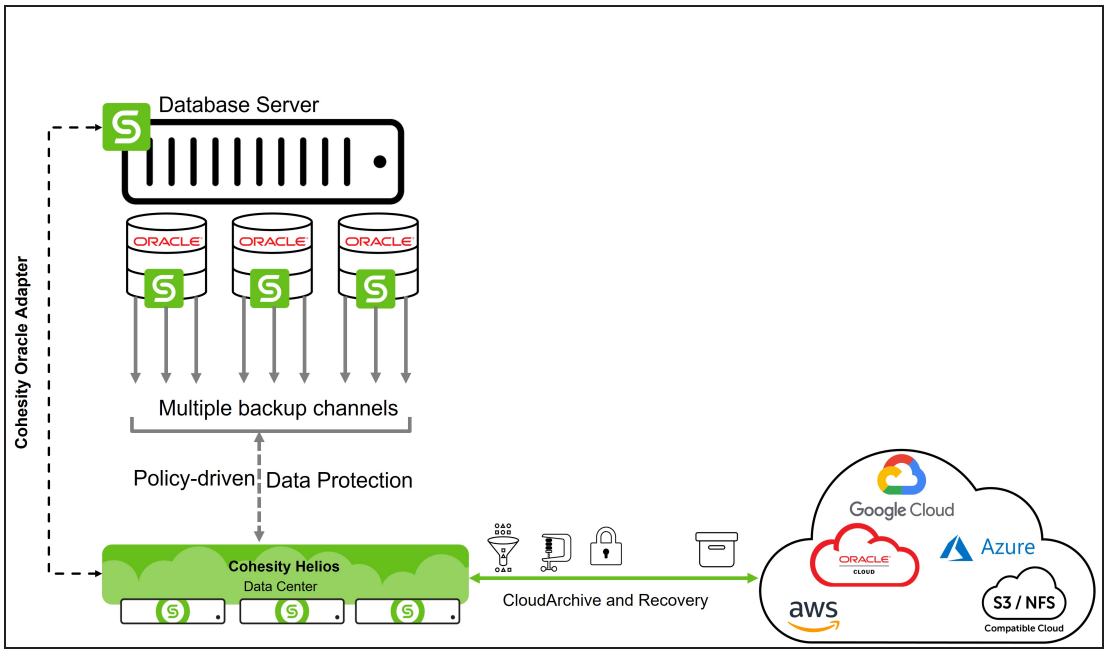

Cohesity provides a policy-based method to archive to public clouds (AWS, Azure, GCP, and Oracle Cloud), to any S3-compatible storage, to tape, and/or to any NFS mount point. Oracle application administrators can use Cohesity CloudArchive to address long-term data retention requirements and can choose to restore an entire database. The archived data is efficiently transferred and stored by sending only deduplicated, compressed incremental backups, thereby reducing network and storage utilization.

Once the data is archived, Oracle application administrators can make use of Cohesity CloudRetrieve features as a cost-effective alternative for disaster recovery and geo-redundancy. If the cluster from which the data is archived becomes unavailable, you can retrieve your data onto a different cluster.

CloudRetrieve for Oracle Database

The CloudRetrieve approach involves the following steps:

1. Register the cloud container (where you archived the Oracle Database) with a new Cohesity cluster.
2. Provide the search parameters, such as cluster name, Protection Group name, and date range.
3. Cohesity retrieves the data that you specify onto the new Cohesity cluster, based on your search of the metadata.

For more information, see CloudRetrieve Searches.

> **Note:** CloudArchive and CloudRetrieve are agnostic of Cohesity Oracle Adapter and Remote Adapter versions.

Oracle Database Troubleshooting

This topic provides the troubleshooting information for the Cohesity Oracle Adapter.

Oracle agent troubleshooting checklist

1. Verify the agent user (cohesityagent by default) is part of the dba and/or the oinstall groups. Use the command ID to check:

```
id cohesityagent
uid=1002(cohesityagent) gid=1001(dba) groups=1001(dba)
```

2. Verify the /etc/sudoers file has the agent configured. For example:

```
cohesityagent       ALL=(ALL)       NOPASSWD:ALL
Defaults:cohesityagent !requiretty
```

> **Note:** Use spaces and not tabs as tabs may lead to syntax errors.

3. Verify the home of the agent user is configured with the drwxr-xr-x minimum permission. For example:

```
ls -ld /home/cohesityagent
drwxr-xr-x 6 cohesityagent cohesityagent 198 Jan 17 08:58
/home/cohesityagent
```

   To change the permissions, execute the following command:

```
chmod 755 /home/cohesityagent
```

4. Verify `.bashrc` in the home of the agent user has `umask 0002`.
5. Verify the agent is running:

```
/etc/init.d/cohesity-agent status
```

6. Verify with the DBA that the database to be backed up is in `/etc/oratab` and has the correct `ORACLE_HOME` path.

> **Note:** Ensure the third field contains Y or N.

```
CohesityOra:/oracle/app/oracle/product/12.2.0/dbhome_1:Y
^database^ ^ORACLE_HOME path                          ^
```

7. If database access will use OS authentication, the Cohesity agent user must have sysdba privileges on the database. If database authentication will be used, the database user also requires sysdba privileges. Verify using sqlplus commands:

```
SQL> column sysdba format a10
SQL> select username,sysdba from v$pwfile_users;
USERNAME   SYSDBA
```

```
---------- ----------
SYS        TRUE
COHBACKUP  TRUE
```

8. Verify the Linux physical server has been registered as a source.
   a. Navigate to **Protection** > **Sources**.
   b. Click **Register Source** located at the top right of the page.
   c. Click **Physical Server**.
   d. Enter the Oracle server Host Address.
   e. Click **Register Server**.
9. Verify the Oracle Database is registered.
   a. Navigate to **Protection** > **Sources**.
   b. Click **Register Source** located at the top right of the page.
   c. Click **Oracle Source**.
   d. Accept the default Single Instance as the source type.
   e. Enter the Oracle server Host Address or browse the registered sources and select one.
   f. Select the **Authentication Type** to use when connecting to the database.
   g. Click **Register**.

Troubleshooting Agent Issues

The following table provides details about Cohesity Oracle adapter issues that you might encounter:

Issue	Cause	Workaround
Oracle Remote Adapter Job error "Execution of script failed with exit code: 255, error msg: Received disconnect from X.X.X.X port 22: Too many authentication failures for oracle Authentication failed"	This can occur if Public Key (Password less) authentication between Cohesity Nodes and the Oracle server is not working.	• Access the Cohesity bash shell, as described in How to access the bash shell on a Cohesity node and run the below command to confirm password less authentication is working to the Oracle server: `ssh oracle@<oracle_host>` • If the above step prompts for password, it means that the Public Key (Password less) authentication between Cohesity Nodes and the Oracle server is not working. Perform the following to resolve the issue: ◦ Attempt to re-copy the RSA key from the Cohesity Job edit page to the oracle server authorized_keys file for the corresponding user. ◦ Check the permissions on the authorized_keys file is owned by the correct user and file permission allows the same user to read. ◦ Refer to the Cohesity Documentation and search for the section Add or Edit a Remote Adapter Protection Group.

Oracle single instance database is not discovered after registering a Linux Oracle server	There is an incorrect `$ORACLE_HOME` directory path entry in the `/etc/oratab` file for the Oracle SID.	To resolve this error, verify that entries in `/etc/oratab` are following the correct format:
		`$ORACLE_SID:$ORACLE_HOME:<N|Y>`
		Example:
		`ORADB:/u01/app/oracle/product/11.2.0.4:N`

Issue	Cause	Workaround
Cohesity workaround for Oracle RMAN database recovery fails (ORA-01180 and ORA-01110)	This error occurs because of known Oracle issues. Refer to ORA-01180 and ORA-01110.	To workaround this issue, open the Cohesity UI and adjust the Point in Time slider to a later date to include more log files. For example, as shown in the following screen shot, move the Point in Time slider from August 19 to August 22. 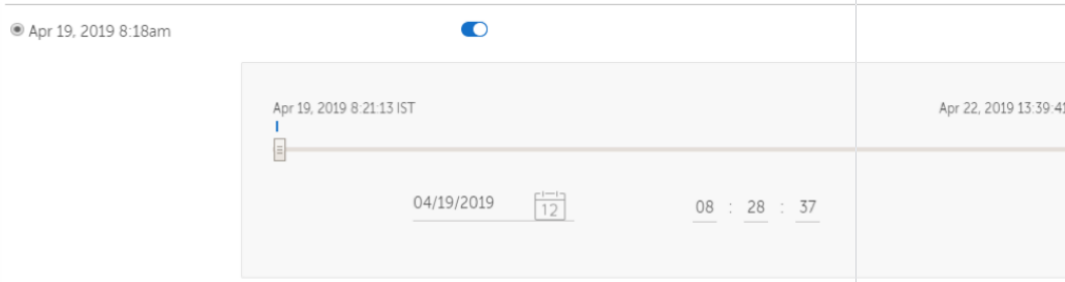
Oracle Remote Adapter recovery job to an alternate location fails with "RMAN-05501: aborting duplication of target database"	If the RMAN script includes the same file_ path as the original server when recovering in the destination server, Oracle will report an error.	Contact Cohesity support and reference this KB article. As a workaround, the script can be updated to ignore this error by adding NOFILENAMECHECK.

Issue	Cause	Workaround
Oracle protection job fails with the RMAN error "ORA-09925: Unable to create audit trail file"	Cohesity creates a new session and in turn a new Oracle trace file when the database backup progress query is run. This behavior is expected with Oracle. Refer to the Oracle documentation.	Cause Cohesity creates a new session and in turn a new Oracle trace file when the database backup progress query is run. This behavior is expected with Oracle. Refer to the Oracle documentation. Resolution Review the Oracle database settings and determine whether SQL_TRACING is enabled. In addition, kcb / krb tracing might be enabled.
Oracle RMAN protection job configured to use source-side deduplication (SBT) fails in environments without IPv6 enabled	By default, the SBT library uses sun_rpc calls which communicate over IPv6.	To resolve this issue: 1. Enable IPv6 on the Oracle database server(s). 2. Contact Cohesity support to download a new library which forces communication over IPv4. The new library requires that additional SBT parameters be used in the RMAN script.

Issue	Cause	Workaround
Cloning an Oracle database fails with the RMAN error "ORA-65093: multitenant container database not set up properly"	As indicated in the error message, Cohesity does not support cloning Oracle databases when multitenant container database (CDB) is enabled.	To request this capability, refer to Cohesity knowledge base article, How to create and view Request for Enhancements.

Issue	Cause	Workaround
Oracle database protection job fails with "[kTimeout]: RPC call MountOracle NAS timed out"	This error has been seen when the time required to set the Oracle Server permissions for the Cohesity user times out or because of a connectivity issue between the Oracle Server and the Cohesity Cluster. For example, this error has been seen when bi-directional communication over port 2049 is not open across network firewalls or when there is a networking issue.	For details on how to resolve this error, see the KB article.

Issue	Cause	Workaround
Oracle protection job fails with "[kAgentError]: ORA-00312: online log 2 thread 1: ORA-19809: limit exceeded for recovery files"	This failure results from a server-side issue.	Refer to the Oracle documentation on the ORA error codes. **ORA-000312: session marked for kill** **Cause**: The session specified in an ALTER SYSTEM KILL SESSION command cannot be killed immediately (because it is rolling back or blocked on a network operation), but it has been marked for kill. This means it will be killed as soon as possible after its current uninterruptible operation is done. **Action**: No action is required for the session to be killed, but further executions of the ALTER SYSTEM KILL SESSION command on this session may cause the session to be killed sooner. **ORA-19809: limit exceeded for recovery files** **Cause**: The limit for recovery files specified by the db_recovery_file_dest_size was exceeded. **Action**: The error is accompanied by 19804. See message 19804 for further details **ORA-19804: cannot reclaim string bytes disk space from string limit** **Cause**: Oracle cannot reclaim disk space of specified bytes from the `DB_RECOVERY_FILE_DEST_SIZE` limit. **Action**: There are five possible solutions: 1) Take frequent backup of recovery area using RMAN. 2) Consider changing RMAN retention policy. 3) Consider changing RMAN archivelog deletion policy. 4) Add disk space and increase `DB_RECOVERY_FILE_DEST_SIZE`. 5) Delete files from recovery area using RMAN.

Issue	Cause	Workaround
Oracle Remote Adapter protection job fails with the RMAN error " [kIOError]: [kAppError]: Dedup write failed with error: kQuotaExceeded"	These errors have been seen when a quota was set on the View used for the RMAN backup, and the View's quota nearly reached the limit.	To resolve this issue, increase the quota assigned to the View and re-run the job.

CockroachDB

Cohesity natively integrates with CockroachDB to provide backup and recovery for clustered CockroachDB deployments. Using Cohesity Connector Agent, you can backup and restore the CockroachDB database. The Cohesity Connector Agent provides a simple workflow to back up and restore your CockroachDB deployment.

Cohesity provides a CockroachDB connector for backup and restore of enterprise CockroachDB deployments. The CockroachDB connector leverages the backup and restore capabilities provided by CockroachDB.

The CockroachDB connector enables you to back up and restore the following:

- Backup:
  - Full backup of the following with or without revision history:
    - Cluster
    - Database
    - Table
  - Incremental backup of the following with or without revision history:
    - Cluster
    - Database
    - Table

- Restore: Full restore of the following:
  - Cluster
  - Database
  - Table
- Point In Time Restore (PITR) of the following:
  - Cluster
  - Database
  - Table

Core Components

The following are the core components for protecting CockroachDB databases:

- CockroachDB cluster
- Cohesity CockroachDB connector (DB Agent)
- Cohesity Connector Agent (Physical Agent)
- Cohesity S3 View

As illustrated, you can integrate Cohesity with your CockroachDB deployments using the Cohesity CockroachDB connector and Cohesity Connector Agent.

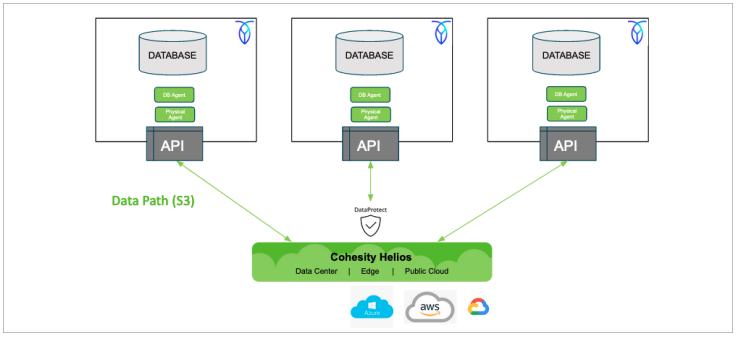

Plan and Prepare for CockroachDB Protection

Before you register your CockroachDB deployment as a source with Cohesity, and protect CockroachDB databases, ensure the following prerequisites:

- Identify the node(s) in your CockroachDB cluster with adequate resources to execute the CockroachDB connector.
- Ensure that the fully qualified domain name (FQDN) of the Cohesity cluster is resolvable from all the nodes in the CockroachDB cluster.
- A database user with database admin user privileges.

- Cohesity CockroachDB connector uses PostgreSQL client (psql) to connect to CockroachDB. Ensure that PostgreSQL client is already installed and available in PATH.
- Download and install the Cohesity Linux Agent on the CockroachDB node.
- Download and install the CockroachDB connector.
- Review and understand the considerations.

Supported Versions

For more information on supported CockroachDB versions, see Supported Software versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	CockroachDB Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
CockroachDB Local Host	CockroachDB Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
CockroachDB Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
CockroachDB Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following considerations before you protect your CockroachDB cluster.

- In a Recovery run, we can restore only one object at a time.
- Backup or restore might fail, if the S3 keys are changed when the backup or restore is in progress.

Download and Install the Cohesity Linux Agent

The Cohesity Linux Agent is available as script-based and RPM-based installers, based on your requirements download the installer and install it accordingly on the node(s) of the CockroachDB cluster . For detailed instructions, see Download and Install the Linux Agent.

> **Note:** Ensure that you install the Cohesity Linux Agent using the RPM-based installer.

Download and install the CockroachDB Connector

Cohesity provides a CockroachDB connector for backup and restore of enterprise CockroachDB deployments. The CockroachDB connector leverages the backup and restore capabilities provided by CockroachDB.

Download the CockroachDB Connector

Cohesity provides a CockroachDB connector to backup and restore CockroachDB databases. The CockroachDB connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for CockroachDB connector, and copy the datasource agent to the CockroachDB host.

Install the CockroachDB Connector

Before you install the CockroachDB Connector, ensure that you have sudo privileges.

To install the CockroachDB connector:

1. Log on to the CockroachDB node where you have copied the CockroachDB connector RPM-based installer, `cohesity-cockroachdb-connector-<version>.x86_64.rpm`.
2. Run the following command to install the CockroachDB connector:

   ```
   sudo rpm -ivh /home/cockroach/cohesity-cockroachdb-connector-
   <version>.x86_64.rpm --prefix <connector_install_directory>
   ```

   Where `<connector_install_directory>` is the CockroachDB node's directory where you want to install the CockroachDB connector. If you do not define the installation directory, then by default, the connector is installed in the `/opt/cohesity/cockroachdb/scripts` directory.

Uninstall the CockroachDB Connector

To uninstall the CockroachDB Connector:

1. Log on to the CockroachDB node where you have installed the CockroachDB connector.
2. Run the following command to uninstall the CockroachDB connector:

```
#sudo rpm -e cohesity-cockroachdb-connector-<version>.x86_64
```

Register and Manage the CockroachDB Source

To protect your CockroachDB databases, you need to register your CockroachDB deployment as a source on Cohesity. After registering your CockroachDB deployment as a source, you can:

- Update the CockroachDB source configuration, such as:
  - Add more node(s) to run the CockroachDB connector. Ensure that you have installed the Cohesity Linux Agent and CockroachDB connector on the node(s).
  - Update the username and password of the database user.
- Refresh the source details.
- Unregister the CockroachDB source from Cohesity.

Register CockroachDB as a source

Before you register CockroachDB as a source, ensure that you:

- Review the requirements and considerations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for CockroachDB Protection.

To register CockroachDB as a source:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** at the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
   1. From the Source Type drop-down menu, select **CockroachDB**.
   2. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the node(s) you have identified to run the CockroachDB connector. For more information, see Plan and Prepare for CockroachDB Protection.
   3. In the **Datasource Agent Installation Path** field, enter the directory path on the CockroachDB node, where you have installed the CockroachDB connector. For more information, see Install the CockroachDB Connector.
   4. In the **APP Authentication** section, enter the **Username** and **Password** of the CockroachDB database admin user.

> **Note:** The **Password** field is optional.

5. Enable **Custom Options**, and perform the following:
   1. Ignore the **Mounts View** option. It is not applicable for CockroachDB source registration.

      > **Note:** The **Mount View** option is not applicable, and ensure you do not enable the option.

   2. Optional. In the **Source Registration Options** field, enter the following source registration parameters based on your requirement:

Parameters	Mandatory?	Description
`--cluster-name`	No	(optional) A unique name to identify the CockroachDB cluster.
`--port`	No	The port used to connect to the CockroachDB cluster. Default : 26257
`--certs-dir`	No	If you want the Cohesity to authenticate to the CockroachDB cluster using client certificate and key authentication then provide the path to the certificate directory containing all certificates and keys in the CockroachDB node. The certificate file names must be in following format: • CA certificate: `ca.crt` • Client certificate for the user: `client.<user>.crt`. Example: The root user certificate file must be `client.root.crt`. • Client key for the user: `client.<user>.key` Example: The root user client key file must be `client.root.key`
`--keytab`	No	If you want the Cohesity to authenticate to the CockroachDB cluster using kerberos authentication then provide the path to the Keytab file with client principals.

6. Click **Register**.

Update the CockroachDB Source Configuration

After registering your CockroachDB as a source on Cohesity, you can:

- Add more node(s) to run the CockroachDB connector. Ensure that you have installed the Cohesity Linux Agent and CockroachDB connector on the node(s).
- Update the username and password of the database user.
- Modify the **Global Source Registration Parameters** based on your requirements.
- Update the **Script Directory Path** only if you have uninstalled and reinstalled the CockroachDB connector in a different directory on the CockroachDB node.

> **Note:** Ensure that the CockroachDB connector is installed on all the node(s) in the Cohesity cluster.

To update the CockroachDB source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the CockroachDB source and select **Edit**.
3. On the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the CockroachDB Source Configuration

You can refresh the CockroachDB source and re-verify the source with the available source registration parameters.

To refresh the CockroachDB source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the CockroachDB source and select **Refresh**.

Unregister the CockroachDB Source

If you plan to stop taking the backup of your CockroachDB databases, you can unregister the CockroachDB source from the Cohesity Cluster.

To unregister the CockroachDB source:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the CockroachDB source and select **Unregister**.

Backup CockroachDB Database

Cohesity provides a simple workflow to back up your CockroachDB databases. After registering the CockroachDB source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides a CockroachDB connector for backup and restore of enterprise CockroachDB deployments. The CockroachDB connector leverages the backup capabilities provided by CockroachDB.

After registering the CockroachDB source on Cohesity, you can create a Protection Group and specify the objects such as databases and tables that you want to backup. In the Protection Group, you need to define the parameters that enable you to provide additional parameters required by the backup statements provided by CockroachDB.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your CockroachDB backup configuration:

Backup Parameters	Description
`--cluster-backup`	Enables you to create a backup of the CockroachDB cluster. If you define this parameter, then the objects specified in the Protection Group are ignored. For example: `--cluster-backup=true`
`--custom-args`	Allows you to use the backup options provided by CockroachDB. For example: `--custom-args=<backup-options>`

Protect CockroachDB Deployments

CockroachDB deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to:

- Specify the objects such as databases and tables that you want to backup.
- Based on your requirement, define the parameters for the backup statements provided by CockroachDB. For more information, see Backup Parameters.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

To protect CockroachDB deployments:

1. Do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the CockroachDB source and then select **Protect**.
2. From the **Source** drop-down list, select the registered CockroachDB source or click **Register Source** to register a CockroachDB source. For more information, see Register CockroachDB as a source.
3. In the **Objects** section, add the objects that you plan to protect. The object can be a database or a table. To back up a database, enter the name of the database. If you want to backup a table, enter the database name and the table name with a dot(**.**). For example, `bank.account`. Where `bank` is the database name and `account` is the table name.
4. Click the plus icon **(+)** to enter multiple objects that you plan to back up. Then, click **Save Selection**.
5. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Protect CockroachDB Deployments.
7. In the Settings section, do the following:
   1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
   2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

      Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
   3. In the **Custom Options** section, perform any one of the following:
      - Apart from the `--custom-args` parameters, do not enter any backup parameters in the **Full Backup** and **Incremental Backup** fields if you plan to protect only the specified object(s) in the Protection Group.
      - Enter the Backup Parameters in the **Full Backup** and **Incremental Backup** fields if you plan to back up the entire CockroachDB cluster.

Ensure that you enter the same backup parameters in the **Full Backup** and **Incremental Backup** fields.

For example:

Use Cases	Sample Parameters
Create a full or incremental backup of the CockroachDB cluster. **Note:** In the case of the CockroachDB cluster backup, the objects you have defined in the Protection Group is ignored.	`--cluster-backup=true`
Create a full or incremental backup of the CockroachDB cluster with full revision history **Note:** Based on your requirements you can also use various backup options provided by CockroachDB. You can use all backup options provided by CockroachDB, except the `detached` option.	`--cluster-backup=true --custom-args=revision_history`

- Ignore the **Mounts** field. It is not applicable for CockroachDB data protection.

8. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click Edit on the right. For more information, see Additional Settings.
9. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.

Additional Settings	Description
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the cluster. By default, maximum scaling is capped at 16. > **Note:** Not applicable for CockroachDB data protection.
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the CockroachDB Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new databases or tables to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for CockroachDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan

to back up. click **Save Selection**.

6. Click **Save**.

Remove Object(s) from the Protection Group

You can remove object(s) from the Protection Group that you do not want to back up in the next protection run. When you remove an object from the protection group, incremental backup of the object is not performed from the next protection run. The existing backup for the object is preserved until the snapshot expires.

To remove object(s) from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for CockroachDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ on the **Objects** section.
5. In the **Objects** section, Click the minus icon (**-**) next to the objects to remove them from the backup. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for CockroachDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ on the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for CockroachDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

When you cancel a protection run, the CockroachDB connector leverages the Cancel Job capabilities provided by CockroachDB.

To cancel a recovery run:

1. Navigate to **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for CockroachDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover CockroachDB Database

You can restore the CockroachDB cluster or object(s) such as databases and tables in CockroachDB. You can restore to the same CockroachDB cluster or to a different CockroachDB cluster.

The CockroachDB connector leverages the restore statements provided by CockroachDB. You can create a Recovery configuration and specify the objects such as databases and tables backed up using the Protection Group.

In the Recovery configuration, you need to specify the object you plan to restore and define the script parameters required by the restore statements provided by CockroachDB.

Recovery Parameters

In the Recovery configuration, you can define the following script parameters for your CockroachDB recovery:

Recovery Parameters	Description
`--cluster-restore`	Enables you to restore the CockroachDB cluster. If you define this parameter, then the object you have specified for recovery is ignored. For example: `--cluster-restore=true`

Recovery Parameters	Description
`--custom-args`	Allows you to use the restore options provided by the CockroachDB. For example: `--custom-args=<restore-options>` **Note:** The `detached` and `into_db` options from the restore options are not supported.
`--pitr-epoch`	Allows you to perform Point In Time Restore (PITR). You need to provide the restore point in the epoch format. For example, on a Linux machine, you can find the epoch for a required time using the following command: `# date "+%s" -d "02/28/2021 15:41:15"` For example: `--pitr-epoch=1632125909` **Note:** You can perform PITR only if the corresponding backup is taken with revision history.

Recover CockroachDB Databases

To recover CockroachDB:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same CockroachDB cluster.
New Object	To restore to a different CockroachDB cluster. Select the **Target** CockroachDB cluster to which the data needs to be restored. Also, enable **Overwrite existing object with the same name** and type **Yes**, if you want to overwrite the existing object with the same name.

8. In the **Recovery Options** section, configure the following options:

Options	Description
Mounts	Not applicable for CockroachDB recovery.
Concurrency	Not applicable for CockroachDB recovery.
Custom Options	Specify the Recovery Parameters based on your requirements. For examples, see Sample Recovery Use Cases
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

9. Click **Start Recovery**.

Sample Recovery Use Cases

Use Case	Sample Recovery Parameters
Restore the CockroachDB cluster	`--cluster-restore=true` **Note:** If you have specified this parameter then the object defined in the Recovery configuration is ignored.

Use Case	Sample Recovery Parameters
Restore a database	Do not specify the parameter: `--cluster-restore=true` **Note:** You can use other parameters such as `--custom-args=<restore-options>` based on your requirements. For example, If you plan to restore a database and want to ignore sequence dependencies during the restore process. You can specify the parameter as follows: `--custom-args=skip_missing_sequences`
Restore a table	Do not specify the parameter: `--cluster-restore=true` **Note:** You can use other parameters such as, `--custom-args=<restore-options>` based on your requirements. For example, if you plan to restore either a table that was previously a sequence owner or a sequence that was previously owned by a table. You can specify the parameter as follows: `--custom-args=skip_missing_sequence_owners`
Point In Time Restore (PITR)	You need to provide the restore point in the epoch format. For example, on a Linux machine, you can find the epoch for a required time using the following command: `# date "+%s" -d "02/28/2021 15:41:15"` To perform PITR, specify the parameter as follows: `--pitr-epoch=value` For example: `--pitr-epoch=1632125909`

Cancel a Recovery Run

When you cancel a recovery run, the CockroachDB connector leverages the Cancel Job capabilities provided by CockroachDB.

To cancel a recovery run:

1. Navigate to **Data Protection > Recoveries**.
2. Click the action icon (⋮) next to the recovery run and click **Cancel Run**.
3. In the Cancel Recovery dialog box, select any one of the following:
   - Click **Continue Recovery** if you want to continue to recovery run
   - Click **Cancel Recovery** if you want to cancel the recovery run.

Epic Database Protection using Cohesity Physical Agents

Epic database administrators face the challenges of prolonged backup durations, rapid data growth, escalating storage costs, and the need for enhanced flexibility and storage management tools within Epic.

In addition to these challenges, today's database protection must encompass more than obtaining a clean copy; it must also include factors such as security, storage efficiency, minimizing impact on production systems, automation, and scalability.

Cohesity protection for Epic provides a solution to these challenges. It simplifies database backups and restores through secure, streamlined workflows. You can protect and manage your workloads and execute available protection and recovery workflows with a few steps. The Cohesity Physical Agents provide flexible deployment options to make Epic backup and restores simple and secure.

The backup and recovery of the Epic Caché database involves snapshotting the LUNs on the storage array. This method is agnostic of the OS on which Epic is hosted. Also, this method doesn't need the mount host but still requires a Proxy server to host the Epic freeze- thaw script.

Plan and Prepare for Epic Database Protection

This topic provides information on minimum permissions, prerequisites, and considerations to protect physical servers with a Cohesity cluster.

Support Matrix

Before you register your physical servers with the Cohesity cluster, ensure that Cohesity supports the physical server versions. For more information, see Supported Software.

Minimum Agent Requirements

For details on hardware and software agent requirements, see Minimum Agent Requirements.

How Cohesity Works with Epic Databases

Using the Mount host method for backing up the Epic database involves utilizing the Cohesity Linux or AIX agents on an independent mount host. The mount host has the Epic LUNs mounted to it from a previous Snap of the primary storage.

Cohesity's Linux or AIX agent integrates with the mount point host to provide a file-based backup and recovery solution for Epic deployments.

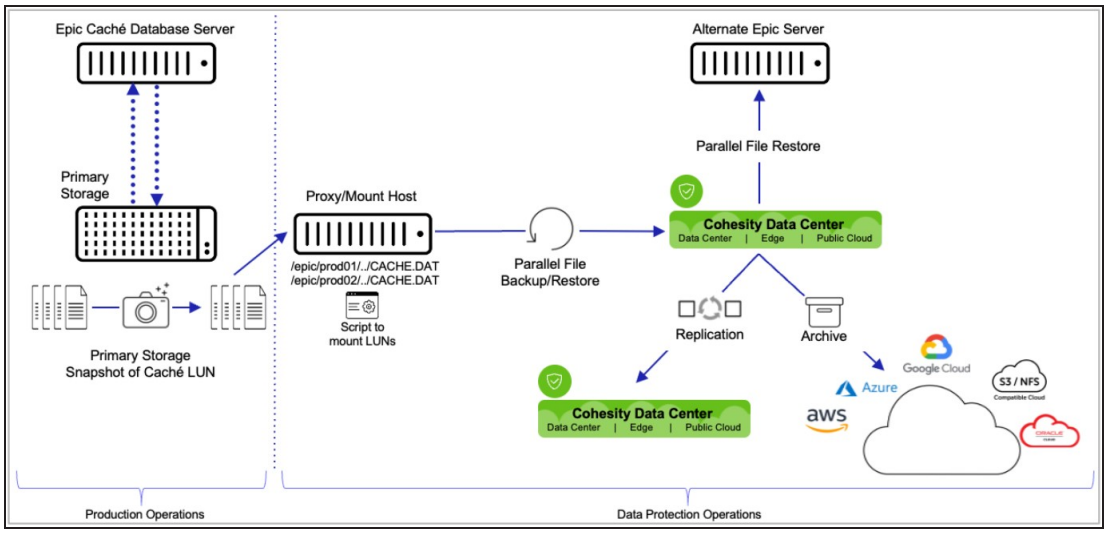

The following steps illustrate the steps that the script handles, and it can be modified to meet your requirements.

1. Script (executed by Cohesity Pre-Script option)
    a. Freeze Caché in production.
    b. Take a SAN/LUN Snapshot on primary storage.
    c. Unfreeze Caché.
    d. Present the snapshot from primary storage to the proxy host.
    e. Mount the File System on the proxy host.
2. Initiate a Cohesity file-based backup of the filesystem
3. Script (continues from Cohesity Pre-Script option)
    a. Unmount the file system from the proxy host.
    b. Delete snapshot in primary storage.

Deployment Steps

Install the Cohesity Physical Agents for Linux or AIX on the mount point host you plan to use.

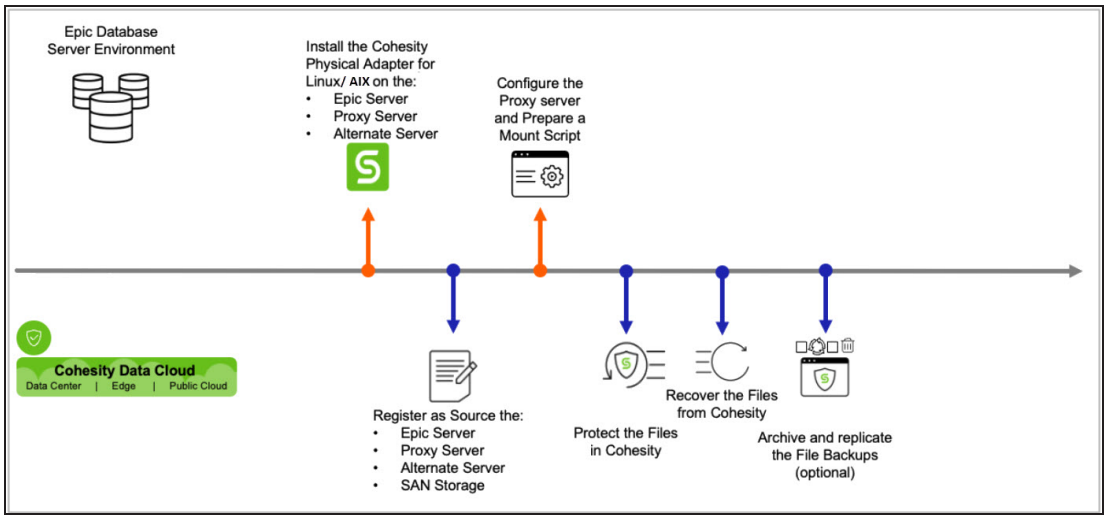

Once you complete the installation of the Physical Agents on each host, you are ready to register them as a source.

You must complete the following steps to configure the Epic environment for protection:

1. Register Epic Database Source
2. Mount Host Server Script Setup for Epic Database

Register Epic Database Source

Downloading and installing the Physical Agents for Linux or AIX on the mount point server allows you to register it as a source with the Cohesity cluster. The Cohesity cluster can then use a file-based Protection Group for backups.

Before you install the Physical Agents, ensure the following prerequisites are met. For details, see Port Requirements and Firewall Settings.

Note: Install the Cohesity Linux or AIX Physical Agents on each host that is potentially a restore target.

Install the Linux or AIX Agent on Selected Hosts

Cohesity recommends the RPM installer for Linux installations. For details on installing the Cohesity Physical Agents for Linux or AIX, refer to the **Linux** or **AIX** sections in Download and Install Agent for Physical Servers.

> **Note:** The Cohesity cluster can be tuned using cluster and agent gflag values to achieve its performance. Your Cohesity team needs to determine recommended flag values based on the size of your cluster. Contact the Cohesity Account team for cluster configuration.

Mount Host Server Script Setup for Epic Database

The mount host is used as the target location where the script mounts the newly snapped LUNs. These LUNs are then backed up by Cohesity, and it is also the location for the existing script.

All Epic customers generally have the Epic freeze and thaw script that snapshots the Epic database and mounts it on a mount host.

This script can also live on a script server (the most common method) or be part of Cohesity agent pre- and post-script (uncommon, but wholly supported).

The examples in the following sections will show you how to incorporate your existing script into the Cohesity agent pre- and post-script options in the protection group.

Your existing script must perform these major tasks:

1. Freeze the Epic database.
2. Execute a snap of the storage LUNs.
3. Unfreeze the Epic database.
4. Mount the newly snapped LUNs on the mount host.

Use the mount host method:

- If granular file folder recovery is required.
- If zero impact to the production server is required.

Existing Script Location and Modification

1. Cohesity will execute your script as part of the protection job. Copy the script to the user_scripts directory on the mount host, located in `/opt/cohesity/agent/software/crux/bin/user_scripts`.
2. Where the script calls out to the existing backup software, comment that step out of the script.

Register Mount Host as a Server Source

The mount point method protects the Epic database at the primary storage level by taking a snap of the LUNs via script execution. To register your mount point host as a Source on a Cohesity cluster, see Register a Physical Server with Cohesity Cluster.

By registering the mount point host, you prepare the server for script execution and Epic file protection by Cohesity.

Protect Epic Database (File-Based)

After registering your mount point host on the Cohesity cluster, you can start configuring your Cohesity Protection Group. For more information, see Protect Physical Server (File-

Based).

When creating a Protection Group for the Epic Databases, configure the **Pre & Postscripts Settings** using the settings listed below:

Cohesity provides you the option to run user-defined scripts as a part of the Protection Group. When creating a Protection Group, you can configure:

- Pre-Scripts: Scripts that will be executed before the protection run.
- Script Path: A path to the script is not needed if the script is located in the directory: `/opt/cohesity/agent/software/crux/bin/user_scripts` on the mount host.
- Script Params: Parameters entered here are appended to the script at runtime and are subject to PowerShell parsing rules.
- Continue Backup if script fails: Disable this option. By default, this option is enabled.

  > **Note:** Cohesity recommends disabling this option so that the protection job can accurately report the success or failure of the run. Additionally the script must have zero exit code to properly report success.

- Postscripts: Scripts that will be executed after the protection run.
- Post Snapshot Scripts: Scripts that will be executed after the snapshot has been created on the source.

  > **Important:** Cohesity does not validate the content of the script and is not responsible for the content of the script or any actions done by the script. You must verify the script actions before a script or executable is run.

You can find a more in-depth discussion of Pre & Post Script in Configure Pre & Post Scripts.

Restore Epic Database

Cohesity provides the ability to restore individual database files. You can restore them to their original location or an alternate location. For more information, see Recover Physical Server Files or Folders.

Once the Epic files are recovered, you must restart the Epic server.

IBM DB2

Cohesity natively integrates with IBM DB2 to provide backup and recovery for IBM DB2 deployments on Linux x86_64 platform. Using Cohesity Connector Agent, you can backup and restore the IBM DB2 databases.

Cohesity provides a DB2 connector that enables you to perform a full backup, incremental backup, log backup, restore, and point in time restore at the database level.

Core Components

The following are the core components for protecting IBM DB2 databases:

- IBM DB2 cluster
- Cohesity IBM DB2 connector
- Cohesity Connector Agent

As illustrated, you can integrate the Cohesity with your IBM DB2 deployments using the Cohesity IBM DB2 connector and Cohesity Connector Agent.

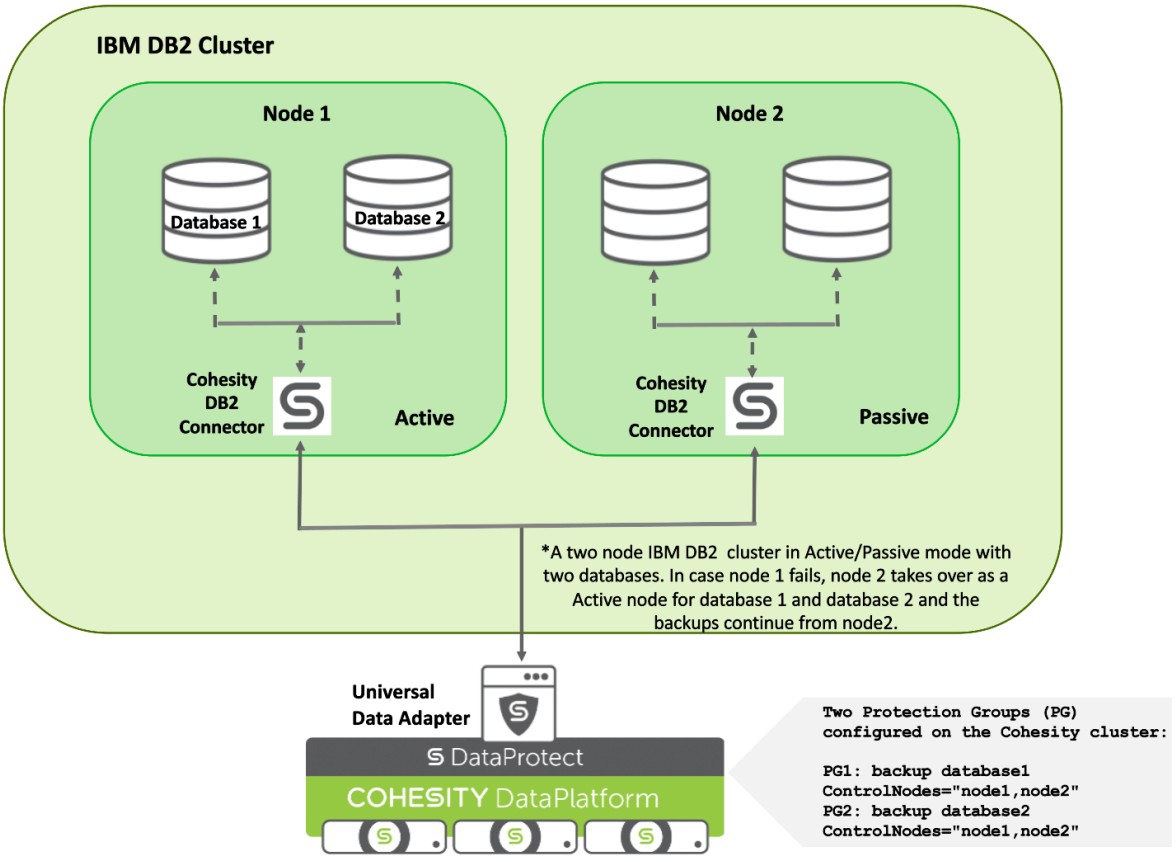

Plan and Prepare for IBM DB2 Protection

Before you register your IBM DB2 deployment as a source with Cohesity Platform, ensure the following prerequisites:

- The user account of the DB2 instance must have SYSADM, SYSCTRL, or SYSMAINT privilege.
- Download and install the Cohesity Linux Agent on the IBM DB2 nodes.
- Download and install the Cohesity DB2 connector.
- Review and understand the prerequisites.

Supported Versions

For more information on supported IBM DB2 versions, see Supported Software versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	IBM DB2 Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
IBM DB2 Local Host	IBM DB2 Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.
IBM DB2 Local Host	IBM DB2 Local Host	58881	TCP/IP	Local data transfer
IBM DB2 Local Host	IBM DB2 Local Host	58882	TCP/IP	Local data transfer

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
IBM DB2 Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
IBM DB2 Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following before you protect your IBM DB2 cluster:

- Only database-level backup is supported.
- Once a database is configured to be protected using Cohesity, users should not initiate a manual backup on the same database outside of the Cohesity cluster.
- The same database should not be part of multiple Protection Groups.
- Backup with compression and encryption options are not supported.
- Using the Cohesity DB2 Connector, you can protect an IBM DB2 HADR active/standby cluster. But, you can backup or restore from the active node only.
- Protection of DB2 databases with Purescale is not supported.
- If the hole is introduced in the backup timeline, then point-in-time recovery (PITR) is possible only after the subsequent successful full or incremental backup is completed.
- Source side deduplication is not supported.
- Log archival method configured for a source during source registration can not be changed. Any change in the log archival method needs re-registration of the source with different archival method.
- Cohesity supports backup and restore of DB2 databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.
- The Cohesity DB2 connector is only compatible with the bash shell and cannot be used with any other shells.
- If you have set a limit on Java heap memory usage, ensure that at least 512 MB is allocated per VIP. For example, if you are using 4 VIPs, the total Java heap memory usage limit should be at least 2 GB.
- Cohesity supports data protection for IBM DB2 deployments with SSL-encrypted communication between the IBM DB2 client and server.

Prerequisites

Ensure that:

- You use a root user or non-root user to install the Cohesity Linux agent and IBM DB2 Connector. If you choose to use a non-root user, you must add the `sudo` privileges for the non-root user using the syntax given below:

  ```
  <user_name> ALL=(ALL) NOPASSWD:SETENV: /path/to/command1,
  /path/to/command2, /path/to/command3
  ```

  The following are the `sudo` commands for IBM DB2 Connector:

  ```
  cohesityagent ALL=NOPASSWD:SETENV:/usr/bin/rpm, /usr/bin/find,
  /usr/bin/stat,/usr/bin/test, /usr/bin/mv, /usr/bin/cat, /usr/bin/mkdir,
  /usr/bin/rm, /usr/bin/cp, /usr/bin/chown, /usr/bin/chmod,
  /usr/bin/kinit, /opt/cohesity/db2/bin/named_pipe_snap_fs_util,
  /usr/bin/su, /usr/bin/java, /usr/bin/netstat, /usr/bin/sed,
  /usr/bin/cmp, /usr/bin/tee,
  ```

```
/usr/bin/wc,/usr/bin/truncate,/usr/sbin/runuser, /usr/bin/ps,
/usr/bin/cat, /usr/bin/true
```

Where, `/opt/cohesity/db2` is the IBM DB2 connector installation directory.

> **Note:**
> - If the **cohesityagent** user entry is already present in **/etc/sudoers**, append the list of commands to the existing entry.
> - The actual paths for each command might vary depending on the Platform and OS distribution. For more information, see Download and Install the Linux Agent.

- Ensure that the user account of the DB2 instance is assigned the required system authorities to perform backup or restore operations on a DB2 instance.
- For Kerberos authentication, the **keytab** file must be available on all nodes at the given location and must be accessible by the DB2 instance and the protection (backup/restore) user.
- Install Java 1.8.0_382 or higher version on all the IBM DB2 nodes.
- The Cohesity DB2 connector is only compatible with the bash shell and cannot be used with any other shells.
- Set "`TRACKMOD`" to "`YES`" in the database configuration. This is a mandatory prerequisites for using Cohesity DB2 connector with online backup (full and incremental). If `TRACKMOD` is not set to `YES`, backup jobs fails before execution.

> **Note:** `TRACKMOD` is not required for offline backups.

- If you are using the Dump and Sweep based log archival (log backup) method, set "`LOGARCHMETH1`" to a specific path.

  For more information, see the BACKUP DATABASE command reference. If you have not performed the mentioned database configurations then you must manually configure the database configuration on each database in the cluster and requires database downtime. The following is the sample steps for the database configuration:

```
db2 connect to <database>
db2 update database configuration using LOGARCHMETH1 'DISK:<log_
archive_path>'
db2 "update db cfg using TRACKMOD YES"
db2 terminate
db2 deactivate database <database>
```

```
db2 backup database <database> to /dev/null
db2 activate database <database>
```

> **Note:** The database goes into the BACKUP PENDING state after applying the required configurations. As shown in the sample steps, ensure that you take an offline backup of the database to prevent the BACKUP PENDING state. For more information, see the IBM Support article on BACKUP PENDING state.

- To securely connect to the Cohesity cluster while reading or writing data, you can use secure gRPC.

  > **Note:**
  > - It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
  > - Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
  > - Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

  To use secure gRPC, perform the following:

  1. Generate the certificate config file on the Cohesity node using the following command:

     ```
     ~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name <clustername> -output_file /home/cohesity/<clustername>.cfg
     ```

  2. Copy `/home/cohesity/<clustername>.cfg` from the Cohesity node to all IBM DB2 nodes at the same location. For example, `/opt/certs/<clustername>.cfg`.

     > **Note:** Ensure that the `root` or `non-root` user used to install the Cohesity Linux agent and IBM DB2 Connector has read access to the copied certificate config file.

  3. While registering the IBM DB2 source, provide the path of the certificate config file in the **Cohesity SSL Settings** field.

Required System Authorities

To perform backup or restore operations on a DB2 instance, the protection user account of the DB2 instance must have one of the following levels of authority:

- **SYSADM**
- **SYSCTRL**
- **SYSMAINT**

When restoring to a new database server, either **SYSADM** or **SYSCTRL** authority is required.

System authorities (**SYSADM**, **SYSCTRL**, and **SYSMAINT**) are determined by group membership and are defined in the database manager configuration file for a given instance:

- **SYSADM Authority**
  - Assigned to the group specified by the `sysadm_group` configuration parameter.
  - Users with **SYSADM** authority can also perform the functions of users with **SYSCTRL** and **SYSMAINT** authorities.
- **SYSCTRL Authority**
  - Assigned to the group specified by the `sysctrl_group` configuration parameter.
  - Users with **SYSCTRL** authority can also perform the functions of users with **SYSMAINT** authority.
- **SYSMAINT Authority**
  - Assigned to the group specified by the `sysmaint_group` configuration parameter.

Download and install the Cohesity Linux Agent

The Cohesity Linux Agent is available as an RPM-based installer. Based on your requirements, download the installer and install it accordingly on the node(s) of the IBM DB2 cluster. For detailed instructions, see Download and Install the Linux Agent.

> **Note:** Ensure that you install the Cohesity Linux Agent using the RPM-based installer.

Download and install the IBM DB2 Connector

Cohesity Platform provides IBM DB2 connector RPM for backup and restore of IBM DB2 deployments.

Download the IBM DB2 Connector

Cohesity Platform provides a IBM DB2 Connector to backup and restore IBM DB2 databases. The IBM DB2 Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for IBM DB2 Connector, and copy the datasource agent to the IBM DB2 node.

Install the Cohesity DB2 Connector RPM

Before you install the Cohesity DB2 connector, ensure that you have `sudo` privileges.

To install the Cohesity DB2 connector:

1. Log on to the IBM DB2 node where you have copied the Cohesity DB2 connector RPM-based installer, `cohesity-db2-connector-<version>.x86_64.rpm`.
2. Run the following command to install the Cohesity DB2 connector:

   Using `yum`:

   ```
   sudo sh -c 'yum localinstall /tmp/cohesity/cohesity-db2-connector-<version>.x86_64.rpm'
   ```

   Using `rpm -ivh`:

   ```
   sudo rpm -ivh cohesity-db2-connector-<version>.x86_64.rpm
   ```

Uninstall the Cohesity DB2 Connector

To uninstall the Cohesity DB2 connector, log on to the IBM DB2 node where you have installed the Cohesity DB2 connector and run the following command:

Using `yum`:

```
sudo yum remove cohesity-db2-connector-<version>.x86_64.rpm
```

Using `rpm`:

```
sudo rpm -e cohesity-db2-connector-<version>.x86_64.rpm
```

Upgrade the Cohesity DB2 Connector

To upgrade the Cohesity DB2 connector, run the following command on IBM DB2 node:

```
sudo rpm -U cohesity-db2-connector-<version>.x86_64.rpm
```

Register and Manage the IBM DB2 Source

To protect your IBM DB2 databases, you need to register your IBM DB2 deployment as a source on the Cohesity cluster. After registering your IBM DB2 deployment as a source, you can:

- Update the IBM DB2 source configuration.
- Unregister the IBM DB2 source from the Cohesity cluster.

Register IBM DB2 as a Source

Before you register IBM DB2 as a source, ensure that you:

- Read and understand the log backup methods provided by Cohesity
- Review the considerations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for IBM DB2 Protection.

To register IBM DB2 as a source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
   a. From the **Source Type** drop-down menu, select **IBM DB2**.
   b. From the **Host OS Type** drop-down menu, select **Linux**.
   c. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the node(s) that you have identified to run the Cohesity DB2 connector. For more information, see Plan and Prepare for IBM DB2 Protection.

      > **Note:** Cohesity supports backup and restore of DB2 databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

   d. In the **Datasource Agent Installation Path** field, update the directory path if you have installed the Cohesity DB2 connector at the non-default location. For more information, see Install the Cohesity DB2 Connector RPM.
   e. Optional. Enable **Kerberos Authentication** Settings, if you want to connect to the IBM DB2 database using Kerberos-based authentication. Enter the following details:

Field	Mandatory?	Description
Kerberos Principal	No	The Kerberos principal that is required to connect to the database.

Field	Mandatory?	Description
Keytab File Path	No	The absolute path of the keytab file for the given Kerberos principal. **Note:** The path must be accessible by the protection(backup/restore) user.
Kerberos credential cache path	No	The absolute path of the Kerberos ticket cache. **Note:** The path must be accessible by the protection(backup/restore) user.

> **Note:** At a time, you can use either Keytab OR Kerberos cache-based authentication.

f. Enable **Cohesity SSL Settings** to enable a secure connection to the Cohesity cluster while reading or writing data. Enter the path to the SSL certificate used to securely connect to the Cohesity cluster.

> **Note:**
> - It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
> - Ensure that you generate and copy the certificate config file from the Cohesity node to all the IBM DB2 nodes. For more information, see Prerequisites.
> - Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
> - Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

g. Under **Source Settings**, enter the details of the IBM DB2 source:

Field	Mandatory?	Description
Datasource Name	Yes	A unique name to identify your IBM DB2 cluster. > **Note:** After the source registration, you cannot change the Datasource name.
Protection Username	Yes	The username of the user account in the IBM DB2 server that has required privileges to perform backup or restore of DB2 databases. > **Note:** The protection user must have SYSADM, SYSCTRL, or SYSMAINT privileges.
Instance Name	Yes	The name of your IBM DB2 instance.
Environment Variables	No	Provide the environment variables if you want to override the environment's default value. For example, `PATH=/sbin:/bin:/usr/bin`
Use Cohesity storage for log archival using LOGARCHMETH2	No	Enable to use Cohesity storage for log archival using LOGARCHMETH2. For more information, see Log Backup Methods. > **Note:** By default the option is disabled and the dump and sweep log backup method is used. For more information, see Log Backup Methods.
Profile Path	Yes	The path to the DB2 Profile of the instance. For example: `/home/db2inst1/sqllib/db2profile`

h. Click **Register**.

Update the IBM DB2 Source Configuration

After registering your IBM DB2 as a source on the Cohesity cluster, you can update the IBM DB2 source configuration.

> **Note:**
> - Ensure that the Cohesity DB2 connector is installed on all the node(s) in the IBM DB2 cluster.
> - You cannot modify the source name.

To update the IBM DB2 source configuration:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the IBM DB2 source and select **Edit**.
3. On the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Unregister the IBM DB2 Source

If you plan to stop protecting your IBM DB2 databases, you can unregister the IBM DB2 source from the Cohesity cluster.

To unregister the IBM DB2 source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the IBM DB2 source and select **Unregister**.

Backup IBM DB2 Database

Cohesity Platform provides a simple workflow to back up your IBM DB2 databases. After registering the IBM DB2 as a source on the Cohesity Platform, you can start configuring your backups.

Backup Process

Cohesity Platform provides a DB2 Connector for backup and restore of IBM DB2 deployments. The IBM DB2 Connector leverages the native backup capabilities provided by IBM DB2.

After registering the IBM DB2 source on the Cohesity Platform, you can create a Protection Group and specify the objects such as databases you want to backup.

Log Backup Methods

Cohesity provides the following two log backup methods:

- **Dump and sweep** - In this backup method:
  - IBM DB2 archives log files to local archive locations configured in the LOGARCHMETH1 option in the database configuration.

    - The Cohesity IBM DB2 connector copies log files from the local archive location to Cohesity storage during scheduled log backup execution.
    - Log files that are copied to Cohesity storage and older than "X" hours are pruned from the local archive location as a part of the scheduled log backup workflow. Where "X" is the prune hours value configured in the Protection Group.
- **Log Archival to Cohesity View using Vendor Library** - In this backup method,
    - The IBM DB2 Vendor library approach is used for log archival.
    - In the DB2 database configuration, the LOGARCHMETH2 option is configured to use the Cohesity DB2 vendor library to archive log files directly to Cohesity storage instead of the local archive location on DB2 nodes.
    - Logs from the Cohesity storage are pruned or deleted based on the retention period configured in the Protection Policy.

Considerations for Using the Log Backup Methods

Review and understand the following before you use the vendor library-based approach for direct log archival to Cohesity storage:

- The `LOGARCHMETH2` option in the database configuration is used to archive logs directly to the Cohesity view.
- During the first full backup for a database, Cohesity IBM DB2 Connector performs the following:
    - Configures `LOGARCHMETH2` and `LOGARCHOPT2` to utilize the Cohesity DB2 log archival library and Cohesity storage configurations.
    - Deactivates the database, performs an offline backup on the database, and then activates the database to change the `LOGATHMETH` option. This results in database downtime, as required by IBM DB2. For more information, see https://www.ibm.com/docs/en/db2/11.5?topic=parameters-logarchmeth2-secondary-log-archive-method.
    - The first full database backup will take extra time because of additional offline backup.

> **Note:** If you want to use only `LOGARCHMETH2` for log archival, make sure to manually turn off `LOGARCHMETH1`-based archival (if configured). The Cohesity IBM DB2 connector will not turn off `LOGARCHMETH1`-based archival.

Review and understand the following before you use `LOGARCHMETH1` and `LOGARCHMETH2` both on a database:

- In IBM DB2 nodes, you can configure `LOGARCHMETH1` and `LOGARCHMETH2` options in the database configuration to backup logs. If there is an archival failure in one of the

archive methods, IBM DB2 will continue to retry the log archival for failed files until the log files are successfully archived using both the methods.

- Using both archival methods on a DB2 database will impact Cohesity log backup or archival performance.

> **Important:**
> - Log Archival to Cohesity View using the vendor library backup method cannot be enabled on a registered source.
> - The backup jobs added before the release of version 7.1.1 will continue to utilize dump and sweep-based log backups.
> - To use the Log Archival to Cohesity View using the vendor library backup method, you must register a new source and enable the backup method during registration.

Offline Backup Best Practices

If you want to periodically perform an online and offline backup on the same database, then the following is the best practice workflow:

1. Create a Protection Group for online backup, for example, **DB1_online_backup**.
2. Create a different Protection Group for offline backup, for example, **DB1_offline_backup**.
3. Always keep the Protection Group, **DB1_offline_backup**, in the **Paused** state.
4. Run the Protection Group, **DB1_offline_backup**, as per the defined schedule in the protection policy.
5. When you want to perform an offline backup, pause the Protection Group for online backup, **DB1_online_backup**, and wait for the current protection run to finish before triggering the offline backup.
6. Perform the offline backup using the **RunNow** option, and bring the database online.
7. Unpause the Protection Group for online backup, **DB1_online_backup**.

Protect IBM DB2 deployments

In Cohesity Platform, IBM DB2 deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects such as databases that you want to backup.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when

creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect IBM DB2 deployments:

1. In the Cohesity Dashboard, do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the IBM DB2 source and then select **Protect**.
2. From the **Source** drop-down list, select the registered IBM DB2 source or click **Register Source** to register an IBM DB2 source. Click **+** to add multiple objects. For more information, see Register IBM DB2 as a Source.
3. In the **Objects** section, add the object that you plan to protect. The object should be a database name. To back up a database, enter the name of the database. For example, `DEMODB`. Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the **Backup Settings** section, select **Online Backup** or **Offline Backup** as per your requirements.

- If you have selected **Online Backup**, configure the following:

Field	Mandatory?	Default	Description
Convert Incremental Backup To Full Backup On failure	No	Disabled	Enable if you want to convert the incremental backup to full backup in case of an error. **Note:** Any failure of an incremental backup will trigger a full backup, increasing the total backup time.
Prune Logs	No	Disabled	Prune the backup history up to the last successful full backup to delete unnecessary entries from the history table. **Note:** Applicable to both the Log Backup Methods.
Prune Logs hour	No	N/A	Enter the number of hours. DB2 log files older than the provided hour will be deleted or pruned from the local archive location on DB2 nodes as a part of the scheduled log backup. **Note:** Applicable only to dump and sweep backup method. For more information, see Log Backup Methods.

- If you have selected **Offline Backup**, perform the following:

Field	Mandatory?	Default	Description
Convert Incremental Backup To Full Backup On failure	No	Disabled	Enable if you want to convert the incremental backup to full backup in case of an error.

Field	Mandatory?	Default	Description
Deactivate database	No	Disabled	Enable to deactivate the database before performing the offline backup.
Activate database	No	Disabled	Enable to activate the database after the backup is completed.

7. Optional. Enable **Advanced Options** to configure advanced settings:

Field	Mandatory?	Default	Description
Number of objects to backup in parallel	No	6	The number of objects you want to backup simultaneously.
RPC Timeout In Seconds	No	120 seconds	The request timeout between the IBM DB2 source and the Cohesity cluster.
Max IO Write Size To Cohesity	No	4194304 bytes	The maximum size of the data packet that can be sent to the Cohesity cluster.
Environment Variables	No	N/A	Provide the environment variables if you want to override the environment's default value. `For example, PATH=/sbin:/bin:/usr/bin`

8. In the **Settings** section, do the following:
   a. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
   b. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

      Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
   c. In the **Mounts** field, enter the number of VIPs you want to use while writing data to the Cohesity storage.

9. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
10. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the cluster. By default, maximum scaling is capped at 16.

Additional Settings	Description
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the IBM DB2 Backup Configuration

In the Cohesity Platform, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for IBM DB2.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ on the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for IBM DB2.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for IBM DB2.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Protection Group Only
   - Delete Protection Group and Snapshots

Recover IBM DB2 Database

You can restore the IBM DB2 object such as databases in IBM DB2. You can restore to the same IBM DB2 cluster or a different IBM DB2 cluster.

The Cohesity DB2 connector leverages the restore tool provided by IBM DB2. You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

> **Note:**
>
> If you plan to restore the backup to a different IBM DB2 cluster, ensure that you add the following registry variable to the environment variables of the target IBM DB2 instance using the following commands:

```
>db2stop
>db2set DB2_RESTORE_GRANT_ADMIN_AUTHORITIES=ON
>db2start
```

Redirected Restore

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Cohesity Platform supports redirected restore operation provided by IBM DB2. Using the redirected restore operation, you can perform the following while restoring the IBM DB2 backup:

- Replace the same IBM DB2 cluster using the existing storage groups and tablespace locations
- Redirect all tablespaces to new locations
- Redirect the subset of tablespaces to new locations
- Redirect all storage groups to new locations
- Redirect the subset of storage groups to new locations

Before performing the redirected restore operation using the Cohesity Platform, you must create a script with the redirection statements and provide the path of the script in the **Redirection script path** field in the IBM DB2 restore configuration. The script is replayed at the time of redirection for storage group and tablespaces containers.

> **Note:** The script should only contain redirection statements, and the restore statements must not be part of the script.

The following is a sample of the script:

```
set stogroup paths for DATA_SG on '/var/ibmdb2/DATABASE1/db1.DMSFILETSP' ,
'/var/ibmdb2/DATABASE1/db1.DMSFILETSP2' , '/var/ib-
mdb2/DATABASE1/db1.DMSFILETSP3';
set stogroup paths for TEMP_SG on '/var/ibmdb2/DB1/db1.TMPFILETSP' ;
set tablespace containers for 0 using ( path
```

```
'/var/ibmdb2/DATABASE1/NODE0000/BACKTST2/SYSCATSPACE' );
set tablespace containers for 3 using automatic storage;
```

Recover IBM DB2 Databases

To restore IBM DB2 databases:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *****.
4. From the list of search results, select the Protection Group that contains the snapshots to recover. By default, the latest recovery point is selected. To select a different recovery point, from the selected lists, mouse over the recovery point, and then click the **edit** icon.

   By default, the option to select the recovery point is displayed in **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page. The **Timeline** mode displays the full backups (recovery points) captured at the scheduled time and the CDP-based data available for point-in-time restore. The **List** tab displays only the full backups captured at a scheduled time.

   - To select a recovery point from the Timeline mode:
     a. From the **Choose a date** drop-down, select a date.

        The recovery points available on the selected date are displayed in blue dots on the timeline bar and the point-in-time for restore is displayed in a solid green line on the timeline bar.

     b. To recover from a full backup, click the corresponding blue dot. For point-in-time restore, specify the exact point-in-time recovery by sliding the timeline bar.

     c. Click **Select Recovery Point** and then go to **step 5**.
   - To select a recovery point from the List mode:
     a. Select the **List** tab from the top-right corner of the *Edit recovery point for <database_name>* page.

        By default, the page displays the recovery points available for the last 30 days.

     b. You can also select a recovery point from a different date range by selecting the date range from the Date drop-down list. The page displays the available recovery points in the specified range.

     c. Select the recovery points to be recovered.

     d. Click **Select Recovery Point** and then go to **step 5**.

5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. Click **+** to add multiple objects.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same IBM DB2 cluster.
New Object	To restore to a different IBM DB2 cluster. Select the **Target** IBM DB2 cluster to which the data needs to be restored. Also, enable **Overwrite existing object with the same name** and type **Yes**, if you want to overwrite the existing object with the same name. > **Note:** Ensure that the **Target** IBM DB2 cluster is registered as a source on the Cohesity cluster.

8. In the **Restore Settings** section, select **Regular Restore** or **Redirected Restore** based on your requirements.

- If you have selected **Regular Restore**, configure the following:

Field	Mandatory?	Default	Description
New Data Path	No	N/A	Enter the path on the target IBM DB2 instance to restore the data directory from the database. You can specify the path(s) as a comma-separated list of directories. **Note:** If the path is not specified, the data is restored to the default database path, DFTDBPATH. You can view the default database path, DFTDBPATH, using the following command: `db2 get dbm cfg | grep DFTDBPATH`
New Log Path	No	N/A	Enter the path on the target IBM DB2 instance to restore the active logs from the restored database.
Local Database Directory	No	N/A	Enter the path of the local database directory.
Rollforward Database	No	Enabled	Roll forward the database to recover the database by applying transaction logs after restore.

- If you have selected **Redirected Restore**, configure the following:

Field	Mandatory?	Default	Description
New Data Path	No	N/A	Enter the path on the target IBM DB2 instance to restore the data directory from the database. You can specify the path(s) as a comma-separated list of directories. **Note:** If the path is not specified, the data is restored to the default database path, DFTDBPATH. You can view the default database path, DFTDBPATH, using the following command: `db2 get dbm cfg | grep DFTDBPATH`
Redirection Script Path	Yes	N/A	Enter the path of the script with the redirection statements. For more information, see Redirected Restore.
New Log Path	Yes	N/A	Enter the path on the target IBM DB2 instance to restore the active logs from the restored database.
Local Database Directory	No	N/A	Enter the path of the local database directory.
Rollforward Database	No	Enabled	Roll forward the database to recover the database by applying transaction logs after restore.

9. Optional. Enable **Advanced Options** to configure advanced settings:

Field	Mandatory?	Default	Description
RPC Timeout In Seconds	No	120 seconds	The request timeout between the IBM DB2 source and the Cohesity cluster.

Field	Mandatory?	Default	Description
Max IO Read Size To Cohesity	No	4194304 bytes	The maximum size of the data packet that can be transferred to the Cohesity cluster.
Environment Variables	No	N/A	Provide the environment variables if you want to override the environment's default value. For example, `PATH=/sbin:/bin:/usr/bin`

10. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	Not applicable for IBM DB2 cluster database recovery.
Rename	Allows you to rename the object, if you are recovering the object to **New Location**. **Note:** If you are overwriting the existing object, do not rename the object using the database alias name.
Task Name	Change the default name of the recovery task.

11. Click **Recover**.

Cancel a Recovery Run

When you cancel a recovery run, the IBM DB2 connector cancels the job by terminating the recovery processes.

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click the action icon (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
   - **Continue Recovery** if you want to continue to recovery run
   - **Cancel Recovery** if you want to cancel the recovery run.

MySQL EE

Cohesity natively integrates with MySQL to provide backup and recovery for MySQL deployments. Using Cohesity Connector Agent, you can backup and restore the MySQL database. The Cohesity Connector Agent provides a simple workflow to back up and restore your MySQL deployment.

The MySQL Connector enables you to back up and restore the following:

- Full and incremental backup of a MySQL database.
- Physical level restore (No option to restore on a per-database level).

Core Components

The following are the core components for protecting MySQL databases:

- MySQL Linux machine
- MySQL Connector
- Cohesity Connector Agent
- Cohesity View

As illustrated, you can integrate Cohesity with your MySQL deployments using the MySQL Connector and Cohesity Connector Agent.

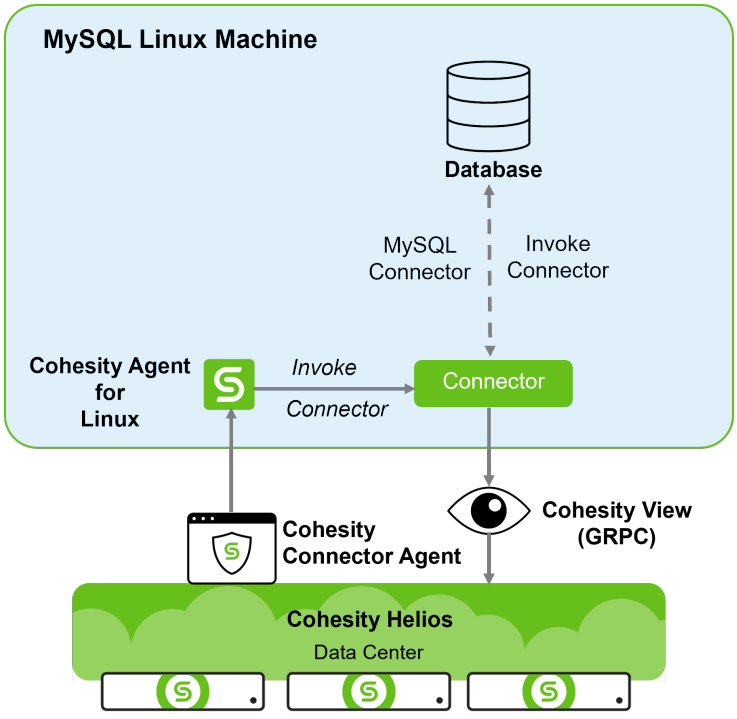

Plan and Prepare for MySQL EE Protection

Before you register your MySQL deployment as a source with Cohesity and protect MySQL instances, ensure the following prerequisites:

- Ensure that the fully qualified domain name (FQDN) of the Cohesity cluster is resolvable from the MySQL host.
- A database user with database admin user privileges. (Only MySQL Enterprise Edition is supported. Ensure mysqlbackup is installed along with the MySQL database.)
- Download and install the Cohesity Linux Agent on the MySQL host.
- Download and install the MySQL Connector.

Supported Versions

For more information on supported MySQL versions, see Oracle MySQL versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	MySQL Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
MySQL Local Host	MySQL Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
MySQL Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
MySQL Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following before you protect your MySQL database.

- Only RHEL is supported.
- Database instance auto-discovery and object search are not supported.
- You cannot protect multiple hosts using a single Protection Group.
- Multiple object restore is not supported.
- Only instance-level backup and restore are supported. Granular-level use cases such as database-level backup and restore are not supported.
- Log backups and PITR recovery are supported.
- Recovery is performed only at the instance level; hence, you cannot rename the database after recovery.
- Only single-stream backup and restore are supported.
- Only backup and restore of a single MySQL database on the source machine is supported. If the source machine has multiple database instances, you cannot backup and restore the database.
- Recovery from a replication copy of the database is not supported.

Download and install the Cohesity Linux Agent

The Cohesity Linux Agent is available as script-based and RPM-based installers, based on your requirements download the installer and install it accordingly on the MySQL host. For detailed instructions, see Download and Install the Linux Agent.

> **Note:** Ensure you meet the following prerequisites below before you install the Cohesity Linux agent.

Prerequisites

- For Cohesity clusters to be able to interact with the Linux agent, ports 50051 and 59999 need to be opened up on the agent host.
- On the Cohesity cluster, enable UDA on the Linux agent (if agent older than 6.6.0) and restart the agent from the cluster.

Download the MySQL Connector

Cohesity provides a MySQL Connector to backup and restore MySQL databases. The MySQL Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for MySQL Connector, and copy the datasource agent to the MySQL host.

Install the MySQL Connector

Before you install the MySQL Connector, ensure that you have sudo privileges.

To install the MySQL Connector:

1. Log on to the MySQL EE host where you have copied the MySQL Connector installer, `cohesity-mysql-java-agent-<version>.x86_64.rpm`.
2. To set the service user to run the installer, run the following command. If not set, the default value will be `cohesityagent`.

   ```
   export SERVICE_USER=user1
   ```

3. Run the following command to install the MySQL Connector:

   ```
   rpm -ivh cohesity-mysql-java-agent-<version> .x86_64.rpm
   ```

After the MySQL Connector is installed, the agent is available at: `/opt/cohesity/agent/uda_scripts`

Uninstall the MySQL Connector

To uninstall the MySQL Connector:

Log on to the MySQL host where you have installed the MySQL Connector. Run the following command to uninstall the MySQL Connector:

```
rpm -e cohesity-mysql-java-agent-<verions>. x86_64
```

Register and Manage the MySQL Source

To protect your MySQL databases, you need to register your MySQL deployment as a source on Cohesity.

After registering your MySQL deployment as a source, you can:

- Update the username and password of the database user.
- Refresh the source details.
- Unregister the MySQL source from Cohesity.

Register MySQL as a Source

Before you register MySQL as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for MySQL EE Protection.

To register MySQL as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:

1. From the **Source Type** drop-down menu, select **MySQL**.
2. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the MySQL host you have identified to run the MySQL Connector. For more information, see Plan and Prepare for MySQL EE Protection.
3. In the **Datasource Agent Installation Path**, enter the directory path on the MySQL host, where you have installed the MySQL Connector. For more information, see Install the MySQL Connector.
4. In the **APP Authentication** section, provide the MySQL database's **Username** and **Password**.
5. Enable **Custom Options**, and perform the following:
    1. Ignore the **Mounts View** option. It is not applicable for MySQL source registration.
    2. In the **Source Registration Options** field, enter the following source registration arguments based on your requirement:

Arguments	Mandatory ?	Description
`source_name`	Yes	A unique name to identify the MySQL source.

Arguments	Mandatory ?	Description
`defaults_file`	Yes	The MySQL's configuration file path (typically `/etc/my.cnf`).

Arguments	Mandatory ?	Description

		Note: Ensure that the `[client]` section is present in the `my.cnf` file and includes the same socket path as defined in the `[mysqld]` section. This is required for proper socket-based communication between the MySQL client and server.
		Example:
		`[mysqld]`
		`datadir=/data/mysql/mysql-8.4`
		`socket=/data/mysql/mysql-8.4/mysql.sock`

```
source_name=<SOURCE_NAME>, defaults_file=/etc/my.cnf
```

6. Click **Register**.

Update the MySQL Source Configuration

After registering your MySQL as a source on Cohesity, you can:

- Update the username and password of the database user.
- Update the **Datasource Agent Installation Path** only if you have uninstalled and reinstalled the MySQL Connector in a different directory on the MySQL host.

To update the MySQL source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the MySQL source and select **Edit**.

> **Note:** When you edit the source, the username and password fields are reset. Ensure that you re-enter the credentials before clicking **Update**.

3. In the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the MySQL Source Configuration

You can refresh the MySQL source and re-verify the source with the available source registration parameters.

To refresh the MySQL Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the MySQL source and select **Refresh**.

Unregister the MySQL Source

If you plan to stop taking the backup of your MySQL databases, you can unregister the MySQL source from Cohesity.

> **Note:** To unregister a source, ensure that no protection jobs are associated with it. The unregistration will fail if active protection jobs exist for the source.

To unregister the MySQL source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the MySQL source and select **Unregister**.

Backup MySQL Database

Cohesity provides a simple workflow to back up your MySQL databases. After registering MySQL as a source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides a MySQL EE connector for backup and restores of enterprise MySQL deployments. The MySQL EE connector leverages the backup capabilities provided by MySQL.

After registering the MySQL source on Cohesity, you can create a Protection Group and specify the MySQL instances that you want to back up. In the Protection Group, you need to define the script parameters that enable you to provide additional parameters required by the backup statements provided by MySQL.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your MySQL backup configuration: (parameters are to be passed comma(,) separated)

Backup Parameters	Description
`number-of-buffers`	Optional. Specify the number of buffers to use during multithreaded operations. you can use a high number of buffers for CPU-intensive processing such as backup, particularly when using compression, and use a low number of buffers for disk-intensive processing such as restoring a backup. This value should be at least as high as the number of read threads or write threads, depending on the type of operation. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.
`read-thread`	Optional. Specify the number of threads to use for reading data from the disk. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.

Backup Parameters	Description
`process-threads`	Optional. Specify the number of threads to use for processing data. This includes compression and uncompression, encryption and decryption, apply-log operations, and packing and extracting backup images. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.
`write-threads`	Optional. Specify the number of threads to use for writing data to disk. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.
`limit-memory`	Optional. Specify maximum memory in megabytes that can be used by mysqlbackup. **Note:** Do not include any suffixes such as MB or KB in the option value. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.

Example

```
number-of-buffers=10, read-threads=5, process-threads=5, write-threads=5,
limit-memory=500
```

Protect MySQL EE Deployments

MySQL deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to:

- Specify the MySQL instance that you want to backup.
- Based on your requirement, define the parameters for the backup statements provided by MySQL. For more information, see Backup Parameters.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect MySQL deployments:

1. Do one of the following:
    1. Select **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
    2. Select **Data Protection** > **Source**. Click the actions menu (⋮) next to the MySQL source and then select **Protect**.
2. From the **Source** drop-down list, select the registered MySQL source or click **Register Source** to register a MySQL source. For more information, see Register and Manage the MySQL Source.
3. In the **Objects** section, specify the MySQL instance name to protect.

    > **Note:** Only instance-level backup and restore are currently supported. Granular-level operations, such as backing up or restoring individual databases or tables, are not supported. As a result, you need to enter a placeholder value in the **Objects** field to continue.

    Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:
    1. Select **New Group** to create a new Protection Group. In the Protection Name field, enter a name from the Protection Group.
    2. Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click Protect.
5. In the **Settings** section, do the following:
    1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
    2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

        Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
6. After you have completed the **Settings** configuration, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click **Edit** on the right. For more information, see **Additional Settings** below.
7. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts before this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	Optional. Enter a brief description of the Protection Group.

Manage the MySQL Backup Deployments

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups, local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details will not be displayed immediately. It will take some time to reflect on Cohesity.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select Delete.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    1. Delete Protection Group Only
    2. Delete Protection Group and Snapshots.

Recover MySQL Database

You can restore the MySQL host or object(s) such as databases in MySQL. You can restore to the same MySQL host or a different MySQL host.

The MySQL EE connector leverages the restore statements provided by MySQL. You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

In the Recovery configuration, you need to specify the object you plan to restore and define the script parameters required by the restore statements provided by MySQL.

Recovery Parameters

In the Recovery configuration, you can define the following script parameters for your MySQL recovery:

> **Note:** You can use the optional recovery parameters based on your requirements.

Recovery Parameters	Mandatory?	Description
`data_dir`	Yes	Specify the path where the MySQL data directory will be restored. The data directory contains all information managed by the MySQL server. **Note:** Ensure that the specified directory is empty. If it is not empty, Cohesity will automatically delete its contents during the restore operation.
`restart_instance`	No	Specifies whether the recovered instance automatically starts after the recovery completes. • If set to `true`, the instance restarts once recovery completes. • If set to `false`, the instance remains powered off. Default value: `true`.
`number-of-buffers`	No	Specify the number of buffers to use during multithreaded operations. you can use a high number of buffers for CPU-intensive processing such as backup, particularly when using compression, and use a low number of buffers for disk-intensive processing such as restoring a backup. This value should be at least as high as the number of read threads or write threads, depending on the type of operation. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.

Recovery Parameters	Mandatory?	Description
`read-thread`	No	Specify the number of threads to use for reading data from the disk. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.
`process-threads`	No	Specify the number of threads to use for processing data. This includes compression and uncompression, encryption and decryption, apply-log operations, and packing and extracting backup images. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.
`write-threads`	No	Specify the number of threads to use for writing data to disk. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.
`limit-memory`	No	Specify maximum memory in megabytes that can be used by mysqlbackup. **Note:** Do not include any suffixes such as MB or KB in the option value. For more information, see MySQL Enterprise Backup 8.0 User's Guide :: 20.10 Performance / Scalability / Capacity Options.

Example

```
data_dir=/var/lib/mysql_restore, restart_instance=false, number-of-buf-
fers=10, read-threads=5, process-threads=5, write-threads=5, limit-memory-
y=500
```

Recover MySQL

To recover a MySQL source:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, select **Replace Original**.

Options	Description
Replace Original	To restore to the same MySQL host.
New Object	To restore to a different MySQL host. Select the **Target** MySQL host to which the data needs to be restored. Also, ignore **Overwrite existing object with the same name toggle** option. It is not applicable for MySQL.

8. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Not applicable for MySQL host recovery.
Custom Options	Specify the Recovery Parameters based on your requirements. See Sample Recovery Use Cases for examples.
Rename	Provide the new name for the database, if you want to restore the database with a new name.
Task Name	Change the default name of the recovery task.

9. Click **Recover**.

Sample Recovery Use Cases

Use Case	Instructions with Sample Recovery Parameters
Restore the MySQL host to Same Location	`data_dir=/var/lib/mysql_restore, restart_instance=true` **Note:** Provide an empty directory as an parameter in the `data_dir` parameter when restoring data to the same location to avoid overwriting of data in the restore data directory.

Cancel a Recovery Run

To cancel a recovery run follow the steps below:

1. Navigate to **Data Protection > Recoveries.**
2. Click the action menu (⋮) next to the recovery run and click **Cancel Run**.
3. In the Cancel Recovery dialog box, select any one of the following:
    1. Click **Continue Recovery** if you want to continue a recovery run
    2. Click **Cancel Recovery** if you want to cancel the recovery run.

Percona MySQL on Linux

Cohesity natively integrates with Percona MySQL to provide backup and recovery for Percona MySQL deployments. Using Cohesity Connector Agent, you can backup and restore the Percona MySQL databases.

Cohesity Platform provides a MySQL connector that enables you to perform full backup, incremental backup, log backup, restore and point in time restore at the instance level.

Core Components

The following are the core components for protecting Percona MySQL databases:

- Percona MySQL Linux machine
- MySQL Connector
- Cohesity Connector Agent
- Cohesity View

As illustrated, you can integrate Cohesity with your Percona MySQL deployments using the MySQL Connector and Cohesity Connector Agent.

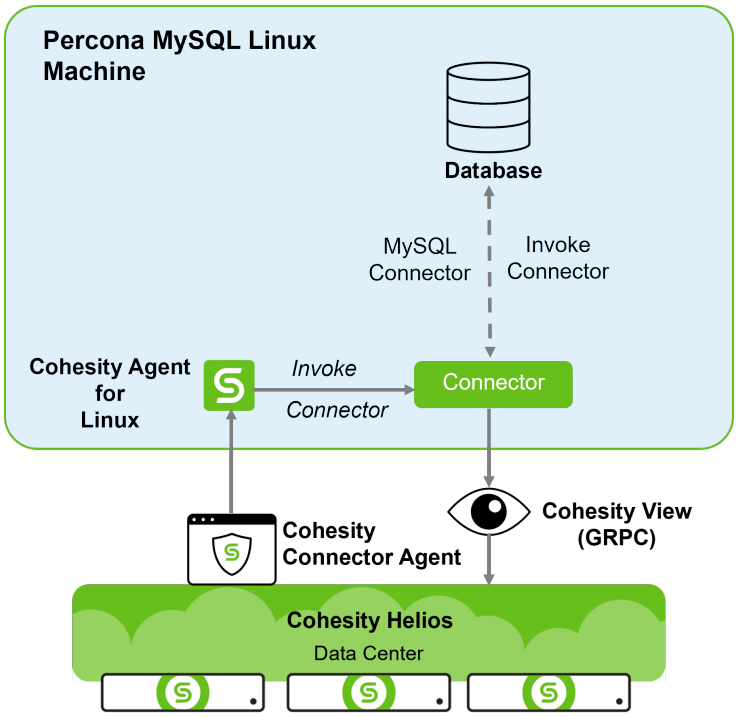

Plan and Prepare for Percona MySQL Protection

Before you register your Percona MySQL deployment as a source with Cohesity and protect Percona MySQL instances, ensure the following prerequisites:

- Ensure that the fully qualified domain name (FQDN) of the Cohesity cluster is resolvable from the MySQL host.
- A database user with database admin user privileges.
- Ensure that the `xtrabackup` and `mysqlbinlog` utilities are installed on the host along with the MySQL database.
- Download and install the Cohesity Linux Agent on the MySQL host.
- Download and install the MySQL connector.

Supported Versions

For more information on supported Percona MySQL versions, see Percona MySQL versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	Percona MySQL Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
Percona MySQL Local Host	Percona MySQL Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
Percona MySQL Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
Percona MySQL Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following before you protect your Percona MySQL database.

- Only RHEL is supported.
- Database instance auto-discovery and object search are not supported.
- You cannot protect multiple hosts using a single Protection Group.
- Multiple object restore is not supported.
- MySQL Connector upgrade is not supported. To use the latest version of the MySQL Connector, uninstall the existing MySQL Connector and install the latest version.
- Only instance-level backup and restore are supported. Granular-level use cases such as database-level backup and restore are not supported.
- Log backups and PITR recovery are supported.
- Recovery is performed only at the instance level; hence, you cannot rename the database after recovery.
- Only backup and restore of a single Percona MySQL database on the source machine is supported. If the source machine has multiple database instances, you cannot backup and restore the database.
- Recovery from a replication copy of the database is not supported.

Download and install the Cohesity Linux Agent

The Cohesity Linux Agent is available as script-based and RPM-based installers, based on your requirements download the installer and install it accordingly on the Percona MySQL host. For detailed instructions, see Download and Install the Linux Agent.

> **Note:** Ensure you meet the following prerequisites below before you install the Cohesity Linux agent.

Prerequisites

- For Cohesity clusters to be able to interact with the Linux agent, ports 50051 and 59999 need to be opened up on the agent host.
- On the Cohesity cluster, enable UDA on the Linux agent (if agent older than 6.6.0) and restart the agent from the cluster.

Download the MySQL Connector

Cohesity provides a MySQL Connector to backup and restore Percona MySQL databases. The MySQL Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for MySQL Connector, and copy the datasource agent to the Percona MySQL host.

Install the MySQL Connector

Before you install the MySQL Connector, ensure that you have sudo privileges.

To install the MySQL Connector:

1. Log on to the Percona MySQL host where you have copied the MySQL Connector installer, `cohesity-mysql-java-agent-<version>.x86_64.rpm`.
2. To set the service user to run the installer, run the following command. If not set, the default value will be `cohesityagent`.

   ```
   export SERVICE_USER=user1
   ```

3. Run the following command to install the MySQL Connector:

   ```
   rpm -ivh cohesity-mysql-java-agent-<version> .x86_64.rpm
   ```

After the MySQL Connector is installed, the agent is available at:
`/opt/cohesity/agent/uda_scripts`

Uninstall the MySQL Connector

To uninstall the MySQL Connector:

Log on to the Percona MySQL host where you have installed the MySQL Connector. Run the following command to uninstall the MySQL Connector:

```
rpm -e cohesity-mysql-java-agent-<verions>. x86_64
```

Register and Manage the Percona MySQL Source

To protect your Percona MySQL databases, you need to register your Percona MySQL deployment as a source on Cohesity.

After registering your Percona MySQL deployment as a source, you can:

- Update the username and password of the database user.
- Refresh the source details.
- Unregister the Percona MySQL source from Cohesity.

Register Percona MySQL as a Source

Before you register Percona MySQL as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for Percona MySQL Protection.

To register Percona MySQL as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
    a. From the **Source Type** drop-down menu, select **MySQL**.
    b. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the Percona MySQL host you have identified to run the MySQL Connector. For more information, see Plan and Prepare for Percona MySQL Protection.
    c. In the **Datasource Agent Installation Path**, enter the directory path on the Percona MySQL host, where you have installed the MySQL Connector. For more information, see Install the MySQL Connector.
    d. In the **APP Authentication** section, provide the Percona MySQL database's **Username** and **Password**.
    e. Enable **Custom Options**, and perform the following:
        1. Ignore the **Mounts View** option. It is not applicable for Percona MySQL source registration.
        2. In the **Source Registration Options** field, enter the following source registration arguments based on your requirement:

Arguments	Mandatory ?	Description
`--source-name`	Yes	A unique name to identify the Percona MySQL source.

Arguments	Mandatory ?	Description
`--defaults-file`	Yes	The Percona MySQL's configuration file path (typically `/etc/my.cnf`).

Arguments	Mandatory ?	Description

		Note: Ensure that the `[client]` section is present in the `my.cnf` file and includes the same socket path as defined in the `[mysqld]` section. This is required for proper socket-based communication between the MySQL client and server.
		Example:
		`[mysqld]`
		`datadir=/data/mysql/mysql-8.4`
		`socket=/data/mysql/mysql-8.4/mysql.sock`

Arguments	Mandatory ?	Description
`--port`	No	The port used to connect to the Percona MySQL source. If not specified, the default port **3306** is used.

Sample:

```
--source-name=test_source --defaults-file=/etc/my.cnf --
port=3306
```

f. Click **Register**.

Update the Percona MySQL Source Configuration

After registering your Percona MySQL as a source on Cohesity, you can:

- Update the username and password of the database user.
- Update the **Datasource Agent Installation Path** only if you have uninstalled and reinstalled the MySQL Connector in a different directory on the Percona MySQL host.

To update the Percona MySQL source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Percona MySQL source and select **Edit**.

> **Note:** When you edit the source, the username and password fields are reset. Ensure that you re-enter the credentials before clicking **Update**.

3. In the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the Percona MySQL Source Configuration

You can refresh the Percona MySQL source and re-verify the source with the available source registration parameters.

To refresh the Percona MySQL Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Percona MySQL source and select **Refresh**.

Unregister the Percona MySQL Source

If you plan to stop taking the backup of your Percona MySQL databases, you can unregister the Percona MySQL source from Cohesity.

> **Note:** To unregister a source, ensure that no protection jobs are associated with it. The unregistration will fail if active protection jobs exist for the source.

To unregister the Percona MySQL source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Percona MySQL source and select **Unregister**.

Backup Percona MySQL Database

Cohesity provides a simple workflow to back up your Percona MySQL deployment. After registering the Percona MySQL as a source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides a MySQL connector to protect Percona MySQL deployments. The MySQL Connector uses the XtraBackup FIFO data sink method from Percona to stream data efficiently during backup.

When a protection job runs, multiple parallel threads write database data into a set of FIFO (First-In-First-Out) pipes on the source system. These FIFO pipes act as streaming interfaces. The Cohesity cluster then reads the data from the pipes and transfers it directly to the Cohesity cluster for protection.

Backup Scope

- Backups are supported at the instance level.
- If the instance contains multiple databases, all databases are included in a single backup workflow.
- Selective database-level backups are not supported.

Backup Types Supported

Cohesity supports the following types of backups:

- **Full Backup**

  Captures a complete copy of the entire MySQL instance, including all databases, tables, indexes, tablespaces, and logs.
- **Incremental Backup**

Captures only the changes made since the last successful full or incremental backup on the Cohesity cluster. **XtraBackup** performs incremental backups based on the last redo log position of the previous backup.

- **Log Backup**

  Captures the binary logs (`binlogs`), which contain all database changes (`INSERT`, `UPDATE`, `DELETE`) with associated timestamps. Log backups enable point-in-time recovery by replaying changes after a full or incremental backup.

Backup Parameters

In the Protection Group, you can define the following parameters for your Percona MySQL backup configuration: (parameters are to be passed space (" ") separated)

> **Note:** The backup parameters are not applicable for log backups.

Backup Parameters	Mandatory?	Description
`--fifo-streams`	No	Number of FIFO streams used for backup. Default: **6**.
`--parallel-threads`	No	Number of parallel threads writing to the FIFO streams. Default: **18**.

Example

```
--fifo-streams=8 --parallel-threads=24
```

Protect Percona MySQL Deployments

In Cohesity Platform, Percona MySQL deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to:

- Specify the MySQL instance that you want to backup.
- Based on your requirement, define the parameters for the backup statements provided by MySQL. For more information, see Backup Parameters.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect Percona MySQL deployments:

1. Do one of the following:
   a. Select **Data Protection** > **Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   b. Select **Data Protection** > **Source**. Click the actions menu (⋮) next to the Percona MySQL source and then select **Protect**.
2. From the **Source** drop-down list, select the registered Percona MySQL source or click **Register Source** to register a Percona MySQL source. For more information, see Register and Manage the Percona MySQL Source.
3. In the **Objects** section, specify the object you want to protect.

   > **Note:** Only instance-level backup and restore are supported. Granular operations, such as backing up or restoring individual databases or tables, are not available. To proceed, you must enter a placeholder value in the **Objects** field.

   Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:
   a. Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   b. Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the **Settings** section, do the following:
   a. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
   b. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

      Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
   c. In the **Custom Options** sections, perform the following:

      Provide the backup parameters in the **Full Backup** and **Incremental Backup** fields if you plan to back up the entire Percona MySQL host. The parameters for

full and incremental backups can differ, but ensure that the value of `--fifo-streams` is always less than the value of `--parallel-threads`.

> **Note:** The **Log Backup** field is optional. Although log backups are supported, any value entered in this field is not used during backup operations.

For example,

Use Cases	Sample Parameters
Create a full backup of the Percona MySQL host	`--fifo-streams=8 --parallel-threads=24`
Create an incremental backup of the Percona MySQL host	`--fifo-streams=4 --parallel-threads=12`

Ignore the **Mounts** field. It is not applicable for Percona MySQL data protection.

7. After you have completed the **Settings** configuration, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click **Edit** on the right. For more information, see Additional Settings below.
8. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts before this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.

Additional Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	Optional. Enter a brief description of the Protection Group.

Manage the Percona MySQL Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for Percona MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Percona MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups, local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details will not be displayed immediately. It will take some time to reflect on Cohesity.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Percona MySQL.
3. Click the actions menu (⋮) next to the Protection Group and select Delete.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    1. Delete Protection Group Only
    2. Delete Protection Group and Snapshots.

Recover Percona MySQL Database

You can restore the protected Percona MySQL data to the same Percona MySQL deployment or to a different Percona MySQL deployment.

Recovery Process

Cohesity Platform supports instance-level recovery for Percona MySQL using the MySQL connector. If multiple databases were included in the backup, all databases are restored together within the same recovery workflow.

Recovery Options

Cohesity provides the following restore options:

- **Full or Incremental Image Restore**

  Restores the database instance using a full backup or a combination of full and incremental backups. Recovery is performed to the state captured in the selected backup image(s).

- **Point-in-Time Restore (PITR)**

  Restores the database instance to an exact point in time defined by the user. PITR combines full backups, incremental backups, and binary logs (binlogs) to replay changes up to the requested timestamp, providing granular recovery.

Recovery Parameters

In the Recovery configuration, you can define the following script parameters for your Percona MySQL recovery:

> **Note:** You can use the optional recovery parameters based on your requirements.

Recovery Parameters	Mandatory?	Description
`--target-dir`	Yes	Specifies the path where the MySQL data directory will be restored. The data directory contains all information managed by the Percona MySQL server. **Note:** Ensure that the directory path provided in the **target-dir** parameter already exists and is empty. The restore workflow will fail if the directory does not exist or contains data.
`--restart-instance`	No	Set to **false** if you do not want the Percona MySQL instance to restart automatically after recovery. The default value is **true**.

Example

```
--target-dir=/var/lib/mysql_restore --restart-instance=true
```

Recover Percona MySQL

To recover a Percona MySQL source:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, select **Replace Original**.

Options	Description
Replace Original	To restore to the same Percona MySQL host.
New Object	Select this option to restore data to a different MySQL host. Choose the **Target MySQL** host where you want to restore the data. **Note:** The **Overwrite existing object with the same name** toggle is not applicable for MySQL and can be ignored.

8. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Not applicable for Percona MySQL host recovery.
Custom Options	Specify the Recovery Parameters based on your requirements. See Sample Recovery Use Cases for examples.
Rename	Not applicable for Percona MySQL. This field can be ignored.
Task Name	Change the default name of the recovery task.

9. Click **Recover**.

Sample Recovery Use Cases

Use Case	Instructions with Sample Recovery Parameters
Restore the Percona MySQL host to Same Location	`--target-dir=/var/lib/mysql_restore --restart-instance=true` **Note:** Ensure that the directory path provided in the **target-dir** parameter already exists and is empty. The restore workflow will fail if the directory does not exist or contains data.

Cancel a Recovery Run

To cancel a recovery run follow the steps below:

1. Navigate to **Data Protection > Recoveries.**
2. Click the action menu (⋮) next to the recovery run and click **Cancel Run**.
3. In the Cancel Recovery dialog box, select any one of the following:
    1. Click **Continue Recovery** if you want to continue a recovery run
    2. Click **Cancel Recovery** if you want to cancel the recovery run.

SAP HANA

Cohesity natively integrates with SAP HANA to provide backup and recovery for clustered SAP HANA deployments. Using Cohesity Connector Agent, you can backup and restore the SAP HANA database. The Cohesity Connector Agent provides a simple workflow to back up and restore your SAP HANA deployment.

> **Tip:** To protect SAP HANA database using the Cohesity BACKINT plug-in Remote Adapter Jobs, see About SAP HANA Database Backups.

Cohesity provides a SAP HANA Connector for backup and restore of enterprise SAP HANA deployments. The SAP HANA Connector leverages the backup and restore capabilities provided by SAP HANA.

The SAP HANA Connector enables you to back up and restore both TenantDB and SystemDB. It provides:

- Backup:
    - Full backup
    - Incremental backup
    - Differential backup
    - Automatic Log backup
- Restore:
    - Restore the database to its most recent state.
    - Restore the database to the point in time.
    - Restore the database to a specific data backup.
    - Alternate restore of SAP HANA database
        - Alternate restore to new database or active database on the same SAP HANA system.
        - Alternate restore to a new or active database on a different SAP HANA system.

Core Components

The following are the core components for protecting SAP HANA databases:

- SAP HANA Linux machine
- SAP HANA Connector
- Cohesity Connector Agent
- Cohesity View

As illustrated, you can integrate the Cohesity cluster with your SAP HANA deployments using the SAP HANA Connector and Cohesity Connector Agent.

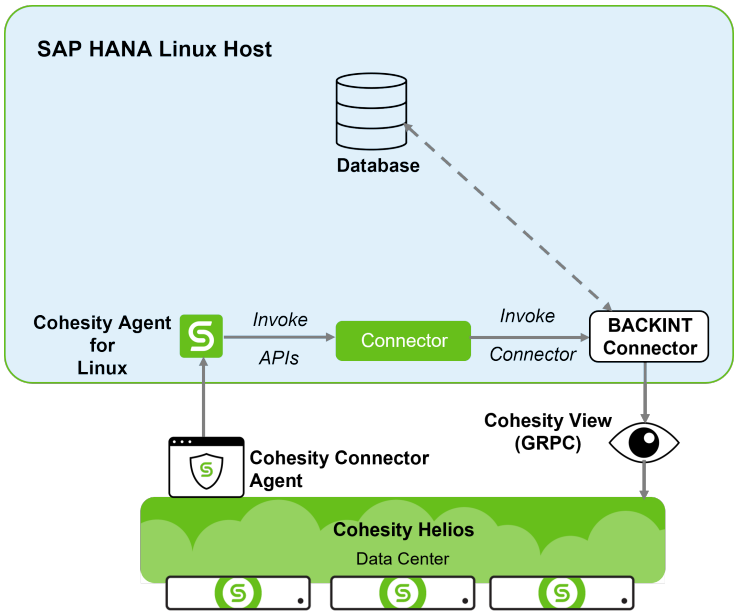

Plan and Prepare for SAP HANA Protection

Before you register your SAP HANA deployment as a source with Cohesity and protect SAP HANA databases, ensure the following prerequisites:

- Identify the node(s) in your SAP HANA cluster with adequate resources to execute the SAP HANA connector.
- Ensure that the fully qualified domain name (FQDN) of the Cohesity cluster is resolvable from all the nodes in the SAP HANA cluster.
- Synchronize the SAP HANA Linux machine time with a Network Time Protocol (NTP) server. For detailed instructions, see How to Sync Linux Server Time with an NTP Server.
- Ensure that the NTP time is in sync between the client and the Cohesity cluster. Perform the following on the client:

- Verify if NTP is enabled and synchronized in the client using the following command:

```
timedatectl
```

- If the NTP time is not synchronized, use the following command to synchronize the NTP time

```
timedatectl set-ntp true
```

- Use "bash" shell to install the SAP HANA Connector and to run backups and restore workflows of SAP HANA databases.
- A database user with database admin user privileges.
- Download and install the Cohesity Linux Agent on the SAP HANA node.
- Download and install the SAP HANA connector.

Supported Versions

For more information on supported SAP HANA versions, see Supported Software versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	SAP HANA Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP HANA Local Host	SAP HANA Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP HANA Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
SAP HANA Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following before you protect your SAP HANA database:

- SAP HANA 1.0 is not supported.
- Source-side deduplication is not supported for SAP HANA on PowerPC platform.
- You cannot protect multiple hosts using a single Protection Group.
- You cannot restore multiple objects, search and recover an object.
- Scheduling log backups in the Cohesity cluster is not supported, as SAP HANA automatically performs log backups.
- Cohesity supports backup and restore of SAP HANA(x86) databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.
- Cohesity supports the recovery of SAP HANA 2.0 backups across both Intel x86-64 and Power PC platforms, allowing you to restore your critical SAP HANA data regardless of the underlying hardware architecture.
- Restoring backups from a higher version of SAP HANA to a lower version is not supported.
- Cohesity supports SAP HANA scale-out deployments that use independent disks for each node.
- Only Pacemaker-based cluster stacks provided by RHEL and SLES are supported for SAP HANA HA cluster data protection.

  > **Note:** The Pacemaker version 0.9.167 is supported.

- It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.

  From 7.2.2 version onwards, it is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
- For a secure connection between the Cohesity cluster and the SAP HANA database, you need to generate an SAP HANA certificate for authentication.

  For a secure connection between the Cohesity cluster and the SAP HANA database, Cohesity cluster certificate is required for authentication. The certificate is auto-generated only after you run the first database UDA protection job. Post running the job, certificate is auto-created. The SAP HANA log backup may fail until the first protection job is running.

  > **Note:** The certificate is not created during source registration.

- Cohesity supports native SAP HANA backup encryption. When encryption is enabled, backup data is encrypted and transferred to the Cohesity cluster for backups created

using the Cohesity SAP HANA Connector Agent.

- Supported encryption methods include:
- SAP HANA Backup Encryption
- System PKI SSFS (for internal communication encryption)
- Data and Log Volume Encryption

Supported secure stores:

- SSFS (Secure Store and File System)
- LSS (Local Secure Store)

Create the Secure User Store (hdbuserstore) Key

Create the SAP HANA hdbuserstore key to connect to SYSTEMDB. For detailed instructions, see Secure User Store (hdbuserstore). Use the following command to create the hdbuserstore key:

```
hdbuserstore SET <KEY> <host:port> <USERNAME> <PASSWORD>
```

> **Note:** Ensure that the user account of the secure user store has the following administrator privileges:
> - BACKUP ADMIN
> - DATABASE ADMIN
> - DATABASE START
> - CATALOG READ
> - MONITORING

Sample:

```
hdbuserstore SET SYS_KEY test-FI:30013 SYSTEM testpassword
```

After you have created the hdbuserstore key, verify if the key has the required SAP HANA connection information, such as hostname or login credentials. Verify the key using the following command:

```
tecadm@test-FI:/usr/sap/TEC/home> hdbsql -AU SYS_KEY
"select DATABASE_NAME,ACTIVE_STATUS from m_databases"
| DATABASE | ACT |
| -------- | --- |
| SYSTEMDB | YES |
| TEC | YES |
| SRC | YES |
| DST | YES |
4 rows selected (overall time 28.150 msec; server time 659 usec)
tecadm@test-FI:/usr/sap/TEC/home>
```

Download and Install the Cohesity Linux Agent

The Cohesity Linux Agent is available as script-based and RPM-based installers, based on your requirements download the installer and install it accordingly on the SAP HANA node (s). For detailed instructions, see Download and Install the Linux Agent.

> **Note:**
> - Use the `root` user account to install the Cohesity Linux Agent.
> - For multi-node SAP HANA deployments, install the Cohesity Linux Agent on the master node.

Download the SAP HANA Connector Installer

Cohesity provides a SAP HANA connector for SAP HANA to backup and restore SAP HANA databases. The SAP HANA connector for SAP HANA is available as a script or RPM-based installer on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective SAP HANA Connector installer for the Cohesity SAP HANA BACKINT plug-in, and the SAP HANA Connector installer to the SAP HANA host.

Install and Manage the SAP HANA Connector on IBM PowerPC Platform

This topic covers instructions on how to install the SAP HANA Connector on IBM PowerPC Platform.

Prerequisites

Before you register your SAP HANA deployment as a source with Cohesity and protect SAP HANA databases, ensure the following prerequisites:

- Identify the node(s) in your SAP HANA cluster with adequate resources to execute the SAP HANA connector.
- SAP HANA node(s) must be using JDK version 1.8 or higher.
- Ensure that the fully qualified domain name (FQDN) of the Cohesity cluster is resolvable from all the nodes in the SAP HANA cluster.
- Synchronize the SAP HANA Linux machine time with a Network Time Protocol (NTP) server. For detailed instructions, see How to Sync Linux Server Time with an NTP Server.
- Ensure that the NTP time is in sync between the client and the Cohesity cluster. Perform the following on the client:
  - Verify if NTP is enabled and synchronized in the client using the following command:

    ```
    timedatectl
    ```

- If the NTP time is not synchronized, use the following command to synchronize the NTP time

  ```
  timedatectl set-ntp true
  ```

- "bash" shell is required for SAP HANA connector installation, backup, and restore workflows.
- A database user with database admin user privileges.
- Download and install the Cohesity Linux Agent on the SAP HANA node.
- Download the SAP HANA connector Installer for the IBM PowerPC Platform.

Install the SAP HANA Connector

On the SAP HANA node where you have copied the SAP HANA connector, perform the following to install the SAP HANA connector.

> **Note:**
> For multi-node SAP HANA deployments, ensure that:
>
> - You install the SAP HANA connector on the machine where you have installed Cohesity Linux Agent.
> - The SAP HANA Connector is installed on the shared NFS path, to which all the SAP HANA nodes have access.

Install the SAP HANA Connector Using Script Based Installer

To install the SAP HANA connector using script-based installer:

1. Log on to the SAP HANA node's bash CLI using `<SID>adm` user.
2. Make the SAP HANA connector installer (`cohesity_secure_connector_service_<version>_sap_hana_java_installer`) executable using the following command:

   ```
   SAP_HANA_SERVER:/home/adminse # chmod +x cohesity_secure_connector_service_<version>_sap_hana_java_installer
   ```

3. Install the SAP HANA connector using the following command:

   ```
   ./cohesity_secure_connector_service_<version>_sap_hana_java_installer -- --install=<backint install directory> --sid=<SID> --systemdb-key=<systemdb_key> --workflow=<workflow>
   ```

> **Note:** You can view the details of SAP HANA Connector installer options using the `--help` option. For example: `./cohesity_secure_connector_service_sap_hana_installer --help`After the SAP HANA connector is installed, the SAP HANA connector agent is available at: `<backint_install_directory>/cohesity_backint_plugin/uda_scripts`

Where:

Options	Mandatory?	Description
`--install-dir`	No	The directory on the SAP HANA node on which you want to install the SAP HANA connector. **Note:** If `--install-dir` is not specified, the SAP HANA connector is installed in the `/usr/sap/${SAPSYSTEMNAME}/home` directory.
`--systemdb-key`	NO	The SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB. • Use the option `--systemdb-key-k <user_store_key>` only if you want to edit the global.ini file and create a symbolic link from `HDBBACKINT` to the SAP HANA connector. • If you plan to backup your SAP HANA databases using the SAP HANA connector, ensure that you use the option `--systemdb-key <user_store_key>`. • If you have not used the option `--systemdb-key <user_store_key>` while installing the SAP HANA connector, specify the option `--systemdb-key <user_store_key>` while upgrading the SAP HANA connector.
`--cluster-port`	No	The ports in the firewall to allow the Cohesity cluster to transmit and receive data. The supported values are 11111 or 11113. By default the port is set as `11117`.

Options	Mandatory?	Description
`--workflow`	No	The type of workflow you intend to install the SAP HANA connector. The supported values are `uda` or `ra`. The `uda` value is for Universal Data Adapter workflow and `ra` is for Remote Adapter workflow.
`--multi-tenancy`	No	The type of deployment. The supported values are `vlan` or `hyx`. Specify `vlan` for 'VLAN' and `hyx` for 'Single Homed NAT' deployments.
`--standby`	No	Specify if you are installing the SAP HANA connector on a standby node.
`--hana-version`	No	The SAP HANA version. Supported values are 1 or 2. By default, the value is set to 1.
`--debug`	No	Option to enable installer debug.
`--catalog-view`	No	The name of the catalog view where the backup metadata resides. **Note:** The catalog view name should not include the view path or cluster name or IP address.

After the SAP HANA connector is installed, the SAP HANA connector agent is available at: `<backint_install_directory>/cohesity_backint_plugin/uda_scripts`

For example:

```
/usr/sap/SID/home/cohesity/cohesity_backint_plugin/uda_scripts
```

4. Log on to the SAP HANA node as root user and create a soft link of `<backint_install_directory>/cohesity_backint_plugin/uda_scripts` folder to `/opt/cohesity/agent/uda_scripts` using the following command:

```
ln -s /usr/sap/SID/home/cohesity/cohesity_backint_plugin/uda_scripts
/opt/cohesity/agent/uda_scripts
```

> **Note:** Before creating the soft link, make sure to remove any existing soft link to the `<backint_install_directory>/cohesity_backint_plugin/uda_scripts` folder.

Install the SAP HANA Connector Using RPM Based Installer

Before you install the SAP HANA Connector Using an RPM-based installer, ensure that you configure the following environment variables on the SAP HANA machine:

Environment Variable	Description	Mandatory?	Command
SID	The <SID>adm user account that must own the backup files on the Cohesity view. If you have not created the user account, then create an Operating System (OS) user account in the following format: <SID>adm. Where <SID> is a unique three-character code called SAP System Identification (SID). It is available for every R/3 installation (SAP system) that consists of a database server and several application servers.	Yes	export SID=<VALUE>
SYSTEMDB_KEY	The SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB.	Yes	export SYSTEMDB_KEY=<HDBUSERSTORE key>
CLUSTER_PORT	The ports in the firewall to allow the Cohesity cluster to transmit and receive data. The supported values are 11111 or 11113. By default the port is set as 11117.	No	export CLUSTER_PORT=<Port_number>
WORKFLOW	The type of workflow you intend to install the SAP HANA connector. The supported values are uda or ra. The uda value is for Universal Data Adapter workflow and ra is for Remote Adapter workflow. Default is uda.	No	export WORKFLOW=ra
MULTI_TENANCY	The type of deployment. The supported values are vlan or hyx. Specify vlan for 'VLAN' and hyx for 'Single Homed NAT' deployments.	No	export MULTI_TENANCY=<VALUE>
STANDBY	Specify if you are installing the SAP HANA connector on a standby node. Supported values are 1 or 0.	No	export STANDBY=<VALUE>
HANA_VERSION	The SAP HANA version. Supported values are 1 or 2.	No	export HANA_VERSION=<VALUE>

Environment Variable	Description	Mandatory?	Command
DEBUG	Option to enable installer debug. Supported values are 1 or 0.	No	`export DEBUG =<VALUE>`
CATALOG_ VIEW	The name of the catalog view where the backup metadata resides. **Note:** The catalog view name should not include the view path or cluster name or IP address.	No	`export CATALOG_VIEW =<VALUE>`

To install the SAP HANA connector using rpm-based installer:

1. Log on to the SAP HANA node's bash CLI using <SID>adm user.
2. Make the SAP HANA connector installer (`cohesity-hana-java-backint-<version>.noarch.rpm`) executable using the following command:

```
> chmod +x cohesity-hana-java-backint-<version>.noarch.rpm
```

3. Install the SAP HANA connector using the following command:

```
> sudo -E rpm -ivh cohesity-hana-java-backint-<version>.noarch.rpm
```

Enable Displaying Log Backups in the Cohesity Cluster

SAP HANA automatically triggers log backups, you can view these log backups on the Cohesity cluster. For viewing the SAP HANA log backups on the Cohesity cluster, you need to authenticate to the Cohesity cluster using the following java binary available in the SAP HANA connector installation directory: `<backint_install_directory>/cohesity_backint_plugin/et_log_backup_tool.jar`

> **Note:** Ensure that you have JDK installed on the SAP HANA node.

After you have successfully installed the SAP HANA connector, run the following command:

```
java -jar <backint_install_directory>/cohesity_backint_plugin/et_log_backup_tool.jar --authenticate
```

Example:

```
java -jar /opt/cohesity/cohesity_backint_plugin/bin/et_log_backup_tool.jar --authenticate
```

Note:

- When prompted, authenticate using the **Cohesity cluster's local admin user credentials**. If the password has been changed, make sure to reauthenticate using the updated password.
- After you authenticate to the Cohesity cluster using the `et_log_backup_tool.jar` java binary, ensure that you set the `-et-log-backup` source registration parameter to `true` while registering the SAP HANA source. For more information, see Register and Manage the SAP HANA Source.
- After registering the source with Cohesity (with`-et-log-backup=true`), a full backup is mandatory. The initial full backup must be completed before any log backups appear on the Cohesity user interface.
- If you upgrade to version 7.4.1, you need to modify the existing registered source and set the `et-log-backup` source registration parameter to `true` in order to enable the display of log backups in the Cohesity cluster. Only the log backups triggered after enabling `et-log-backup` will be shown on the Cohesity cluster.

Uninstall the SAP HANA Connector

To uninstall the SAP HANA Connector using the root user account, use the following command:

Using the Script installer:

```
./cohesity_secure_connector_service_<version>_sap_hana_java_installer -- --
uninstall --install-dir=<backint_install_directory> --workflow=<workflow>
```

Using the RPM installer:

```
> sudo -E rpm -evh cohesity-hana-java-backint-<version>.noarch.rpm
```

Upgrade the SAP HANA Connector

To upgrade the SAP HANA connector, run the following command on all SAP HANA nodes.

Using the Script installer:

```
./cohesity_secure_connector_service_<version>_sap_hana_java_installer -- --
upgrade --install-dir=<backint_install_directory> --workflow=<workflow>
```

Using the RPM installer:

```
sudo -E rpm -Uvh cohesity-hana-java-backint-<version>.noarch.rpm
```

Register and Manage the SAP HANA Source

To protect your SAP HANA databases, you need to register your SAP HANA deployment as a source on Cohesity. After registering your SAP HANA deployment as a source, you can:

- Add more SAP HANA node(s) to run the SAP HANA database agent. Ensure that you have installed the Cohesity Linux Agent and SAP HANA database agent on the node (s).
- Refresh the source details.
- Unregister the SAP HANA source from Cohesity.

Register SAP HANA as a Source

Before you register SAP HANA as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP HANA Protection.

To register SAP HANA as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
   a. From the **Source Type** drop-down menu, select **SAP HANA**.
   b. From the **Host OS Type** drop-down menu, select **Linux**.
   c. In the **Hostname/IP Addresses** field, enter the hostname or IP address (IPv4 or IPv6) of the node(s) you have identified to run the SAP HANA database agent. For more information, see Plan and Prepare for SAP HANA Protection.

      > **Note:** Cohesity supports backup and restore of SAP HANA(x86) databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

   d. In the **Datasource Agent Installation Path** field, enter the directory path on the SAP HANA node, where you have installed the SAP HANA database agent. For more information, see Plan and Prepare for SAP HANA Protection.
   e. Ignore the In the **APP Authentication** section. It is not applicable to SAP HANA source registration.
   f. Enable **Custom Options**, perform the following:

i. Ignore the **Mounts View** option. It is not applicable for SAP HANA source registration.

ii. Optional. In the **Source Registration Options** field, enter the following source registration parameters based on your requirements:

The parameters are optional. If you do not explicitly specify the values then the values you had provided while installing the SAP HANA database agent are applied.

Parameters	Description
`--sourcename`	A unique name to identify the SAP HANA source. Using this parameter, you can override the unique name that you provided while installing the SAP HANA database agent. For more information, see Plan and Prepare for SAP HANA Protection. For example: `--sourcename=SAPSalesProd`
`--hana-ha-vips`	Optional. Provide the IP address configured for "**Virtual IP address resource**" on the SAP HANA HA cluster. For example: `--hana-ha-vips=10.11.23.112`
`--backint-protocol-version`	Optional. The BACKINT protocol version 1.5 is supported by default from SAP HANA 2.0 SPS 05 or later versions on the IBM PowerPC Platform. By default, the SAP HANA Connector agent supports BACKINT protocol version 1.0. Using this parameter, you can specify the SAP HANA Connector agent to use the BACKINT protocol version 1.5 based on your requirement. For example: `--backint-protocol-version=1.5`
`--et-log-backup`	Set the value to `true`, if you want to view the SAP HANA log backups on the Cohesity cluster. **Note:** Before registering the source, ensure that you authenticate to the Cohesity cluster using the **Cohesity cluster's local admin user credentials** and the `et_log_backup_tool.jar` java binary available in the SAP HANA connector installation directory. For more information, see Enable Log Backups Replication.

Parameters	Description
`--entity-hierarchy`	Optional. Set the value to `true` to enable auto-discovery and entity hierarchy that lets you browse or search and select SAP HANA objects for backup and restore. **Note:** If you are upgrading to version 7.2 or later, you need to update the existing SAP HANA source configuration to enable auto-discovery and entity hierarchy. Set the `--entity-hierarchy` source registration parameter to "`true`." After updating the source configuration (with `--entity-hierarchy=true`), a full backup is mandatory. Cohesity supports the following authentication methods: • **Cluster credentials authentication**: using the <SID>adm user and the et_log_backup_tool.jar java binary available in the SAP HANA connector installation directory. • **API Key authentication**: Uses a Cohesity API key and avoids dependency on cluster username and password rotation.
`--tcp-streams`	Optional. Specify the number of TCP connections you want to make with Cohesity Cluster. **Note:** This parameter is primarily useful when your source machine is behind a Hybrid Extender.

g. Click **Register**.

Update the SAP HANA Source Configuration

After registering your SAP HANA as a source on Cohesity, you can:

- Add more node(s) to run the SAP HANA database agent. Ensure that you have installed the Cohesity Linux Agent and SAP HANA database agent on the node(s).
- Modify the **Source Registration Options** based on your requirements.
- Update the **Datasource Agent Installation Path** only if you have uninstalled and reinstalled the SAP HANA database agent in a different directory on the SAP HANA node.

> **Note:** Ensure that the SAP HANA database agent is installed on all the node(s) in the SAP HANA cluster.

To update the SAP HANA source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP HANA source and select **Edit**.
3. In the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the SAP HANA Source Configuration

You can refresh the SAP HANA Source and re-verify the source with the available source registration parameters.

To refresh the SAP HANA Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP HANA Source and select **Refresh**.

Unregister the SAP HANA Source

If you plan to stop taking the backup of your SAP HANA databases, you can unregister the SAP HANA source from the Cohesity.

To unregister the SAP HANA source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP HANA source and select **Unregister**.

Backup SAP HANA Database

Cohesity provides a simple workflow to back up your SAP HANA databases. After registering the SAP HANA as a source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides a SAP HANA database agent for the backup and restore of enterprise SAP HANA deployments. The SAP HANA database agent leverages the backup capabilities provided by SAP HANA.

After registering the SAP HANA source on Cohesity, you can create a Protection Group and specify the objects such as databases and tables that you want to backup. In the Protection Group, you need to define the parameters that enable you to provide additional parameters required by the backup statements provided by SAP HANA.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your SAP HANA backup configuration:

Backup Parameters	Mandatory?	Description
`--user-store-key`	No	The hdbuserstore key to connect to the SAP HANA SYSTEMDB. For example: `--user-store-key=H00_KEY`
`--backup-delta`	No	The parameter is only applicable for Incremental or Differential backups. The supported values are INCREMENTAL or DIFFERENTIAL. If you do not explicitly specify a value, by default, incremental backup is considered. For example: `--backup-delta=DIFFERENTIAL`
`--use-secure-grpc=<true/false>`	No	Set the parameter to true, if you want to use secure gRPC communication during backup. For example: `--use-secure-grpc=true` > **Note:** It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
`--sap-certificate-path=<cert_file_path>`	No	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during backups. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`
`-r, --retention-days`	Yes	Specify the number of days to retain full or incremental backups. This argument helps manage storage by automatically removing backups older than the defined retention period. > **Note:** Ensure that the number of days specified in the argument matches the days defined in the backup policy.

Log Backups

SAP HANA automatically triggers log backups, you can view these log backups on the Cohesity cluster and perform the following:

- View the SAP HANA log backups on the Cohesity clusters.
- Replicate the SAP HANA log backups across the Cohesity clusters in your deployment.
- Archive the SAP HANA log backups.

To view SAP HANA log backups on the Cohesity cluster, perform the following:

1. After you have installed or upgraded the SAP HANA Connector, enable SAP HANA log backup replication. For more information, see Enable Log Backups Replication.
2. While registering or updating the SAP HANA source to the Cohesity cluster, set the `–et-log-backup source` registration parameter to `true`. For more information, see Register SAP HANA as a Source.
3. Create a new Protection Group and select a Protection Policy with log backups job in the Protection Group for SAP HANA backups. For more information, see Protect SAP HANA Deployments.

   > **Note:**
   > - To replicate jobs, select the replication option, while creating the Protection Group.
   > - In the source, after setting the source parameter `–et-log-backup` to `true`, the log backups will not be displayed on the Cohesity UI or be replicated for older Protection Groups. You must create a new Protection Group with the updated source.

4. Create a Full Backup for the first time.

Replication to the secondary Cohesity cluster will be enabled after the first successful full backup, and log backups will be displayed on the Cohesity UI.

> **Note:** Catalog backups are not displayed on the Cohesity UI.

Protect SAP HANA Deployments

SAP HANA deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to:

- Specify the objects such as databases and tables that you want to backup.
- Based on your requirement, define the parameters for the backup statements provided by SAP HANA. For more information, see Backup Parameters.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP HANA deployments:

1. Do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the SAP HANA source and then select **Protect**.
2. From the **Source** drop-down list, select the registered SAP HANA source or click **Register Source** to register a SAP HANA source. For more information, see Register SAP HANA as a Source.
3. In the **Objects** section, select the database or databases that you plan to backup. Then, click **Save Selection**.

   Important:
   - To select the database or databases from the entity hierarchy, make sure to enable entity hierarchy (`--entity-hierarchy`) during source registration. If it is not enabled, you must manually enter the object name for backup.
   - If you are upgrading to version 7.2 or later, you need to update the existing SAP HANA source configuration to enable auto-discovery and entity hierarchy. Set the `--entity-hierarchy` source registration parameter to "`true`." After updating the source configuration (with `--entity-hierarchy=true`), a full backup is necessary.
4. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.

> **Note:** SAP HANA automatically schedules log backups, and you can view these log backups on the Cohesity cluster. Ensure that you add a log backup job in the Protection Policy. The Protection Policy with log backup job retains the SAP HANA log backups and is not used for scheduling SAP HANA log backups.

6. In the **Settings** section, do the following:
    1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
    2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

       Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
    3. In the **Custom Options** section, perform any one of the following:
        - Do not enter any backup parameters in the **Full Backup** and **Incremental Backup** fields if you plan to perform full backup or incremental backup respectively.
        - Enter the Backup Parameters in the **Full Backup** and **Incremental Backup** fields based on your requirements.

          For example:

Use Cases	Sample Parameters
Create a full backup of the SAP HANA cluster by overriding the user store key. **Note:** In the case of SAP HANA cluster backup, it takes the user store key as `H00_KEY` which is overridden with the backup parameter.	In the **Full Backup** field, specify the following parameter: `--user-store-key=H00_KEY`

Use Cases	Sample Parameters
Create an incremental backup of the SAP HANA cluster by overriding the user store key. **Note:** In the case of SAP HANA cluster backup, it takes the user store key as `H00_KEY` which is overridden with the backup parameter.	In the **Incremental Backup** field, specify the following parameter: `--user-store-key=H00_KEY`
Create a differential backup of the SAP HANA cluster. **Note:** In the case of SAP HANA cluster backup, it takes the backup delta as DIFFERENTIAL which is overridden with the backup parameter.	In the **Incremental Backup** field, specify the following parameter: `--backup-delta=DIFFERENTIAL`

- In the **Mounts** field, enter the number of VIPs you want to use while writing data to the Cohesity storage.

  Note: Cohesity recommends that the number of mounts is equal to the number of nodes or Hybrid Extenders.

7. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click Edit on the right. For more information, see Additional Settings.
8. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	The maximum number of restore streams that is created to exchange data with the cluster. By default, maximum scaling is set to empty. > **Note:** The same value is set to the `parallel_data_backup_backint_channels` parameter in the **global.ini** file for the protected database. Also, the same value is set for the `max_restore_streams` parameter in the SAP HANA parameter file of the database selected for protection.
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the SAP HANA Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP HANA.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.

5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP HANA.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity Platform provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.

- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP HANA.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Protection Group Only
   - Delete Protection Group and Snapshots

Recover SAP HANA Database

You can restore the SAP HANA cluster or object(s) such as databases in SAP HANA. You can restore to the same SAP HANA cluster or a different SAP HANA cluster.

The SAP HANA database agent leverages the restore statements provided by SAP HANA. You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

In the Recovery configuration, you need to specify the object you plan to restore and define the script parameters required by the restore statements provided by SAP HANA.

Recovery Parameters

In the Recovery configuration, you can define the following script parameters for your SAP HANA recovery:

> **Note:** The recovery parameters are optional parameters, you can use these parameters based on your requirements.

Recovery Parameters	Mandatory ?	Description

`--source-sid`	Mandatory for alternate restore use cases.	Provide the SAP HANA source system SID for the SAP HANA host used for the alternate restore use case.
		For example:
		`--source-sid=<sid>`
		`--source-sid=H00`
		Note:
		• The `--source-sid` parameter is applicable only for alternate restore use cases.
		• Ignore the `--source-sid` parameter, if you have enabled entity hierarchy (`--entity-hierarchy`) during source registration.

`--user-store-key`	Optional	The hdbuserstore key to connect to the SAP HANA SYSTEMDB. For example: `--user-store-key=H00_KEY`
`--start-database`	Optional	Set the value to `true` if you want to activate the database after the recovery is completed. Supported values are `true` or `false`. By default the value is set to `true`.
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during restore. For example: `--use-secure-grpc=true`
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during restore. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`

Prerequisites

Ensure that:

- The Protection Group associated with the object that you plan to restore is paused or in not running state.
- The target database exists, or you manually create the target database before you start the restore process.

- The SAP HANA version used for recovery must always be the same or newer than the version used for backup of data and logs.
- If any of the selected snapshots were backed up with SAP HANA backup encryption enabled, ensure that the encryption keys are recovered to the SAP HANA source system before starting the recovery process.
- All the required SAP HANA services are up and running. You can use the following sample command to verify the required services:

```
sap-hana-2:cohadm> sapcontrol -nr ${TINSTANCE} -function GetProcessList
03.09.2024 14:58:23
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2024 09 02 16:49:20, 22:09:03,
14346
hdbcompileserver, HDB Compileserver, GREEN, Running, 2024 09 02
16:49:26, 22:08:57, 14488
hdbindexserver, HDB Indexserver-COH, GREEN, Running, 2024 09 03
13:26:33, 1:31:50, 24735
hdbnameserver, HDB Nameserver, GREEN, Running, 2024 09 02 16:49:21,
22:09:02, 14374
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2024 09 02 16:49:26,
22:08:57, 14491
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2024 09 02
16:50:31, 22:07:52, 15062
hdbxsengine, HDB XSEngine-COH, GREEN, Running, 2024 09 03 13:26:53,
1:31:30, 25060
hdbindexserver, HDB Indexserver-TESTDB, GREEN, Running, 2024 09 03
13:30:39, 1:27:44, 31263
hdbxsengine, HDB XSEngine-TESTDB, GREEN, Running, 2024 09 03 13:31:03,
1:27:20, 31567
```

To stop and start the required services, use the following command:

```
sap-hana-2:cohadm> HDB stop
sap-hana-2:cohadm> HDB start
```

Recover SAP HANA Databases

To restore SAP HANA:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.

4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, select the database that you had backed up using the selected Protection Group.

   > **Note:** Using a recovery job, you can recover only a single database.

   By default, the latest snapshot is selected for recovery. You can also select a point-in-time snapshot or specific data backup that you want to recover. Click the pencil icon and perform the following:

   - For point-in-time snapshots, in the **Edit recovery point for <Protection_Group_Name>** dialog, click the **Timeline** tab, choose the **Date** and select the **point-in-time** from the **timeline** slider and click **Select Recovery Point**.
   - For specific data backup snapshots, in the **Edit recovery point for <Protection_Group_Name>** dialog, click the **list** tab and select the specific **Snapshot Time** that you want and click **Select Recovery Point**.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same SAP HANA database.
New Object	To restore to a different database on the same SAP HANA cluster or a different SAP HANA cluster. Select the **Target** SAP HANA cluster to which the data needs to be restored. Also, ignore **Overwrite existing object with the same name** toggle option. It is not applicable for SAP HANA.

8. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Enter the number of VIPs you want to use while reading data from the Cohesity storage.

Options	Description
Concurrency	The maximum number of restore streams that is created to exchange data with the cluster. By default, maximum scaling is set to empty. **Note:** The same value is set to the `parallel_data_backup_backint_channels` parameter in the **global.ini** file for the protected database. Also, the same value is set for the `max_restore_streams` parameter in the SAP HANA parameter file of the database selected for protection.
Custom Options	Specify the Recovery Parameters based on your requirements. For examples, see Sample Recovery Use Cases.
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

9. Click **Start Recovery**.

Sample Recovery Use Cases

Use Case	Sample Recovery Parameters
Overwrite Point In Time Restore (PITR) to the same SAP HANA database.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, do not specify any parameters in the **Custom Options** field.
Overwrite Point In Time Restore (PITR) the same SAP HANA database overriding the user store key.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, specify `--user-store-key=H00_KEY` in the **Custom Options** field.

Use Case	Sample Recovery Parameters
Restore specific data backup from the snapshot to the same SAP HANA database.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, specify `--user-store-key=H00_KEY --is-specific-data-backup-restore=true` in the **Custom Options** field.
Point In Time Restore (PITR) to a different database on the same SAP HANA system.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the same source name in the **Target** drop-down. 2. Under the **Recovery Options** section, specify `--user-store-key=H00_KEY --source-sid=H00` in the **Custom Options** field and specify the name for the new database in the **Rename** field.
Point In Time Restore (PITR) to a different database on a different SAP HANA system	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** drop-down. 2. Under the **Recovery Options** section, specify `--user-store-key=H00_KEY --source-sid=H00` in the **Custom Options** field and specify the name for the new database in the **Rename** field.

Use Case	Sample Recovery Parameters

Restore specific data backup from the snapshot to a new database on the same SAP HANA system.	Perform the following:
	1. In the **Recover As** section, select **New Object** and choose the same source name in the **Target** drop-down.
	2. Under the **Recovery Options** section, specify the following in the **Custom Options** field and specify the name for the new database in the **Rename** field:
	`--user-store-key=H00_KEY`
	`--source-sid=H00`
	`--is-specific-data-backup-restore=true`

Cancel a Recovery Run

When you cancel a recovery run, the recovery process is terminated. If the recovery run is canceled during an ongoing restore operation, the SAP HANA database may enter an offline or inconsistent state. To bring the SAP HANA database back online, you must initiate a new recovery run.

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click the action icon (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
   - **Continue Recovery**. if you want to continue to recovery run
   - **Cancel Recovery**. if you want to cancel the recovery run.

SAP HANA Certificate Authentication

The SAP HANA Certification provides the necessary authentication and also verifies whether the SAP HANA database, which performs backups to the Cohesity cluster, is a trusted entity. This certification provides the capability to set verification which allows only trusted sources to write to the Cohesity cluster through the SpanFS.

Create an SAP HANA Authentication Certificate

This section lists the steps for creating an SAP HANA authentication certificate.

1. Start the Cohesity CLI locally and log in to the Cohesity cluster using the ssh command with the `support` user account and specify the IP address. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

```
ssh support@198.51.100.12
```

2. If prompted, enter the password.

```
[support@198.51.100.12's password:
```

The cohesity_shell prompt is displayed.

3. Log in to the CLI:

```
cohesity shell# iris_cli
```

The admin prompt is displayed.

4. Optionally, run the help command for a list of CLI commands.

```
admin@198.51.100.12> help
```

5. Run the cluster gen-src-cert command to generate a new SAP HANA certificate.

```
admin@198.51.100.12 cluster gen-src-cert cert-name="newdemocert" cert-
validity=100 host-name="10.2.146.175" host-target="/tmp"
Success: Certificate "newdemocert" generation and deployment
successful.
```

Parameter Name	Description
cert-name	Name of the certificate file. Its naming format is ${cert-name}_<random_number>.cfg.
cert-validity	The validity of the source certificate generated. The time is displayed in days.
host-name	The IP address/host name of the target system where you want to move the certificate. This is a mandatory parameter because the certificate generation process needs details of the host machine to generate public and private keys.
host-target	The file directory path on the Cohesity cluster where the certificate is generated.

> **Note:** After the SAP HANA certificate is generated on the Cohesity cluster, manually copy the certificate to your SAP HANA machine.

6. Update the following parameters in the parameter file.

Parameter Name	Description
cluster_port	Port number on Cohesity cluster nodes to which the BACKINT executable connects. Use the port 11117 for secure connection (certificate config path should be a valid path). Use the port 11111 for an insecure connection (certificate config path should be an empty string). Contact Cohesity Support to enable these ports.
certificate_ config_path	The path to which the certificate is generated and deployed from the Cohesity Dashboard to the SAP HANA Node. You can use this path as the default config path: `/usr/sap/ORC/home/install/newdemocert.cfg`

7. Run the `list-source-certificates` command to check all the global certificates deployed from the cluster.

Deploy an SAP HANA Certificate in a Multi-node Setup

This section lists the steps to deploy an SAP HANA authentication certificate in a multi-node SAP HANA setup.

1. Generate and deploy the SAP HANA authentication certificate to any one of the nodes in a multi-node SAP HANA system using the instructions provided in the previous section.
2. Use SCP to securely transfer the certificate to a target location on all the remaining (N-1) nodes.

   ```
   scp <certificate_config_path> ip_address:<certificate_config_path>
   ```

   Example

   ```
   scp /usr/sap/ORC/home/install/cert.cfg 10.2.146.72:
   /usr/sap/ORC/home/install/cert.cfg
   ```

3. Change the `cluster_port parameter` value to 11117 and update the `certificate_ config_path` parameter value to the correct path of the certificate in the respective paramfile of each of the N nodes.

SAP Sybase IQ

Cohesity Platform natively integrates with SAP Sybase IQ to provide backup and recovery for clustered SAP Sybase IQ deployments. Using Cohesity Connector Agent, you can backup and restore the SAP Sybase IQ database. The Cohesity Connector Agent provides a simple workflow to back up and restore your SAP Sybase IQ deployment.

Cohesity Platform provides an SAP Sybase IQ Plug-in for backup and restore of the SAP Sybase IQ database. The SAP Sybase IQ plug-in leverages the third-party dynamic link library (DLL) based capabilities provided by SAP Sybase IQ.

Core Components

The following are the core components for protecting SAP Sybase IQ databases:

- SAP Sybase IQ cluster
- Cohesity SAP Sybase IQ plug-in
- Cohesity Connector Agent

As illustrated, you can integrate the Cohesity Platform with your SAP Sybase IQ deployments using the SAP Sybase IQ Plug-in and Cohesity Connector Agent.

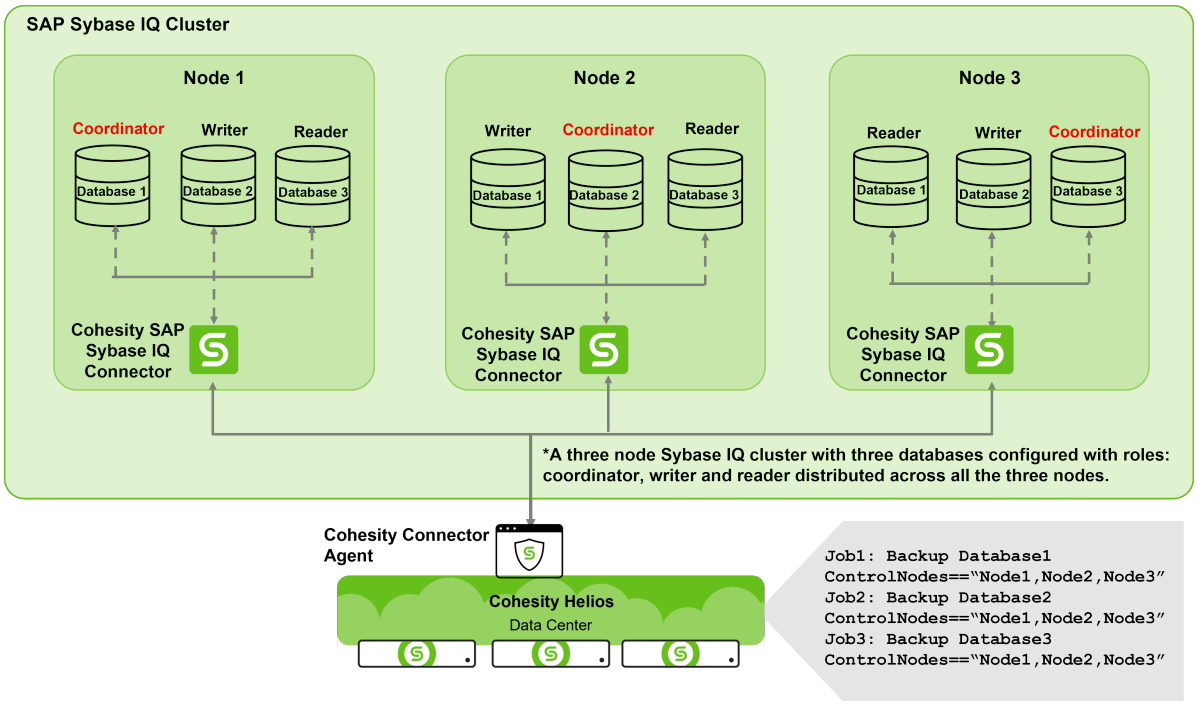

Plan and Prepare for SAP Sybase IQ Protection

Before you register your SAP Sybase IQ deployment as a source with Cohesity Platform, ensure the following prerequisites:

- Identify the Coordinator Node of the database that you plan to protect. The backup or restore scripts run on the coordinator node. If you have a multiplex cluster with failover capability for the coordinator node, choose all those nodes where the coordinator node can failover. To identify the failover node for the coordinator, use

the following DB command:

```
select connection_info ,coordinator_failover from sp_iqmpxinfo() where
coordinator_failover=server_name;
```

- A database user with database admin user privileges.
- Download and install the Cohesity Linux Agent on the SAP Sybase IQ nodes.
- Download and install the SAP Sybase IQ connector RPM-based installer.
- Review and understand the prerequisites.

Supported Versions

For more information on supported SAP Sybase IQ versions, see Supported Software versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	SAP Sybase IQ Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP Sybase IQ Local Host	SAP Sybase IQ Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP Sybase IQ Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
SAP Sybase IQ Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Supported Backup and Restore Matrix

The following table provides the support matrix for backup and restore operations:

		Supported SAP Sybase IQ Database Restore Versions			
		16.0 SP11 PL28	**16.1 SP03 PL9**	**16.1 SP04 PL01**	**16.1 SP04 PL09**
SAP Sybase IQ Database Backup Version	**16.0 SP11 PL28**	Supported	Supported	Supported	Supported
	16.1 SP03 PL9	Not Supported	Supported	Supported	Supported
	16.1 SP04 PL01	Not Supported	Not Supported	Supported	Supported
	16.1 SP04 PL09	Not Supported	Not Supported	Supported	Supported

Considerations

Review and understand the following before you protect your SAP Sybase IQ database:

- You can back up only one database per Protection Group.
- Currently only online backup of databases is supported.
- The same database should not be part of multiple Protection Groups.
- Log backups are not supported.
- SAP Sybase IQ does not support Point in time recovery (PITR) from backups taken using a third-party dynamic link library (DLL).
- In the case of multiplex database restore, data is restored to only one node (coordinator). You need to set the multiplex roles and start the multiplex database manually.
- Cohesity supports backup and restore of SAP Sybase IQ databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

Prerequisites

- For backups, the database server must be up and running.
- For Kerberos authentication, the keytab file must be available on all nodes at the given location and accessible by the user used to run the Cohesity Linux Agent.
- SSL certificates must be used for connecting to the database servers. The SSL certificates must be available in the same path on all servers and accessible by the user used to run the Cohesity Linux Agent.

- To securely connect to the Cohesity cluster while reading or writing data, you can use secure gRPC.

  > **Note:**
  > - From 7.2.2 version onwards, it is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
  > - Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
  > - Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

  To use secure gRPC, perform the following:

  1. Generate the certificate config file on the Cohesity node using the following command:

     ```
     ~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name <clustername> -output_file /home/cohesity/<clustername>.cfg
     ```

  2. Copy `/home/cohesity/<clustername>.cfg` from the Cohesity node to all your SAP Sybase IQ nodes at the same location. For example, `/opt/certs/<clustername>.cfg`.

     > **Note:** Ensure that the `root` or `non-root` user used to install the Cohesity Linux agent and your SAP Sybase IQ Connector has read access to the generated certificate config file.

  3. While registering the your SAP Sybase IQ source, provide the following as a part of **Source Registration Options**:

     ```
     --certificate-config-path=opt/certs/<clustername>.cfg
     ```

Download and install the Cohesity Linux Agent

The Cohesity Linux Agent is available as script-based and RPM-based installers. Based on your requirements, download the installer and install it accordingly on the node(s) of the SAP Sybase IQ cluster. For detailed instructions, see Download and Install the Linux Agent.

> **Note:** Ensure that you install the Cohesity Linux Agent using the RPM-based installer.

Download and Install the SAP Sybase IQ Connector Installer

Cohesity Platform provides an SAP Sybase IQ connector for the backup and restore of SAP Sybase IQ deployments. Sybase IQ supports backup and restore operations with third-party products using a dynamic link library loaded at run time(.dll or .so).

Download the SAP Sybase IQ Connector

Cohesity Platform provides a SAP Sybase IQ Connector to backup and restore SAP Sybase IQ databases. The SAP Sybase IQ Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for SAP Sybase IQ Connector, and copy the datasource agent to the SAP Sybase IQ host.

Install the SAP Sybase IQ Connector

Before you install the SAP Sybase IQ Connector, ensure that you have `sudo` privileges.

To install the SAP Sybase IQ connector:

1. Log on to the SAP Sybase IQ node where you have copied the SAP Sybase IQ connector RPM-based installer, `cohesity-sybaseiq-connector-<version>.x86_64.rpm`.
2. Run the following command to install the SAP Sybase IQ connector:

   ```
   sudo sh -c 'export SYBASEIQ_ENV_PATH=/path/to/SAP Sybase IQ/IQ-16_0/IQ-16_0.sh ; yum localinstall /home/cohesity/cohesity-sybaseiq-connector-<version>.x86_64.rpm'
   ```

   Where `/path/to/SAP Sybase IQ/` is the directory where the SAP Sybase IQ connector is installed.

   > **Note:** If you do not explicitly specify the details, the hostname is applied as the value for control nodes and source name. Also, a unique hash value is applied for the source UUID.

Upgrade the SAP Sybase IQ Connector

To upgrade the SAP Sybase IQ Connector, run the following command with appropriate parameters on SAP Sybase IQ node's CLI:

```
$ sudo sh -c 'export SYBASEIQ_ENV_PATH=/opt/sybase/iq16/IQ-16_1/IQ-16_1.sh ; yum update cohesity-sybaseiq-connector-<version>-1.x86_64.rpm'
```

Uninstall the SAP Sybase IQ Connector

To uninstall the SAP Sybase IQ Connector, run the following command with appropriate parameters on SAP Sybase IQ node's CLI:

Using `yum`:

```
sudo yum remove cohesity-sybaseiq-connector-<version>.x86_64
```

Using `rpm`:

```
sudo rpm -e cohesity-sybaseiq-connector-<version>.x86_64.rpm
```

Register and Manage the SAP Sybase IQ Source

To protect your SAP Sybase IQ databases, you need to register your SAP Sybase IQ deployment as a source on the Cohesity Platform. After registering your SAP Sybase IQ deployment as a source, you can:

- Update the SAP Sybase IQ source configuration, such as:
  - Add more coordinator node(s) to run the SAP Sybase IQ Connector. Ensure that you have installed the Cohesity Linux Agent and SAP Sybase IQ Connector on the node(s).
  - If you connect to the database using password-based authentication then you can modify the username and password of the database user.
- Unregister the SAP Sybase IQ source from the Cohesity Platform.

Register SAP Sybase IQ as a Source

Before you register SAP Sybase IQ as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP Sybase IQ Protection.

To register SAP Sybase IQ as a source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
   a. From the **Source Type** drop-down menu, select **SAP Sybase IQ**.
   b. From the **Host OS Type** drop-down menu, select **Linux**.
   c. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the coordinator node for a given database that you have identified to run the SAP Sybase IQ connector. For more information, see Plan and Prepare for SAP Sybase IQ Protection.

> **Note:** Cohesity supports backup and restore of SAP Sybase IQ databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

d. In the **Datasource Agent Installation Path** field, enter the directory path on the SAP Sybase IQ Linux machine, where you have installed the SAP Sybase IQ Connector. For more information, see Install the SAP Sybase IQ Connector.

e. In the **APP Authentication** section, enter the **Username** and **Password** of the SAP Sybase IQ database admin user to connect to the database using password-based authentication.

> **Note:** If you do not want to use password-based authentication, ignore the **OS APP Authentication** section.

f. Enable **Custom Options**, and perform the following:

   a. Ignore the **Mounts View** option. It is not applicable for SAP Sybase IQ source registration.

   > **Note:** The **Mount View** option is not applicable, and ensure you do not enable the option.

   b. In the **Source Registration Options** field, enter the following source registration parameters based on your requirements:

   For example:

```
--source-name=sybaseDBCluster1
```

Parameters	Mandatory?	Default	Description
`--source-name`	Yes	N/A	A unique name to identify the SAP Sybase IQ source. For example: `--sourcename=SalesProd` **Note:** Once you configure the source name, you cannot change it later.
`--use-secure-grpc`	No	False	Set the parameter to `true`, to enable secure connection to the Cohesity cluster while reading or writing data.

Parameters	Mandatory?	Default	Description
`--certificate-config-path`	No	N/A	The path to the certificate used to securely connect to the Cohesity cluster.

Parameters	Mandatory?	Default	Description
			Note: • It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data. • Ensure that you generate and copy the certificate config file from the Cohesity node to all the SAP Sybase IQ nodes. For more information, see Prerequisites. • Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communicatio n. • Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure

g. Click **Register**.

Update the SAP Sybase IQ Source Configuration

After registering your SAP Sybase IQ as a source on the Cohesity Platform, you can:

- Add more coordinator node(s) to run the SAP Sybase IQ Connector. Ensure that you have installed the Cohesity Linux Agent and SAP Sybase IQ Connector on the node(s).
- If you connect to the database using password-based authentication then you can modify the username and password of the database user.
- Modify the **Source Registration Options** based on your requirements.
- Update the **Datasource Agent Installation Path** only if you have uninstalled and reinstalled the SAP Sybase IQ Connector in a different directory on the SAP Sybase IQ node.

> **Note:** Ensure that the SAP Sybase IQ Connector is installed on all the node(s) in the SAP Sybase IQ cluster.

To update the SAP Sybase IQ source configuration:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase IQ source and select **Edit**.
3. In the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Unregister the SAP Sybase IQ Source

If you plan to stop taking the backup of your SAP Sybase IQ databases, you can unregister the SAP Sybase IQ source from the Cohesity Platform.

To unregister the SAP Sybase IQ source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase IQ source and select **Unregister**.

Backup SAP Sybase IQ Database

Cohesity Platform provides a simple workflow to back up your SAP Sybase IQ databases. After registering the SAP Sybase IQ as a source on the Cohesity Platform, you can start configuring your backups.

Backup Process

Cohesity Platform provides an SAP Sybase IQ connector for the backup and restore of SAP Sybase IQ deployments. The SAP Sybase IQ connector leverages the third-party dynamic link library (DLL) based capabilities provided by SAP Sybase IQ.

After registering the SAP Sybase IQ source on the Cohesity Platform, you can create a Protection Group and specify the objects such as databases you want to backup. In the Protection Group, you need to define the parameters that enable you to provide additional parameters required by the backup statements provided by SAP Sybase IQ.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your SAP Sybase IQ backup configuration:

Backup Parameters	Default Value	Mandatory?	Description
`--database-server`	N/A	Yes	The database server name that is used to connect to the database. You can retrieve the database server name using the following command: `dblocate -dn <dbname>`
`--database-port`	N/A	No	If database broadcast is disabled, specify the database port.
`--database-host`	N/A	No	If database broadcast is disabled, specify the database host.
`--database-config-path`	N/A	No	The path of the database configuration file that is used to start the database. Cohesity SAP SybaseIQ Connector checks database configuration file at location `<.db path>/params.cfg` or `<.db path>/dbname.cfg`. If the file is not present, you have to explicitly provide the absolute path of the database configuration file or the parameters file.
`--sybase-block-factor`	25	No	Specifies the number of blocks to write to the archive device at a time.

Backup Parameters	Default Value	Mandatory?	Description
`--check-crc`	true	No	Allows you to verify the checksums computed on the backup during any subsequent restore operation.
`--user-principal`		Yes (If you are using Kerberos authentication)	The Kerberos principal required to connect to the database.
`--keytab-path`			The absolute path of the keytab file for a given Kerberos principal. **Note:** The path must be accessible by the Cohesity Agent user.

Protect SAP Sybase IQ Deployments

In Cohesity, SAP Sybase IQ deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to:

- Specify the objects such as databases and tables that you want to backup.
- Based on your requirement, define the parameters for the backup statements provided by SAP Sybase IQ. For more information, see Backup Parameters.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP Sybase IQ deployments:

1. In the Cohesity Dashboard, do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Sybase IQ and then select **Protect**.
2. From the **Source** drop-down list, select the registered SAP Sybase IQ source or click **Register Source** to register a SAP Sybase IQ source. For more information, see Register SAP Sybase IQ as a Source.

3. In the **Objects** section, add the object that you plan to protect. The object should be a database name. To back up a database, enter the name of the database. For example, `DEMODB`. Click the plus icon (+) to enter multiple objects that you plan to back up. Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the Settings section, do the following:
   1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
   2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

      Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
   3. In the **Custom Options** section, enter the backup parameters in the **Full Backup**, **Incremental Backup** and **Log Backup** fields based on your requirements. For more information, see Backup Parameters.

      > **Note:** Ensure that you enter the same backup parameters in the **Full Backup** and **Incremental Backup** fields.

      For example: To perform backup of an SSL-enabled database with Kerberos authentication. In the Full Backup and Incremental Backup fields, enter the following parameters:

      ```
      --database-server=bankdb_server --user-principal=username@REALM_
      NAME --keytab-path=/data/bankdb/username.keytab --cert-
      path=/data/bankdb/certs/CA.crt
      ```
   4. In the **Mounts** field, enter the number of VIPs you want to use while writing data to the Cohesity storage.

> **Note:** Cohesity recommends that the number of mounts is equal to the number of nodes or Hybrid Extenders.

7. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
8. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the cluster. By default, maximum scaling is capped at 16.

Additional Settings	Description
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the SAP Sybase IQ Backup Configuration

In the Cohesity Platform, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase IQ.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ on the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase IQ.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase IQ.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Protection Group Only
   - Delete Protection Group and Snapshots

Recover SAP Sybase IQ Database

You can restore the SAP Sybase IQ object(s) such as databases in SAP Sybase IQ. You can restore to the same SAP Sybase IQ cluster or a different SAP Sybase IQ cluster.

The SAP Sybase IQ connector leverages the restore statements provided by SAP Sybase IQ. You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

In the Recovery configuration, you need to specify the object you plan to restore and define the parameters required by the restore statements provided by SAP Sybase IQ.

> **Note:** In the case of multiplex database restore, data is restored to only one node (coordinator). Therefore, you must set the multiplex roles and start the multiplex database manually. For details on how to setup multiplex roles, see SAP IQ Administration Guide,

Recovery Parameters

In the Recovery configuration, you can define the following parameters for your SAP Sybase IQ recovery:

> **Note:** The recovery parameters are optional parameters, you can use these parameters based on your requirements.

Recovery Parameters	Default Value	Mandatory?	Description
`--catalog-store-path`	N/A	Yes	The absolute path of the directory to which the catalog store must be restored. **Note:** The directory must already exist.
`--main-store-path`	The path provided in `--catalog-store-path`	No	The absolute path of the directory to which the IQ main store must be restored. If you do not define the directory path, the IQ main store is restored to the directory provided in the "`--catalog-store-path`" parameter. **Note:** The directory must already exist.
`--temp-store-path`	N/A	No	The absolute path of the directory to which the temp store must be restored. **Note:** The directory must already exist.

Recovery Parameters	Default Value	Mandatory?	Description
`--shared-temp-store-path`	N/A	No	The absolute path of the directory to which the IQ shared temp store must be restored. If you do not define the directory path, the IQ shared temp store is restored to the directory provided in the "`--catalog-store-path`" parameter. **Note:** The directory must already exist.
`--coordinator`	False	No	The hostname or IP address of the node where you want to restore data in case of multi-node cluster. **Note:** Applicable only for multi-node cluster.
`--start-database`	False	No	Specify if you want to start the database at the end of the restore operation.
`--stop-database`	True	No	Suppose the database with the same name is already running on the cluster where you are restoring the database. Specify if you want to stop the database.
`--service-user`	Backup instance user	No	After the data is restored, you can start the database with the specified user. If the user is not specified, the same user with which the database was backed up is used to start the database.

Recover SAP Sybase IQ Databases

To recover SAP Sybase IQ databases:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.

5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same SAP Sybase IQ database.
New Object	To restore to a different SAP Sybase IQ cluster. Select the **Target** SAP Sybase IQ cluster to which the data needs to be restored. Also, enable **Overwrite existing object with the same name** and type **Yes**, if you want to overwrite the existing object with the same name.

8. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	Not applicable for SAP Sybase IQ database recovery.
Custom Options	Specify the Recovery Parameters based on your requirements. For examples, see Sample Recovery Use Cases.
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

9. Click **Start Recovery**.

Sample Recovery Use Cases

Use Case	Sample Recovery Parameters
Restore a database to the location: `/mnt/database1`	`--catalog-store-path=/mnt/database1`

Use Case	Sample Recovery Parameters
Restore a database with: • Catalog store at `/data/restore/catalog` • IQ_SYSTEM_MAIN at `/data/restore/main` • IQ_SYSTEM_TEMP at `/data/restore/temp`	`--catalog-store-path=/data/restore/catalog --main-store-path=/data/restore/main --temp-store-path=/data/restore/temp`

Cancel a Recovery Run

When you cancel a recovery run, the SAP Sybase IQ connector cancels the job by terminating the recovery processes.

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click the action icon (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
    - **Continue Recovery** if you want to continue to recovery run
    - **Cancel Recovery** if you want to cancel the recovery run.

SAP MaxDB on Linux

Cohesity natively integrates with SAP MaxDB to provide backup and recovery for clustered SAP MaxDB deployments. Using Cohesity Connector Agent, you can backup and restore the SAP MaxDB database.

Cohesity provides a Cohesity BACKINT plug-in for the backup and restore of enterprise SAP MaxDB deployments. The Cohesity BACKINT plug-in leverages the backup and restore capabilities provided by SAP MaxDB.

The Cohesity BACKINT plug-in enables you to back up and restore both TenantDB and SystemDB. It provides:

- Backup:
    - Full backup
    - Incremental backup
    - Scheduled Log backup
- Restore:
    - Restore the database to the point in time.
    - Restore the database to a specific data backup.

- Alternate restore of SAP MaxDB database:
  - Alternate restore to a new database or active database on the same SAP MaxDB system.
  - Alternate restore to a new or active database on a different SAP MaxDB system.

Core Components

The following are the core components for protecting SAP MaxDB databases:

- SAP MaxDB Linux machine
- Cohesity BACKINT plug-in
- Cohesity Connector Agent

As illustrated, you can integrate the Cohesity Cluster with your SAP MaxDB deployments using the Cohesity BACKINT plug-in and Cohesity Connector Agent.

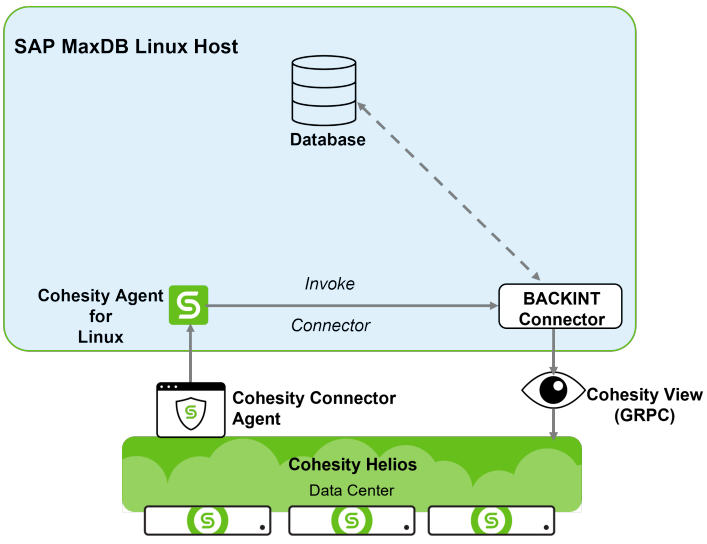

Plan and Prepare for SAP MaxDB on Linux Protection

Before you register your SAP MaxDB deployment as a source with Cohesity Platform and protect SAP MaxDB databases, ensure the following prerequisites:

- Ensure that the fully qualified domain name (FQDN) of the Cohesity cluster is resolvable from the SAP MaxDB Linux machine.
- Synchronize the SAP MaxDB Linux machine time with a Network Time Protocol (NTP) server. For detailed instructions, see How to Sync Linux Server Time with an NTP Server

- Ensure that the NTP time is in sync between the client and the Cohesity cluster. Perform the following on the client:
  - Verify if NTP is enabled and synchronized in the client using the following command:

    ```
    timedatectl
    ```

  - If the NTP time is not synchronized, use the following command to synchronize the NTP time

    ```
    timedatectl set-ntp true
    ```

- A database user with database admin user privileges.
- Download and install the Cohesity Linux Agent on the SAP MaxDB Linux machine.
- Download and install the SAP MaxDB Connector agent.
- Create a XUSER key.
- Create a backup template.

Supported Versions

For more information on supported SAP MaxDB versions, see Supported Software versions.

SAP MaxDB Connector Requirements

To install and run the Cohesity SAP MaxDB Connector agent, ensure that the SAP MaxDB host meets the following specifications:

Components	Specification
SAP MaxDB Version	For details, see Supported Software.
Operating System	
Platform	x86_64
Disk Space	• Minimum 256 MB. Required for installing the SAP MaxDB Connector agent. • 1 GB. Required for the log directory.
Memory	8GB of free memory to run the SAP MaxDB Connector agent.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	SAP MaxDB Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP MaxDB Local Host	SAP MaxDB Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP MaxDB Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
SAP MaxDB Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following before you protect your SAP MaxDB database:

- SAP MaxDB with liveCache technology is not supported.
- NFS datapath is not supported.
- Only PIPE backups are supported. File backups are not supported, and hence only scheduled log backups are supported, and automatic log backups are not supported.
- In a protection run, you cannot back up multiple objects.
- Only the single-node MAXDB database system is supported.
- Canceling restore operations is not supported.
- For alternate restore, the target SAP MaxDB database must be created before starting the restore operation.
- For alternate restore to a different SAP MaxDB host, the path specified in IndepDataPath must be the same.
- Cohesity replication is supported.

- After a successful database restore, Cohesity recommends to trigger a full backup to ensure that incremental and log backups do not fail. For more information, see Data backup missing (backup history lost).
- If the database system administration (SUPERDBA) password is different in both the source and the target SAP MaxDB hosts, then when you perform an alternate restore of a specific data backup to the target SAP MaxDB host, the recovery run completes with the following errors:

  ```
  -24907,ERR_DBAWRONG: Wrong database system administrator
  -24988,ERR_SQL: SQL error
  -4008,Unknown user name/password combination
  ```

  To avoid the issue in the future, you must change the database system administrator password on the target SAP MaxDB host before the system copy to be consistent with the source SAP MaxDB host's password. Once the system copy completes, you can change the password.
- Ensure that the SAP MaxDB database versions in both the source and target systems are compatible.
- Cohesity supports backup and restore of SAP MaxDB databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.
- It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.

Download and install the Cohesity Linux Agent

The Cohesity Linux Agent is available as script-based and RPM-based installers, based on your requirements download the installer and install it accordingly on the SAP MaxDB host. For detailed instructions, see Download and Install the Linux Agent.

Download and Install the SAP MaxDB Connector Installer

Cohesity Platform provides a SAP MaxDB Connector agent for the backup and restore of enterprise SAP MaxDB deployments. The SAP MaxDB Connector agent leverages the backup and restore capabilities provided by SAP MaxDB.

Prerequisites

- Before you install the SAP MaxDB Connector agent, make sure that you:
- Install java runtime environment (JRE) version 1.8 or above on the SAP MaxDB Linux machine.
- Download and install the Cohesity Linux Agent on the SAP MaxDB Linux machine. For more information, see Download and install the Cohesity Linux Agent.
- Download and copy the SAP MaxDB Connector agent (cohesity-maxdb-java-backint-<version>.noarch.rpm) to the SAP MaxDB Linux machine. For more information, see Download the SAP MaxDB Connector agent.

- Use a user with sudo privileges or root user to install or uninstall the SAP MaxDB Connector agent.
- Use "bash" shell to install the SAP MaxDB Connector and to run backups and restore workflows of SAP MaXDB databases.
- Configure the following environment variables on the SAP MaxDB Linux machine:

Environment Variable	Description	Mandatory?	Command
`SERVICE_USER`	The Operation System (OS) user account that has access to the `dbmcli` binaries and can connect to the database.	Yes	`export SERVICE_USER=<VALUE>`
`CLUSTER_IPS`	A comma-separated list of virtual IP addresses of the servers in the Cohesity cluster.	No	`export CLUSTER_IPS="<VIP1>,<VIP2>,<VIP3>"`
`VIEW_NAME`	The name of the Cohesity view that you have created for SAP MaxDB. For more information, see Plan and Prepare for SAP MaxDB on Linux Protection.	No	`export VIEW_NAME=<cohesity-view-name>`
`XUSER_KEY`	The database key to connect to the SAP MaxDB database. You can create the database key using the following command: `xuser -U <KEY_NAME> -u <username>,<password> -d <MaxDB_Name>` For more information, see Creating or Changing XUSER entries.	Yes	`export XUSER_KEY=<KEY_NAME>`
WORKFLOW	Type of workflow that you want to use for SAP MaxDB data protection. For example, UDA or RA. The default value is **UDA**.	No	

Download the SAP MaxDB Connector

Cohesity Platform provides a SAP MaxDB Connector agent for SAP MaxDB to back up and restore SAP MaxDB databases. The SAP MaxDB Connector agent for SAP MaxDB is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for the Cohesity SAP MaxDB connector, and copy the datasource agent to the SAP MaxDB host.

Install the SAP MaxDB Connector

On the SAP MaxDB database host where you have copied the SAP MaxDB Connector agent (`cohesity-maxdb-java-backint-<version>.noarch.rpm`), perform the following to install the SAP MaxDB Connector agent.

To install the SAP MaxDB Connector agent:

1. Log on to the SAP MaxDB Linux machine using the user with `sudo` privileges.
2. Navigate to the directory where you have copied the RPM-based installer.
3. Install the SAP MaxDB Connector agent using the following command:

   ```
   sudo -E rpm -ivh cohesity-maxdb-java-backint-<version>.noarch.rpm
   ```

The following is a sample output of the successful installation:

```
$sudo -E rpm -ivh cohesity-maxdb-java-backint-<version>.noarch.rpm
Preparing...                          ##################################
[100%]
Pre-install/Upgrade section
==> Service user is vinadm
==> Cluster ips is 10.2.100.106
==> View Name is maxdb_data_view
==> Xuser key is XUSER_KEY
Precheck: linux_agent_exec pids are
23222 8352
==>
==> User vinadm does not exist in /etc/passwd
==> Checking if User vinadm exists using getent command.
==> User vinadm found with getent command
==> Package installation is happening at /opt/cohesity/maxdb_java_plugin
Updating / installing...
   1:cohesity-maxdb-java-backint-0.0.0##################################
[100%]
==> Post-install/Upgrade section
Service user is set to vinadm
------------------------------------------------------------------------
--
------------------------------------------------------------------------
--
```

```
==> Customizing the required files(updating templates).
CONTROL_NODES env is not set, hence using default value sap
SOURCE_NAME env is not set, hence using default value sap
------------------------------------------------------------------------
--
==> Adding 'cohesityagent' to sudoers list
==> Whitelisting following command /usr/bin/su
Creating soft link from /opt/cohesity/maxdb_java_plugin/uda_scripts to /op-
t/cohesity/agent/uda_scripts
Successful Installation of UDA MaxDB Java Agent
```

After the SAP MaxDB Connector agent is installed, the plug-ins are available at: `<backint_install_directory>/cohesity_backint_plug-in/uda_scripts`

For example:

```
/usr/sap/TEC/home/cohesity/cohesity_backint_plug-in/uda_scripts
```

Upgrade the SAP MaxDB Connector

Before upgrading the connector, ensure the following:

- Version 6.x uses Bash scripts and hence does not require Java.
- Version 7.x onwards uses Java and hence requires Java runtime.

To upgrade the SAP MaxDB Connector agent, run the following command with the appropriate parameters on SAP MaxDB host CLI:

```
sudo -E rpm -Uvh cohesity-maxdb-java-backint-<version>.noarch.rpm
```

Uninstall the SAP MaxDB Connector

To uninstall the SAP MaxDB Connector agent, run the following command with appropriate parameters on SAP MaxDB host CLI:

```
sudo -E rpm -evh cohesity-maxdb-java-backint-<version>.noarch.rpm
```

The following is a sample output of the successful uninstallation:

```
$sudo -E rpm -evh cohesity-maxdb-java-backint-<version>.noarch.rpm

Preparing...                          ##################################
[100%]
==> Pre-uninstall section
Cleaning up / removing...
   1:cohesity-maxdb-java-backint-ENG_1##################################
[100%]
==> Post-uninstall section
==> Deleting RPM install files at /opt/cohesity/maxdb_java_plugin
```

```
UDA MaxDB Java Agent is successfully uninstalled
```

Register and Manage the SAP MaxDB on Linux Source

To protect your SAP MaxDB databases, you need to register your SAP MaxDB deployment as a source on Cohesity. After registering your SAP MaxDB deployment as a source, you can:

- Update the username and password of the database user.
- Refresh the source details.
- Unregister the SAP MaxDB source from Cohesity.

Register SAP MaxDB as a Source

Before you register SAP MaxDB as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP MaxDB on Linux Protection.

To register SAP MaxDB as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
   a. From the **Source Type** drop-down menu, select **SAP MaxDB**.
   b. From the **Host OS Type** drop-down menu, select **Linux**.
   c. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the node(s) you have identified to run the Cohesity SAP MaxDB connector. For more information, see Plan and Prepare for SAP MaxDB on Linux Protection.

      > **Note:** Cohesity supports the backup and restore of SAP MaxDB databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

   d. In the **Datasource Agent Installation Path** field, enter the directory path on the SAP MaxDB node, where you have installed the Cohesity SAP MaxDB connector. For more information, see Install the SAP MaxDB Connector.
   e. Ignore the In the **APP Authentication** section. It is not applicable for SAP MaxDB source registration.
   f. Enable **Custom Options**, perform the following:

i. Ignore the **Mounts View** option. It is not applicable for SAP MaxDB source registration.

ii. Ignore the **Source Registration Options** field.

g. Click **Register**.

Update the SAP MaxDB Source Configuration

After registering your SAP MaxDB as a source on Cohesity, you can:

- Update the username and password of the database user.
- Modify the **Source Registration Options** based on your requirements.
- Update the **Datasource Agent Installation Path** only if you have uninstalled and reinstalled the Cohesity SAP MaxDB connector in a different directory on the SAP MaxDB node.

To update the SAP MaxDB source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP MaxDB source and select **Edit**.
3. On the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the SAP MaxDB Source Configuration

You can refresh the SAP MaxDB Source and re-verify the source with the available source registration parameters.

To refresh the SAP MaxDB Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP MaxDB Source and select **Refresh**.

Unregister the SAP MaxDB Source

If you plan to stop taking the backup of your SAP MaxDB databases, you can unregister the SAP MaxDB source from the Cohesity.

To unregister the SAP MaxDB source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP MaxDB source and select **Unregister**.

Backup SAP MaxDB on Linux Database

Cohesity provides a simple workflow to back up your SAP MaxDB databases. After registering the SAP MaxDB as a source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides a SAP MaxDB connector for backup and restore of enterprise SAP MaxDB deployments. The SAP MaxDB connector leverages the backup capabilities provided by SAP MaxDB.

After registering the SAP MaxDB source on Cohesity, you can create a Protection Group and specify the objects such as databases and tables that you want to backup. In the Protection Group, you need to define the parameters that enable you to provide additional parameters required by the backup statements provided by SAP MaxDB.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your SAP MaxDB backup configuration:

Backup Parameters	Mandatory?	Description
`--xuser-key=<key>`	Optional	By default, SAP MaxDB uses the XUSER key provided during the SAP MaxDB connector installation. You can override the XUSER key using the parameter based on your requirements. For example: `--xuser-key=XUSER_KEY`
`--backup-template`	Mandatory	Specify the backup template that you have created for the backup. Ensure that you create specific backup templates for COMPLETE DATA, INCREMENTAL DATA, AND LOG DATA. For detailed instructions to create a backup template, see backup_template_create on SAP documentation. For example: `--backup-template=BACKData`
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during backup. For example: `--use-secure-grpc=true` **Note:** It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.

Backup Parameters	Mandatory?	Description
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during backups. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg` Use the following command to generate the certificate: `~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name <clustername> -output_file /home/cohesity/<clustername>.cfg`

Protect SAP MaxDB on Linux Deployments

SAP MaxDB deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to:

- Specify the objects such as databases and tables that you want to backup.
- Based on your requirement, define the parameters for the backup statements provided by SAP MaxDB. For more information, see Backup Parameters.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP MaxDB deployments:

1. Do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the SAP MaxDB source and then select **Protect**.
2. From the **Source** drop-down list, select the registered SAP MaxDB source or click **Register Source** to register a SAP MaxDB source. For more information, see Register SAP MaxDB as a Source.
3. In the **Objects** section, add the object that you plan to protect. The object should be a database name. To back up a database, enter the name of the database. For example, `H00`. Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:

- Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
- Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.

5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the **Settings** section, do the following:
    1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
    2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

       Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
    3. In the **Backup Script Arguments** section, enter the Backup Parameters in the **Full Backup**, **Incremental Backup**, and **Log Backup** fields based on your requirements.

       For example:

Use Cases	Sample Parameters
Create a full backup of the SAP MaxDB data.	In the **Full Backup** field, specify the following parameter: `--backup-template=BACKData` Where, `BACKData` is the backup template you have created for complete data backup.
Create a full backup of the SAP MaxDB data by overriding the xuser key.	In the **Full Backup** field, specify the following parameter: `--xuser-key=XUER_KEY --backup-template=BACKData` Where `USER_KEY` is the new xuser key and `BACKData` is the backup template you have created for complete data backup.

Use Cases	Sample Parameters
Create an incremental backup of the SAP MaxDB data.	In the **Incremental Backup** field, specify the following parameter: `--backup-template=BACKPages` Where, `BACKPages` is the backup template you have created for incremental data backup.
Create an incremental backup of the SAP MaxDB data by overriding the xuser key.	In the Incremental Backup field, specify the following parameter: `--xuser-key=XUSER_KEY --backup-template=BACKPages` Where `XUSER_KEY` is the new `xuser` key and `BACKPages` is the backup template you have created for complete data backup.
Create a backup of the SAP MaxDB log data.	In the **Log Backup** field, specify the following parameter: `--backup-template=BACKLog` Where, `BACKLog` is the backup template you have created for log data backup.

    - Ignore the **Mounts** field. It is not applicable for SAP MaxDB data protection.
7. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click Edit on the right. For more information, see Additional Settings.
8. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Not applicable for SAP MaxDB data protection.
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs

Additional Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the SAP MaxDB on Linux Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP MaxDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP MaxDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP MaxDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover SAP MaxDB on Linux Database

You can restore the SAP MaxDB object(s) such as databases in SAP MaxDB. You can restore to the same SAP MaxDB cluster or a different SAP MaxDB cluster.

The SAP MaxDB connector leverages the restore statements provided by SAP MaxDB. You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

In the Recovery configuration, you need to specify the object you plan to restore and define the parameters required by the restore statements provided by SAP MaxDB.

Recovery Parameters

In the Recovery configuration, you can define the following parameters for your SAP MaxDB recovery:

Recovery Parameters	Mandatory?	Description
`--xuser-key=<key>`	Mandatory	SAP MaxDB uses the xuser key provided during the SAP MaxDB connector installation. You can override the xuser key using the parameter based on your requirements. For example: `--xuser-key=XUSER_KEY`
`--data-backup-template`	Mandatory	The template used to restore the data from the backup.
`--pages-backup-template`	Mandatory	The template used to restore the data from the incremental backup.
`--log-backup-template`	Mandatory for PITR use cases	The template used to restore the log data for the backup.

Recovery Parameters	Mandatory?	Description
`--source-sid=<SOURCE_SID>`	Mandatory for alternate restore to SAP MaxDB host with Different SID usecases	The source SAP MaxDB host SID.
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during restore. For example: `--use-secure-grpc=true`
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during restore. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`

Prerequisites

Ensure that the Protection Group associated with the object that you plan to restore is paused or in not running state.

Recover SAP MaxDB on Linux Databases

To recover SAP MaxDB databases:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character ***.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery. You can also select a point-in-time snapshot or specific data backup that you want to recover. Click the pencil icon and perform the following:
   - For point-in-time snapshots, in the **Edit recovery point for <Protection_Group_Name>** dialog, click the **Timeline** tab, choose the **Date** and select the **point-in-time** from the **timeline** slider and click **Select Recovery Point**.
   - For specific data backup snapshots, in the **Edit recovery point for <Protection_Group_Name>** dialog, click the **list** tab and select the specific **Snapshot Time** that you want and click **Select Recovery Point**.
6. Click **Next: Recover Options**.

7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same SAP MaxDB database.
New Object	To restore to a different SAP MaxDB database. Select the **Target** SAP MaxDB host to which the data needs to be restored. Also, enable **Overwrite existing object with the same name** and type **Yes**, if you want to overwrite the existing object with the same name. > **Note:** Ensure that you register the target SAP MaxDB host as source in Cohesity. For more information, see Register SAP MaxDB as a Source.

8. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Select the number of VIPs you want to use while writing data to the Cohesity storage.
Concurrency	Not applicable for SAP MaxDB data protection.
Custom Options	Specify the Recovery Parameters based on your requirements. For examples, see Sample Recovery Use Cases.
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

9. Click **Start Recovery**.

Sample Recovery Use Cases

Use Case	Sample Instructions
Restore specific data backup from the snapshot to the same SAP MaxDB database overriding the xuser key.	In the **Recover As** section, select **Replace Original**. Under the Recovery Options section, specify `--xuser-key=XUSER_KEY --data-backup-template=BACKData --pages-backup-template=BACKPages` in the **Custom Options** field. Where, • `XUSER_KEY` is the new xuser key. • `BACKData` is the backup template you have created for complete data backup. • `BACKPages` is the backup template you have created for the incremental data backup.
Point In Time Restore (PITR) to a new database.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** drop down. 2. Under the **Recovery Options** section, specify `--xuser-key=XUSER_KEY --data-backup-template=BACKData --pages-backup-template=BACKPages --log-backup-template=BACKPages` in the **Custom Options** field and specify the name for the new database in the **Rename** field. > **Note:** Ensure that the new database exists on the target SAP MaxDB host. Where, • `XUSER_KEY` is the `xuser` key on the target SAP MaxDB host. • `BACKData` is the back template you have created for complete data backup. • `BACKPages` is the back template you have created for incremental data backup. • `BACKLog` is the back template you have created for log data backup.

Use Case	Sample Instructions
Restore to a new database on a different SAP MaxDB host.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown. 2. Under the **Recovery Options** section, specify `--xuser-key=XUSER_KEY --data-backup-template=BACKData --pages-backup-template=BACKPages` in the **Custom Options** field and specify the name for the new database in the **Rename** field. > **Note:** Ensure that the new database exists on the target SAP MaxDB host. Where: • `XUSER_KEY` is the `xuser` key on the target SAP MaxDB host. • `BACKData` is the back template you have created for complete data backup. • `BACKPages` is the back template you have created for incremental data backup.

Alternate Restore to SAP MaxDB Host with Different SIDs

You can perform an alternate restore to SAP MaxDB hosts with different SIDs.

Prerequisites

Ensure that you configure the source and target SAP MaxDB host's Cohesity BACKINT parameter files (`/opt/cohesity/maxdb_java_plugin/profile/init<SID>.utl`) with the same Cohesity cluster and Cohesity View details.

To perform an alternate restore to SAP MaxDB host with different SIDs:

1. Using the Command-Line Interface (CLI) or SAP MaxDB Studio, take backups on the source SAP MaxDB host.
2. Copy the recover catalog files from the source SAP MaxDB host to a temporary directory (/tmp/logs) on the target SAP MaxDB host.

   You can locate the recover catalog files using the param_directget command.

   For example:

   ```
   >dbmcli – u <username>,<password> –d <DBname> param_directget
   RunDirectoryPath
   ```

3. On the target SAP MaxDB host, update the temporary directory path in the `ORIGINAL_RUNDIRECTORY` parameter in the `/opt/cohesity/maxdb_java_plugin/profile/bsi.env` file.

   For example:

```
bash-4.2$ cat /opt/cohesity/maxdb_java_plugin/profile/bsi.env
BACKINT /opt/cohesity/maxdb_java_plugin/bin/backint.sh
INPUT /opt/cohesity/maxdb_java_plugin/workflow/input
OUTPUT /opt/cohesity/maxdb_java_plugin/workflow/output
ERROROUTPUT /opt/cohesity/maxdb_java_plugin/workflow/error
PARAMETERFILE /opt/cohesity/maxdb_java_plugin/profile/initSID.utl
TIMEOUT_SUCCESS 600
TIMEOUT_FAILURE 300
CHECK_AVAILABILITY YES
ORIGINAL_RUNDIRECTORY /tmp/logs
```

4. Restore to the target SAP MaxDB host using the restore commands.

   Sample for full database backup:

```
>db_activate recover BACKData DATA Externalbackupid "MAXDB 1630306452_
v48g /tmp/maxdb/pipe_data"
OK
Returncode 0
…
…
```

   Sample for incremental backup:

```
>recover_start BACKPages PAGES EBID "MAXDB 1630306540_bzx6
/tmp/maxdb/pipe_pages" [18/979]
OK
Returncode 0
…
…
```

   Sample for log backups:

```
>recover_start  BACKLog LOG EBID "MAXDB 1632394917_0af3 /tmp/maxdb_log_
pipe" UNTIL 20210923 163300
OK
Returncode              -8020
…

>recover_replace  BACKLog  ExternalBackupID "MAXDB 1632395036_yxg0
/tmp/maxdb_log_pipe"
```

```
OK
Returncode                -8020
```

> **Note:** Execute the following command, if the database is not online:
>
> ```
> >recover_ignore
> OK
> Returncode 0
> ```

5. Verify if the restored database is online using the following command:

```
>db_online
```

SAP MaxDB on Windows

Cohesity DataPlatform natively integrates with SAP MaxDB to provide backup and recovery for clustered SAP MaxDB deployments. Using Cohesity Connector Agent, you can backup and restore the SAP MaxDB database. The Cohesity Connector Agent provides a simple workflow to back up and restore your SAP MaxDB deployment.

Cohesity DataPlatform provides an SAP MaxDB Connector for backup and restore of enterprise SAP MaxDB deployments. The SAP MaxDB Connector leverages the backup and restore capabilities provided by SAP MaxDB.

The SAP MaxDB Connector enables you to back up the various databases. It provides:

- **Backup**:
  - Full backup
  - Incremental backup
  - Scheduled Log backup
- **Restore**:
  - Restore the database to a specific data backup
  - Restore the database to the point in time
  - Alternate restore of SAP MaxDB database
    - Alternate restore to a new database or active database on the same SAP MaxDB system.
    - Alternate restore to a new or active database on a different SAP MaxDB system.

Core Components

The following are the core components for protecting SAP MaxDB databases:

- SAP MaxDB Windows machine
- SAP MaxDB Connector
- Cohesity Connector Agent

As illustrated, you can integrate the Cohesity DataPlatform with your SAP MaxDB deployments using the SAP MaxDB Connector and Cohesity Connector Agent.

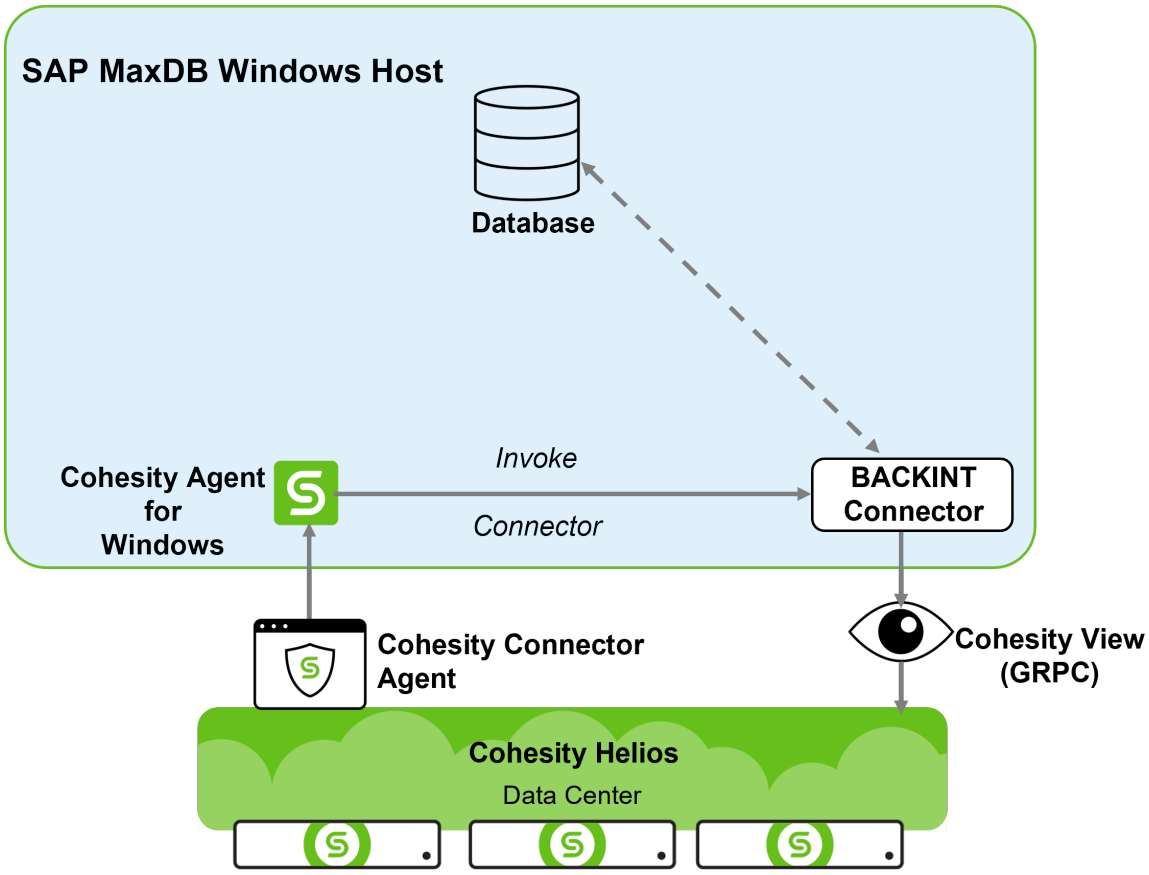

Plan and Prepare for SAP MaxDB on Windows Protection

Before you register your SAP MaxDB deployment as a source with Cohesity and protect SAP MaxDB databases, ensure the following prerequisites:

- Ensure that the fully qualified domain name (FQDN) of the Cohesity cluster is resolvable from the SAP MaxDB Windows machine.
- Synchronize the SAP MaxDB Windows machine time with a Network Time Protocol (NTP) server.
- A database user with database admin user privileges.
- Download and install the Cohesity Windows Agent on the SAP MaxDB Windows machine.

- Download and install the SAP MaxDB Connector.
- Review and understand the limitations.

Supported Versions

For more information on supported SAP MaxDB versions, see Supported Software versions.

SAP MaxDB Connector Requirements

To install and run the Cohesity SAP MaxDB Connector, ensure that the SAP MaxDB host meets the following specifications:

Components	Specification
SAP MaxDB Version	For details, see Supported Software.
Operating System	
Powershell Version	5.1.14393.0
Platform	x86_64
Disk Space	• Minimum 256 MB. Required for installing the SAP MaxDB Connector. • 1 GB. Required for the log directory.
Memory	8GB of free memory to run the SAP MaxDB Connector agent.

Port Requirements

You must open certain ports in the firewall to allow the Cohesity cluster to transmit and receive data. Ensure that you open the following port in the firewall to allow data transfer between SAP MaxDB system and the Cohesity cluster:

Port	Source	Target	Direction	Network Protocol
11113	SAP MaxDB	Cohesity cluster	Bidirectional	TCP/IP

Considerations

Review and understand the following before you protect your SAP MaxDB database:

- SAP MaxDB with liveCache technology is not supported.
- Spaces in the installation path are not supported.

- Only PIPE backups are supported. File backups are not supported, and hence only scheduled log backups are supported, and automatic log backups are not supported.
- Only backup and restore to the same Operating System is supported.
- Only the single-node SAP MaxDB database system is supported.
- Cohesity replication is supported.
- For alternate restore, the target SAP MaxDB database must be created before starting the restore operation.
- If the database system administration (SUPERDBA) password is different in both the source and the target SAP MaxDB hosts, then when you perform an alternate restore of a specific data backup to the target SAP MaxDB host, the recovery run completes with the following errors:

  ```
  -24907,ERR_DBAWRONG: Wrong database system administrator
  -24988,ERR_SQL: SQL error
  -4008,Unknown user name/password combination
  ```

  To avoid the issue in the future, you must change the database system administrator password on the target SAP MaxDB host before the system copy to be consistent with the source SAP MaxDB host's password. Once the system copy is completed, you can change the password.
- Ensure that the SAP MaxDB database versions in both the source and target systems are compatible.
- It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.

Download and install the Cohesity Windows Agent

The Cohesity Windows Agent is available as script-based and RPM-based installers, based on your requirements download the installer and install it accordingly on the SAP MaxDB host. For detailed instructions, see Install and Manage the Agent on Windows Servers.

Download and Install the SAP MaxDB Connector Installer

Cohesity DataPlatform provides an SAP MaxDB Connector for backup and restore of enterprise SAP MaxDB deployments. The SAP MaxDB Connector leverages the backup and restore capabilities provided by SAP MaxDB.

> **Note:** Ensure that you install the Cohesity Windows Agent using the ***<SID>adm*** user account.

Prerequisites

Before you install the SAP MaxDB Connector, make sure that you:

- Install java runtime environment (JRE) version 1.8 or above on the SAP MaxDB Windows machine.
- Download and install the Cohesity Windows Agent on the SAP MaxDB Windows machine. For more information, see Install and Manage the Agent on Windows Servers.
- Download and copy the SAP MaxDB Connector (`Cohesity_SAP_MaxDB_connector_<version>_Win_x64_Installer.exe`) to the SAP MaxDB Windows machine. For more information, see Download the SAP MaxDB Connector.
- Use an **<sid>adm** user to install or uninstall the SAP MaxDB Connector.

Download the SAP MaxDB Connector

Cohesity provides a SAP MaxDB Connector to backup and restore SAP MaxDB databases. The SAP MaxDB Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for SAP MaxDB Connector, and copy the datasource agent to the SAP MaxDB host.

Install the SAP MaxDB Connector

On the SAP MaxDB database host where you have copied the SAP MaxDB Connector (`Cohesity_SAP_MaxDB_connector_<version>_Win_x64_Installer.exe`), perform the following to install the SAP MaxDB Connector.

To install the SAP MaxDB Connector:

1. Log on to the SAP MaxDB Windows machine using the **<sid>adm** user. Ensure that you are logged as the same user with which you installed the Cohesity Windows agent.
2. Navigate to the directory where you have copied the installer.
3. Double-click on the Cohesity SAP MaxDB Connector installer and follow the instructions.

Silently Install the SAP MaxDB Connector

After you have downloaded the SAP MaxDB Connector, you can perform a silent installation of the Cohesity SAP MaxDB Connector.

To perform the silent installation:

1. Log on to the SAP MaxDB node where you have copied the Cohesity SAP MaxDB Connector installer executable, `Cohesity_SAP_MaxDB_connector_<version>_Installer.exe`. Ensure that you are logged as the same user with which you installed the Cohesity Windows agent.
2. Using the command-line interface, navigate to the directory where you have copied the Cohesity SAP MaxDB Connector installer executable, `Cohesity_SAP_MaxDB_connector_<version>_Installer.exe`.
3. Use the following command for the silent installation:

```
<InstallerExeFile> [/log="<absolute-path-to-log-file>"]
[/dir="<absolute-path-to-install-location>"] [/silent, /verysilent]
[/suppressmsgboxes] [/norestart]
```

Where:

- `/log="<absolute-path-to-log-file>"` - Records the events during installation and helps to debug the errors, if any, during the installation process.
- `/silent` - Suppresses the installation wizard and background window. Instead, it only displays the installation progress window.
- `/verysilent` - Suppresses the installation progress window.
- `/suppressmsgboxes` - Suppresses any message boxes if `/verysilent` is specified.
- `/norestart` - Prevents the setup from restarting the system following a successful installation, or after a Preparing to Install failure that requests a restart.
- `/dir="<absolute-path-to-install-location>"` - Specify the directory path where the SAP MaxDB Connector for Windows should be installed.

  > **Note:** The installer creates the necessary folders if they do not exist. If this parameter is not specified, the SAP MaxDB Connector is installed in the default directory path: "`%SystemDrive%\Program Files\Cohesity\SAP MaxDB`".

Sample:

```
Cohesity_SAP_MaxDB_connector_<version>_Installer.exe
/log="C:\tmp\install_logs.txt" /verysilent /suppressmsgboxes /norestart
```

Upgrade the SAP MaxDB Connector

Before upgrading the connector, ensure the following:

- Version 6.x uses Bash scripts and hence does not require Java.
- Version 7.x onwards uses Java and hence requires Java runtime.

On the SAP MaxDB database host where you have copied the SAP MaxDB Connector (`Cohesity_SAP_MaxDB_connector_<version>_Win_x64_Installer.exe`), perform the following to upgrade the SAP MaxDB Connector.

Before you upgrade the SAP MaxDB Connector, ensure that you pause the existing protection runs.

To upgrade the SAP MaxDB Connector:

1. Log on to the SAP MaxDB Windows machine using the **<sid>adm** user. Ensure that you are logged as the same user with which you installed the Cohesity Windows agent.
2. Navigate to the directory where you have copied the installer.
3. Double-click on the Cohesity SAP MaxDB Connector installer, and when prompted, select **Upgrade** and follow the instructions.

Silently Upgrade the SAP MaxDB Connector

Before you upgrade the SAP MaxDB Connector, ensure that you pause the existing protection runs.

To silently upgrade the SAP MaxDB Connector, use the following command:

```
<InstallerExeFile> [/log="<absolute-path-to-log-file>"] [/dir="<absolute-path-to-install-location>"] [/silent, /verysilent] [/suppressmsgboxes] [/norestart]
```

1. Where:
   - `/log="<absolute-path-to-log-file>"` - Records the events during the upgrade and helps to debug the errors, if any, during the upgrade process.
   - `/silent` - Suppresses the upgrade wizard and background window. Instead, it only displays the upgrade progress window.
   - `/verysilent` - Suppresses the upgrade progress window.
   - `/suppressmsgboxes` - Suppresses any message boxes if `/verysilent` is specified.
   - `/norestart` - Prevents the setup from restarting the system following a successful upgrade or after a Preparing to upgrade failure that requests a restart.
   - `/dir="<absolute-path-to-install-location>"` - Specify the directory path where the SAP MaxDB Connector is installed.

   Sample:

   ```
   Cohesity_SAP_MaxDB_connector_<version>_Installer.exe /log="C:\tmp\upgrade_logs.txt" /verysilent /suppressmsgboxes /norestart
   ```

Uninstall the SAP MaxDB Connector

To uninstall the Cohesity SAP MaxDB Connector:

1. Log on to the SAP on MaxDB host.
2. Select the **Start** menu and look for the SAP MaxDB Connector in the list shown.
3. Press and hold (or right-click) on the SAP MaxDB Connector, then select **Uninstall**.

Silently Uninstall the SAP MaxDB Connector

To perform silent uninstallation:

1. Log on to the SAP MaxDB host.
2. Using the command-line interface, navigate to the directory where you have installed the Cohesity SAP MaxDB Connector, For example, `C:\Program Files\Cohesity\SAP MaxDB\Installer\`.

   > **Note:** The details of uninstaller executable are available in the registry value `"HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\{CCEC1E9A-895C-4038-8CC9-31BBD31B9CA7}_is1\UninstallString"`. The default value is `"C:\Program Files\Cohesity\SAP MaxDB\Installer\unins000.exe"`, assuming `C:` is the default operating system installation volume and `"C:\Program Files\Cohesity\SAP MaxDB"` is the default SAP MaxDB Connector installation path.

3. Use the following command for the silent uninstallation:

   ```
   <UninstallerExeFile>  [/log="<absolute-path-to-log-file>"] [/silent, /verysilent] [/suppressmsgboxes] [/norestart] [/preservesettings]
   ```

   Where:

   - `/log="<absolute-path-to-log-file>"` - Records the events during the uninstallation process and helps to debug the errors, if any, during the uninstallation process.
   - `/silent` - Suppresses the uninstallation wizard and background window. Instead, it only displays the uninstallation progress window.
   - `/verysilent` - Suppresses the uninstallation progress window.
   - `/suppressmsgboxes` - Suppresses any message boxes if `/verysilent` is specified.
   - `/norestart` - Prevents the setup from restarting the system following a successful uninstallation.
   - `[/preservesettings]` - Preserve the registry settings when uninstalling the SAP MaxDB Connector.

> **Note:** Uninstalling the SAP MaxDB Connector clears registration settings from the registry. You can preserve the registration settings in the registry by using the `/preservesettings` flag when uninstalling from the command line. Preserving registry settings are not supported when uninstalling from the GUI.

Sample:

```
unins000.exe /log="C:\tmp\uninstall_logs.txt" /verysilent
/suppressmsgboxes /norestart
```

Register and Manage the SAP MaxDB on Windows Source

To protect your SAP MaxDB databases, you need to register your SAP MaxDB deployment as a source on Cohesity. After registering your SAP MaxDB deployment as a source, you can:

- Update the username and password of the database user.
- Refresh the source details.
- Unregister the SAP MaxDB source from Cohesity.

Register SAP MaxDB as a Source

Before you register SAP MaxDB as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP MaxDB on Windows Protection.

To register SAP MaxDB as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
    a. From the **Source Type** drop-down menu, select **SAP MaxDB**.
    b. From the **Host OS Type** drop-down menu, select **Windows**.
    c. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the node(s) you have identified to run the SAP MaxDB connector. For more information, see Plan and Prepare for SAP MaxDB on Windows Protection.
    d. In the **Datasource Agent Installation Path** field, enter the directory path on the SAP MaxDB node, where you have installed the SAP MaxDB connector. For

more information, see Install the SAP MaxDB Connector.

e. Ignore the **APP Authentication** section. It is not applicable for SAP MaxDB source registration.

f. Enable **Custom Options**, perform the following:

i. Ignore the **Mounts View** option. It is not applicable for SAP MaxDB source registration.

ii. Ignore the **Source Registration Options** field.

g. Click **Register**.

Update the SAP MaxDB Source Configuration

After registering your SAP MaxDB as a source on Cohesity, you can:

- Update the username and password of the database user.
- Modify the **Source Registration Options** based on your requirements.
- Update the **Datasource Agent Installation Path** only if you have uninstalled and reinstalled the SAP MaxDB connector in a different directory on the SAP MaxDB node.

To update the SAP MaxDB source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP MaxDB source and select **Edit**.
3. On the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the SAP MaxDB Source Configuration

You can refresh the SAP MaxDB source and re-verify the source with the available source registration parameters.

To refresh the SAP MaxDB Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP MaxDB source and select **Refresh**.

Unregister the SAP MaxDB Source

If you plan to stop taking the backup of your SAP MaxDB databases, you can unregister the SAP MaxDB source from Cohesity.

To unregister the SAP MaxDB source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP MaxDB source and select **Unregister**.

Backup SAP MaxDB on Windows Database

Cohesity provides a simple workflow to back up your SAP MaxDB databases. After registering the SAP MaxDB as a source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides an SAP MaxDB Connector for the backup and restore of enterprise SAP MaxDB deployments. The SAP MaxDB Connector leverages the backup capabilities provided by SAP MaxDB.

After registering the SAP MaxDB source on Cohesity, you can create a Protection Group and specify the objects such as databases and tables that you want to backup. In the Protection Group, you need to define the parameters that enable you to provide additional parameters required by the backup statements provided by SAP MaxDB.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your SAP MaxDB backup configuration:

Backup Parameters	Mandatory?	Description
`--xuser-key=<key>`	Optional	By default, SAP MaxDB uses the xuser key provided during the SAP MaxDB connector installation. You can override the xuser key using the parameter based on your requirements. For example: `--xuser-key=XUSER_KEY`
`--backup-template`	Mandatory	Specify the backup template that you have created for the backup. Ensure that you create specific backup templates for COMPLETE DATA, INCREMENTAL DATA, AND LOG DATA. For detailed instructions to create a backup template, see backup_template_create on SAP documentation. By default, SAP MaxDB uses the 'BACKData' template for full backup, 'BACKPages' template name for incremental backup, 'BACKLog' template name for log backups. For example: `--backup-template=BACKData`

Backup Parameters	Mandatory?	Description
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during backup. For example: `--use-secure-grpc=true` **Note:** It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during backups. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg` Use the following command to generate the certificate: `~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name <clustername> -output_file /home/cohesity/<clustername>.cfg`

Protect SAP MaxDB on Windows Deployments

SAP MaxDB deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to:

- Specify the objects such as databases and tables that you want to backup.
- Based on your requirement, define the parameters for the backup statements provided by SAP MaxDB. For more information, see Backup Parameters.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP MaxDB deployments:

1. Do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the SAP MaxDB source and then select **Protect**.

2. From the **Source** drop-down list, select the registered SAP MaxDB source or click **Register Source** to register a SAP MaxDB source. For more information, see Register SAP MaxDB as a Source.
3. In the **Objects** section, add the object that you plan to protect. The object should be a database name. To back up a database, enter the name of the database. For example, `H00`. Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the Settings section, do the following:
   1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
   2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

      Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
   3. In the **Custom Options** section, enter the Backup Parameters in the **Full Backup**, **Incremental Backup**, and **Log Backup** fields based on your requirements.

      For example:

Use Cases	Sample Parameters
Create a full backup of the SAP MaxDB data.	In the **Full Backup** field, specify the following parameter: `--backup-template=BACKData` Where, `BACKData` is the back template you have created for complete data backup.

Use Cases	Sample Parameters
Create a full backup of the SAP MaxDB data by overriding the xuser key.	In the **Full Backup** field, specify the following parameter: `--xuser-key=XUER_KEY --backup-template=BACKData` Where `USER_KEY` is the new xuser key and `BACKData` is the back template you have created for complete data backup.
Create an incremental backup of the SAP MaxDB data.	In the **Incremental Backup** field, specify the following parameter: `--backup-template=BACKPages` Where, `BACKPages` is the back template you have created for incremental data backup.
Create an incremental backup of the SAP MaxDB data by overriding the xuser key.	In the Incremental Backup field, specify the following parameter: `--xuser-key=XUSER_KEY --backup-template=BACKPages` Where `XUSER_KEY` is the new `xuser` key and `BACKPages` is the back template you have created for complete data backup.
Create a backup of the SAP MaxDB log data.	In the **Log Backup** field, specify the following parameter: `--backup-template=BACKLog` Where, `BACKLog` is the back template you have created for log data backup.

- Ignore the **Mounts** field. It is not applicable for SAP MaxDB data protection.

7. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click Edit on the right. For more information, see Additional Settings.
8. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	The maximum number of streams that are created to exchange data with the cluster. By default, maximum scaling is set to empty.
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the SAP MaxDB on Windows Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP MaxDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ on the **Objects** section.

5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP MaxDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. Navigate to **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.

- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP MaxDB.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Protection Group Only
   - Delete Protection Group and Snapshots

Purge SAP MaxDB on Windows Backups

Cohesity cluster provides an in-built purge functionality that you can use to automatically remove the expired data and log backups from the Cohesity cluster.

The purge operation works based on the retention period that you have configured on the Protection policy associated with the data backup. For information on how to set the retention period for a data backup, see About Policies and Protection Groups.

> **Note:** The retention period you have configured for log backups is not honored because only the last successful full backup is considered based on the retention period. Any backups older than the full backup are deleted.

For example, In the protection policy, if you have configured to take backup every day, the retention period is 14 days. Assuming that the current date is November 26th, in that case, the last successful full backup on or before November 12th is retained, and any backups older than the full backup are deleted.

Recover SAP MaxDB on Windows Database

You can restore the SAP MaxDB object(s) such as databases in SAP MaxDB. You can restore to the same SAP MaxDB cluster or a different SAP MaxDB cluster.

The SAP MaxDB Connector leverages the restore statements provided by SAP MaxDB. You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

In the Recovery configuration, you need to specify the object you plan to restore and define the parameters required by the restore statements provided by SAP MaxDB.

Recovery Parameters

In the Recovery configuration, you can define the following parameters for your SAP MaxDB recovery:

Recovery Parameters	Mandatory?	Description
`--xuser-key=<key>`	Mandatory	SAP MaxDB uses the xuser key provided during the SAP MaxDB connector installation. You can override the xuser key using the parameter based on your requirements. For example: `--xuser-key=XUSER_KEY`
`--data-backup-template`	Mandatory	The template used to restore the data from the backup.
`--pages-backup-template`	Mandatory	The template used to restore the data from the incremental backup.
`--log-backup-template`	Mandatory for PITR use cases	The template used to restore the log data for the backup.
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during restore. For example: `--use-secure-grpc=true`
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during restore. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`

Prerequisites

Ensure that the Protection Group associated with the object that you plan to restore is paused or in not running state.

Recover SAP MaxDB on Windows Databases

To restore SAP MaxDB databases:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery. You can also select a point-in-time snapshot or specific data backup that you want to recover. Click the pencil icon and perform the following:
   - For point-in-time snapshots, in the **Edit recovery point for <Protection_Group_Name>** dialog, click the **Timeline** tab, choose the **Date** and select the **point-in-time** from the **timeline** slider and click **Select Recovery Point**.
   - For specific data backup snapshots, in the **Edit recovery point for <Protection_Group_Name>** dialog, click the **list** tab and select the specific **Snapshot Time** that you want and click **Select Recovery Point**.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same SAP MaxDB database.
New Object	To restore to a different SAP MaxDB database. Select the **Target** SAP MaxDB host to which the data needs to be restored. Also, enable **Overwrite existing object with the same name** and type **Yes**, if you want to overwrite the existing object with the same name. **Note:** Ensure that you register the target SAP MaxDB host as source in Cohesity. For more information, see Register SAP MaxDB as a Source.

8. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Select the number of VIPs you want to use while writing data to the Cohesity storage.
Concurrency	The maximum number of streams that are created to exchange data with the cluster. By default, maximum scaling is set to empty.
Recovery Script Arguments	Specify the Recovery Parameters based on your requirements. For examples, see Sample Recovery Use Cases.
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

9. Click **Start Recovery**.

Sample Recovery Use Cases

Use Case	Sample Recovery Parameters
Restore specific data backup from the snapshot to the same SAP MaxDB database overriding the xuser key.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, specify `--xuser-key=XUSER_KEY --data-backup-template=BACKData --pages-backup-template=BACKPages` in the **Recovery Script Arguments** field. Where, • `XUSER_KEY` is the new xuser key. • `BACKData` is the backup template you have created for complete data backup. • `BACKPages` is the backup template you have created for the incremental data backup.

Use Case	Sample Recovery Parameters
Point in Time Restore (PITR) same SAP MaxDB database.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, specify `--xuser-key=XUSER_KEY --data-backup-template=BACKData --pages-backup-template=BACKPages --log-backup-template=BACKLog` in the **Recovery Script Arguments** field. Where: • `XUSER_KEY` is the new xuser key. • `BACKData` is the backup template you have created for complete data backup. • `BACKPages` is the backup template you have created for the incremental data backup. • `BACKLog` is the back template you have created for log data backup.
Restore specific data backup to a new database on a different SAP MaxDB host.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown. 2. Under the **Recovery Options** section, specify `--xuser-key=XUSER_KEY --data-backup-template=BACKData --pages-backup-template=BACKPages` in the **Recovery Script Arguments** field and specify the name for the new database in the **Rename** field. > **Note:** Ensure that the new database exists on the target SAP MaxDB host. Where, • `XUSER_KEY` is the `xuser` key on the target SAP MaxDB host. • `BACKData` is the back template you have created for complete data backup. • `BACKPages` is the back template you have created for incremental data backup. • `BACKLog` is the back template you have created for log data backup.

Use Case	Sample Recovery Parameters
Point In Time Restore (PITR) to a new database.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown. 2. Under the **Recovery Options** section, specify `--xuser-key=XUSER_KEY --data-backup-template=BACKData --pages-backup-template=BACKPages --log-backup-template=BACKPages` in the **Recovery Script Arguments** field and specify the name for the new database in the **Rename** field. > **Note:** Ensure that the new database exists on the target SAP MaxDB host. Where, • `XUSER_KEY` is the `xuser` key on the target SAP MaxDB host. • `BACKData` is the back template you have created for complete data backup. • `BACKPages` is the back template you have created for incremental data backup. • `BACKLog` is the back template you have created for log data backup.

SAP Sybase ASE on Linux

Cohesity Platform natively integrates with SAP Sybase ASE to provide backup and recovery for SAP Sybase ASE deployments. Using Cohesity Connector Agent, you can backup and restore the SAP Sybase ASE databases.

Cohesity Platform provides a SAP Sybase ASE connector that enables you to perform a full backup, differential backup, log backup, restore and point in time restore at the database level. The SAP Sybase ASE connector leverages the third-party dynamic link library (DLL) based capabilities provided by SAP Sybase ASE for backup and restore functionality.

Core Components

The following are the core components for protecting SAP Sybase ASE databases:

- SAP Sybase ASE Linux Machine
- Cohesity SAP Sybase ASE connector
- Cohesity Connector Agent

As illustrated, you can integrate the Cohesity Platform with your SAP Sybase ASE deployments using the SAP Sybase ASE connector and Cohesity Connector Agent.

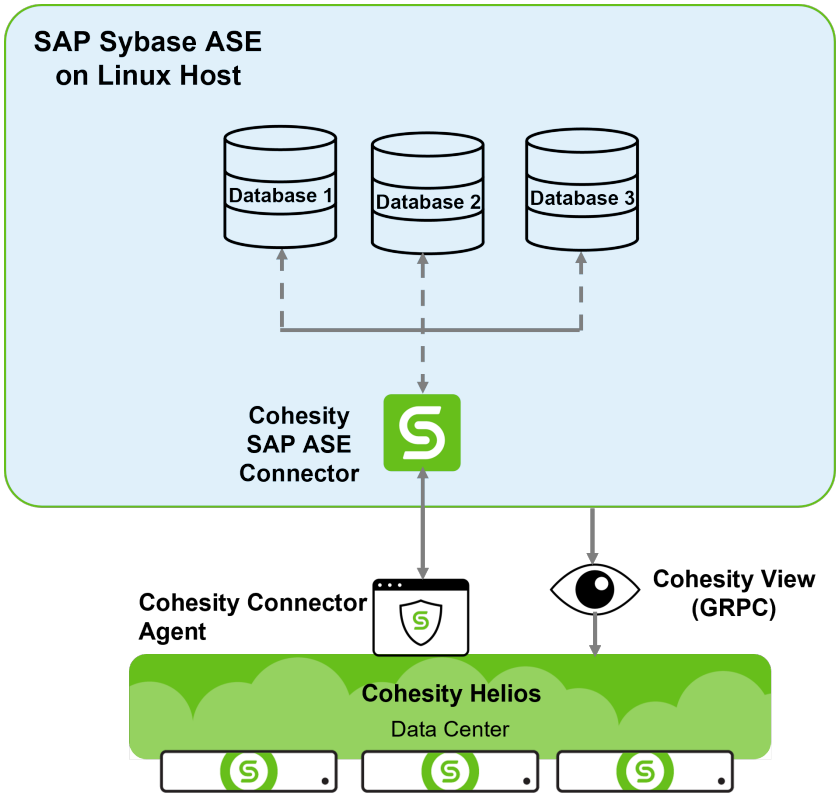

Plan and Prepare for SAP Sybase ASE on Linux Protection

Before you register your SAP Sybase ASE deployment as a source with Cohesity Platform, ensure the following prerequisites:

- The user account for the SAP Sybase ASE instance must have either database administrator or operator privileges. If administrator privileges cannot be granted, the following permissions are required along with operator privileges:
  - Create Database
  - Own Any Database
  - Manage Disk
- Download and install the Cohesity Linux Agent RPM on the SAP Sybase ASE host.
- Download and install the Cohesity SAP Sybase ASE Connector RPM on the SAP Sybase ASE host

Supported Versions

For more information on supported SAP Sybase ASE versions, see SAP Sybase ASE versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	SAP Sybase ASE Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP Sybase ASE Local Host	SAP Sybase ASE Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP Sybase ASE Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
SAP Sybase ASE Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following requirements before you protect your single node SAP Sybase ASE deployment.

- Backup or restore for SAP Sybase ASE HA cluster deployments are not supported.
- Ensure that the database is not in use (active connections) during recovery.
- Once Cohesity protects the Sybase ASE database, it is not recommended to use the SAP backup and/or restore tools outside of Cohesity, including native dump or load tools provided by SAP Sybase ASE.
- The same database should not be part of multiple Protection Groups.
- If the hole is introduced in the backup timeline, then point-in-time recovery (PITR) is possible only after the subsequent successful full or incremental backup is completed.
- Cohesity supports backup and restore of SAP Sybase ASE databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

Prerequisites

Ensure that:

- You use root user or non-root user to install the Cohesity Linux agent and SAP ASE Connector. If you choose to use a non-root user, you must add the `sudo` privileges for the non-root user using the syntax given below:

  ```
  <user_name> ALL=(ALL) NOPASSWD:SETENV: /path/to/command1,
  /path/to/command2, /path/to/command3
  ```

  The following is the `sudo` commands for SAP ASE Connector:

  ```
  cohesityagent ALL=
  NOPASSWD:SETENV:/usr/bin/mkdir,/usr/bin/chmod,/usr/bin/chown,/usr/bin/
  cmp,/opt/cohesity/sapase/.temp.*/run_ase_
  connector.java,/opt/sap/shared/SAPJRE-8_1_097_
  64BIT/bin/java,/usr/bin/cp,/usr/bin/find,/usr/bin/kill,/usr/bin/kinit,
  /usr/bin/mv,/usr/bin/rm,/usr/bin/tee,/usr/bin/test
  ```

  Where, **/opt/cohesity/sapase** is the SAP ASE connector installation directory and **/opt/sap/shared/SAPJRE-8_1_097_64BIT** is the JRE-8 environment directory within the SAP ASE database installation directory.

  > **Note:**
  > - If the **cohesityagent** user entry is already present in **/etc/sudoers**, append the list of commands to the existing entry.
  > - The actual paths for each command might vary depending on the platform and distribution. For more information, see Download and Install the Linux Agent.

- You enable incremental backups on the SAP Sybase ASE host using the following database command:

  ```
  sp_dboption '<DBNAME>','allow incremental dumps', 'true'
  ```

  > **Note:** After you have enabled incremental backups, ensure that you first initiate a full backup through the Cohesity Dashboard.

- If you want to authenticate to the Cohesity cluster using key-based authentication, ensure that you install the SAP Sybase ASE Connector using the same user running the ASE server and Cohesity Agent.
- Ensure that the SAP ASE environment file path (`SAPASE_ENV_PATH`) has following environment variables:

Environment Variable	Description
`SYBASE`	The path of the Sybase installation directory.
`SYBASE_ASE`	The subdirectory of the Adaptive Server component.
`SAP_JRE8_64`	The location of the Java Runtime Environment.
`SAP_JCONNECT`	The Sybase JDBC driver.

- To securely connect to the Cohesity cluster while reading or writing data, you can use secure gRPC.

  > **Note:**
  > - It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
  > - Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
  > - Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

  To use secure gRPC, perform the following:

  1. Generate the certificate config file on the Cohesity node using the following command:

     ```
     ~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name <clustername> -output_file /home/cohesity/<clustername>.cfg
     ```

  2. Copy `/home/cohesity/<clustername>.cfg` from the Cohesity node to all SAP ASE nodes at the same location. For example, `/opt/certs/<clustername>.cfg`.

     > **Note:** Ensure that the `root` or `non-root` user used to install the Cohesity Linux agent and SAP ASE Connector has read access to the generated certificate config file.

  3. While registering the SAP Sybase ASE source, provide the path of the certificate config file in the **Cohesity SSL Settings** field.

4. Increase the SAP Sybase ASE backup server memory to 1GB. For more information, see Utility Commands Reference in SAP Documentation.

Download and install the Cohesity Linux Agent

The Cohesity Linux Agent is available as an RPM installer. Download this installer and install it accordingly on the SAP Sybase ASE hosts to be protected. For detailed instructions, see Download and Install the Linux Agent.

> **Note:** Ensure that you install the Cohesity Linux Agent using the RPM-based installer.

Download and Install the SAP Sybase ASE Connector Installer

Cohesity Platform provides a SAP Sybase ASE Connector for backup and restore of enterprise SAP Sybase ASE deployments.

Download the SAP Sybase ASE Connector

Cohesity Platform provides a SAP Sybase ASE Connector to backup and restore SAP Sybase ASE databases. The SAP Sybase ASE Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for SAP Sybase ASE Connector, and copy the datasource agent to the SAP Sybase ASE host.

Install the SAP Sybase ASE Connector

On the SAP Sybase ASE database host where you have copied the SAP Sybase ASE Connector, perform the following to install the SAP Sybase ASE Connector.

To install the SAP Sybase ASE connector:

1. Log on to the SAP Sybase ASE Linux machine using the user with sudo privileges.
2. Navigate to the directory where you have copied the RPM-based installer.
3. Install the Cohesity SAP Sybase ASE connector using the following command:

   ```
   sudo sh -c 'export SAPASE_ENV_PATH=; rpm -ivh <rpm_file_name>'
   ```

   The `ENV-PATH` path can be found in the following file:

   ```
   $SYBASE/SYBASE.sh
   ```

   The **SYBASE.sh** file should have at least the following mandatory environment variables:

   - `SYBASE=<The SYBASE ASE Installation Path>`
   - `SYBASE_ASE=<The SYBASE ASE Version>`

- `SAP_JRE8_64=<The JRE directory path>`
- `SAP_JCONNECT=<The directory path for the JDBC driver>`

If you do not want to use the default values for the above parameters, then create a new file with these environment variables and point them to custom locations, and provide this file path using the following command:

```
export SAPASE_ENV_PATH=<ENV-PATH>
```

Uninstall the SAP Sybase ASE Connector

Before you uninstall the Cohesity SAP Sybase ASE connector, ensure that you set the following environment file:

```
export SAPASE_ENV_PATH=<ENV-PATH>
```

The `ENV-PATH` path can be found in the following file:

```
$SYBASE/SYBASE.sh
```

The **SYBASE.sh** file should have at least the following mandatory environment variables:

- `SYBASE=<The SYBASE ASE Installation Path>`
- `SYBASE_ASE=<The SYBASE ASE Version>`
- `JRE_PATH=<The JRE directory path>`
- `SAP_JCONNECT=<The directory path for the JDBC driver>`

If you do not want to use the default values for the above parameters, then create a new file with these environment variables and point them to custom locations, and provide this file path using the following command:

```
export SAPASE_ENV_PATH=<ENV-PATH>
```

To uninstall the Cohesity SAP Sybase ASE connector, run the following command with appropriate parameters on SAP Sybase ASE host CLI:

Using `yum`:

```
sudo yum remove <package>
```

Using `rpm`:

```
sudo rpm -e cohesity-sapase-connector-<version>.x86_64.rpm
```

Upgrade the SAP Sybase ASE Connector

To upgrade the Cohesity SAP Sybase ASE connector, run the following command on SAP Sybase ASE host:

```
sudo sh -c 'export SAPASE_ENV_PATH=<ENV-PATH>; rpm -U cohesity-sapase-con-
nector-<version>.rpm'
```

The `ENV-PATH` path can be found in the following file:

```
$SYBASE/SYBASE.sh
```

The **SYBASE.sh** file should have at least the following mandatory environment variables:

- `SYBASE=<The SYBASE ASE Installation Path>`
- `SYBASE_ASE=<The SYBASE ASE Version>`
- `JRE_PATH=<The JRE directory path>`
- `SAP_JCONNECT=<The directory path for the JDBC driver>`

If you do not want to use the default values for the above parameters, then create a new file with these environment variables and point them to custom locations, and provide this file path using the following command:

```
export SAPASE_ENV_PATH=<ENV-PATH>
```

Register and Manage the SAP Sybase ASE on Linux Source

To protect your SAP Sybase ASE databases, you need to register your SAP Sybase ASE deployment as a source on the Cohesity Platform. After registering your SAP Sybase ASE deployment as a source, you can:

- Update the SAP Sybase ASE source configuration.
- Refresh the source details.
- Unregister the SAP Sybase ASE source from the Cohesity Platform.

Register SAP Sybase ASE as a Source

Before you register SAP Sybase ASE as a source, ensure that you:

- Review the considerations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP Sybase ASE on Linux Protection.

To register SAP Sybase ASE as a source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
    a. From the **Source Type** drop-down menu, select **SAP Sybase ASE**.
    b. From the **Host OS Type** drop-down menu, select **Linux**.
    c. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the node(s) you have identified to run the SAP Sybase ASE Connector. For more

information, see Plan and Prepare for SAP Sybase ASE on Linux Protection.

> **Note:** Cohesity supports backup and restore of SAP Sybase ASE databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

d. In the **Datasource Agent Installation Path** field, enter the directory path on the SAP Sybase ASE node, where you have installed the SAP Sybase ASE Connector. For more information, see Install the SAP Sybase ASE Connector.

e. In the **Authentication Settings**, configure any one of the following authentication methods based on your requirements:

Authentication Method	Fields	Description
Password	Select if you want to use password-based authentication.	
	Username	The SAP ASE backup or restore user name.
	Password	The SAP ASE backup or restore user password.
Kerberos	Select if you want to use Kerberos-based authentication.	
	Kerberos Principal	The Kerberos principal that is required to connect to the database.
	Keytab File Path	The absolute path of the keytab file for the given Kerberos principal. **Note:** The path must be accessible by the Cohesity Agent user.
	Kerberos Credential Cache Path	The absolute path of the Kerberos ticket cache.

Authentication Method	Fields	Description
SecureStoreKey	Select if you want to use key-based authentication.	

	SecureStoreKey	The secure store key for key-based authentication.
		
		Note: For key-based authentication, ensure that Cohesity Agent and ASE server is running with the same user, aseuserstore.
		The following is the command to create the key:
		`aseuserstore set <KEY> <HOST>:<PORT> <USERNAME> <PASSWORD>`
		
		Sample key:
		`> aseuserstore list`
		`KEY: SYSKEY`
		`CON_ENV : COH6XTXXT2:5000`
		`USERNAME : sa`
		`PASSWORD : ***`

f. Optional. Enable **SAP ASE Client SSL Setting** and provide the absolute SSL certificate path on the SSL-enables SAP ASE cluster to enable a secure connection between the SAP ASE client and the SAP ASE Server.

> **Note:** Ensure that the Cohesity Agent user has access to the SSL Certificate path.

g. Enable **Cohesity SSL Settings** to enable a secure connection to the Cohesity cluster while reading or writing data. Enter the path to the SSL certificate used for secure data transfer between the Cohesity cluster and the SAP ASE cluster.

Note:

- It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
- Ensure that you generate and copy the certificate file from the Cohesity node to all the SAP ASE nodes. For more information, see Prerequisites.
- Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
- Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

h. Under **Source Settings**, enter the details of the SAP Sybase ASE source:

Field	Mandatory?	Description
Datasource Name	Yes	A unique name to identify your SAP Sybase ASE host.

SAP ASE Server Name	Yes	The SAP Sybase ASE server name. The server name can be found with the following SQL command. Connect to the database and execute the following commands:
		`>SELECT @@SERVERNAME`
		`>go`

SAP ASE Server Port	No	If you are not using the default port, specify the port on which the SAP Sybase ASE server is listening. The default port is 4091.
SAP ASE Server Hostname/IP	No	The hostname or IP address where the SAP ASE server is listening.

4. Click **Register**.

Update the SAP Sybase ASE Source Configuration

After registering your SAP Sybase ASE as a source on the Cohesity Platform, you can update the SAP Sybase ASE source configuration.

> **Important:** You cannot modify the source name.

To update the SAP Sybase ASE source configuration:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase ASE source and select **Edit**.
3. On the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the SAP Sybase ASE Source Configuration

You can refresh the SAP Sybase ASE Source and re-verify the source with the available source registration parameters.

To refresh the SAP Sybase ASE Source Configuration:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase ASE Source and select **Refresh**.

Unregister the SAP Sybase ASE Source

If you plan to stop protecting your SAP Sybase ASE databases, you can unregister the SAP Sybase ASE source from the Cohesity Platform.

To unregister the SAP Sybase ASE source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase ASE source and select **Unregister**.

Backup SAP Sybase ASE on Linux Database

Cohesity Platform provides a simple workflow to back up your SAP Sybase ASE databases. After registering the SAP Sybase ASE as a source on the Cohesity Platform, you can start configuring your backups.

Backup Process

Cohesity Platform provides a SAP Sybase ASE connector for backup and restore of SAP Sybase ASE deployments. The SAP Sybase ASE connector leverages the native backup capabilities provided by SAP Sybase ASE. Cohesity SAP Sybase ASE connector uses JDBC interface to connect the SAP Sybase ASE Server.

After registering the SAP Sybase ASE source on the Cohesity Platform, you can create a Protection Group and specify the objects such as databases you want to backup.

Protect SAP Sybase ASE on Linux Deployments

In Cohesity Platform, SAP Sybase ASE deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects such as databases that you want to backup.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select the standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP Sybase ASE deployments:

1. In the Cohesity Dashboard, do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the SAP Sybase ASE source and then select **Protect**.
2. From the **Source** drop-down list, select the registered SAP Sybase ASE source or click **Register Source** to register a SAP Sybase ASE source. For more information, see Register SAP Sybase ASE as a Source.
3. In the **Objects** section, add the object that you plan to protect. The object should be a database name. To back up a database, enter the name of the database. For example, `H00`. Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the **Backup Settings** section, enable the following configuration based on your requirements:

Options	Mandatory?	Description
Convert Incremental Backup To Full Backup On failure	No	Enable if you want to convert the incremental backup run to full backup in case of an error.

Options	Mandatory?	Description
Enable Incremental dumps for databases	No	Enable incremental/differential backups on the SAP Sybase ASE host. **Note:** After you have enabled incremental backups, ensure that you first initiate a full backup through Cohesity Dashboard.

7. Optional. Enable **Advanced Options** to configure advanced settings:

Field	Mandatory?	Default	Description
Parallel Objects	No	6	The number of databases for parallel data movement.
IO Thread Count	No	0	The number of threads that can be used for processing parallel I/O operations. The default value is calculated by the SAP ASE connector.
RPC Timeout In Seconds	No	120	The RPC request timeout between the SAP ASE source and the Cohesity cluster.
Hang Timeout in Seconds	No	1200	Define the duration in seconds, after which if there is no data movement, the protection run is marked as hung and killed by the SAP ASE Connector.
Max IO Write Size To Cohesity	No	4194304 bytes	The maximum size of the data packet that can be sent to the Cohesity cluster.
Environment Variables	No		The custom environment variables used by SAP ASE connector. **Expected value**: `KEY=VALUE` **Valid Environment Variable**: `ASE_BLOCK_SIZE`

8. In the **Settings** section, do the following:
    1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.

2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

   Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.

3. In the **Mounts** field, enter the number of VIPs you want to use while writing data to the Cohesity storage.

   > **Note:** Cohesity recommends that the number of mounts is equal to the number of nodes or Hybrid Extenders.

9. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click Edit on the right. For more information, see Additional Settings.
10. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Additional Settings	Description
Concurrency	Maximum number of backup streams that will be created to exchange data with the Cohesity cluster. By default, maximum scaling is 16. Available concurrency is distributed across multiple databases to be backed up or restored.
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the SAP Sybase ASE on Linux Backup Configuration

In the Cohesity Platform, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase ASE.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase ASE.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase ASE.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover SAP Sybase ASE on Linux Database

You can restore to the same SAP Sybase ASE deployment or to a different SAP Sybase ASE deployment. The Cohesity SAP Sybase ASE connector leverages the restore tool provided by SAP Sybase ASE. You can create a Recovery configuration and specify the object such as a database to recover.

> **Note:** Ensure that the database is not in use (active connections) during recovery.

Recover SAP Sybase ASE on Linux Databases

To restore SAP Sybase ASE databases:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery. You can also select a point-in-time snapshot or specific data backup that you want to recover.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same SAP Sybase ASE database on the same SAP Sybase ASE host.
New Object	To restore to a different SAP Sybase ASE host or to the same SAP Sybase host as different objects. Select the **Target** SAP Sybase ASE host to which the data needs to be restored. Also, enable **Overwrite existing object with the same name** and type **Yes**, if you want to overwrite the existing object with the same name.

8. In the **Restore Settings** section, enable **Create Database** if you want the restore job to create a database schema on the target cluster. For the database creation during restoration, the details provided in the following table are honored:

	Database already exists	Overwrite	Create Database	Action
1	False	-	-	Create
2	True	Yes	Enabled	Drop and create
3	True	Yes	Disabled	Load
4	True	no	-	Error

> **Note:** For option (1) and (2), schema from source database will be replayed as it is; Hence "disk init" will use the same location as source database to create a physical device

9. Optional. Enable **Advanced Options** to configure advanced settings:

Field	Mandatory?	Default	Description
Parallel Objects	No	6	The number of databases for parallel data movement.
IO Thread Count	No	0	The number of threads that can be used for processing parallel I/O operations. The default value is calculated by the SAP ASE connector.
RPC Timeout In Seconds	No	120	The RPC request timeout between the SAP ASE source and the Cohesity cluster.
Hang Timeout in Seconds	No	1200	Define the duration in seconds, after which if there is no data movement, the protection run is marked as hung and killed by the SAP ASE Connector.
Max IO Write Size To Cohesity	No	4194304 bytes	The maximum size of the data packet that can be sent to the Cohesity cluster.

Field	Mandatory?	Default	Description
Environment Variables	No		The custom environment variables used by SAP ASE connector. **Expected value**: `KEY=VALUE` **Valid Environment Variable**: `ASE_BLOCK_SIZE`

10. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	The maximum number of restore streams that is created to exchange data with the cluster.
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

11. Click **Recover**.

Restore SAP Sybase ASE System Databases

The Cohesity SAP Sybase ASE connector enables you to restore system databases such as `sybsystemdb` and `master` database.

Restore Sybsystemdb Database

Before you restore the `sybsystemdb` database, ensure that you perform the following on the SAP ASE host:

1. Execute the following SQL command:

```
sp_configure "enable xact coordination",0
```

2. Stop and start the SAP ASE service.

To restore the `sybsystemdb` database, follow the steps in the SAP Sybase ASE Databases restore procedure, and ensure that you disable the **Create Database** option in the **Restore Settings** section.

Restore Master Database

Before you restore the master database, ensure that you perform the following on the SAP ASE host:

1. Start Adaptive Server in single-user mode, for detailed instructions, see How to Start Adaptive Server in Single-User Mode in SAP documentation.
2. Restart the SAP ASE server.

To restore the master database, follow the steps in the SAP Sybase ASE Databases restore procedure, and ensure that you disable the **Create Database** option in the **Restore Settings** section.

SAP Sybase ASE on Microsoft Windows

Cohesity Platform natively integrates with SAP Sybase ASE to provide backup and recovery for SAP Sybase ASE deployments on Microsoft Windows. Using Cohesity Connector Agent, you can backup and restore the SAP Sybase ASE databases.

Cohesity Platform provides a SAP Sybase ASE connector that enables you to perform a full backup, differential backup, log backup, restore and point in time restore at the database level.

Core Components

The following are the core components for protecting SAP Sybase ASE databases:

- Cohesity SAP Sybase ASE connector
- Cohesity Connector Agent

As illustrated, you can integrate the Cohesity Platform with your SAP Sybase ASE deployments using the Cohesity SAP Sybase ASE connector and Cohesity Connector Agent.

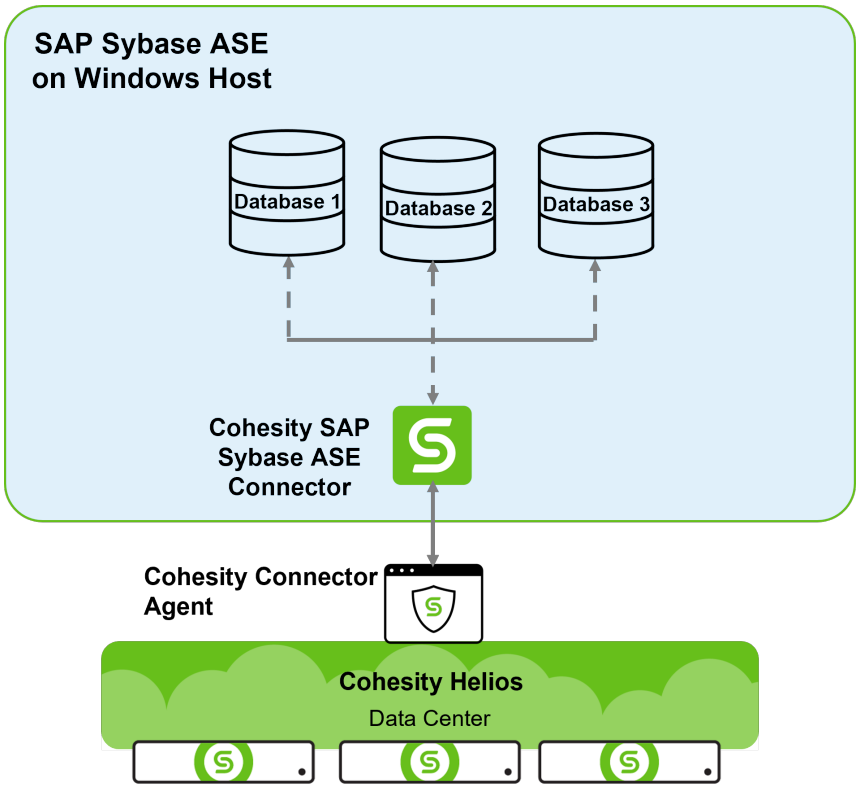

Plan and Prepare for SAP Sybase ASE on Windows Protection

Before you register your SAP Sybase ASE deployment as a source with Cohesity Platform, ensure the following prerequisites:

- The user account of the SAP Sybase ASE instance must have database admin privileges or operator privileges.
- Download and install the Cohesity Windows Agent on the SAP Sybase ASE host.
- Download and install the Cohesity SAP Sybase ASE Connector on the SAP Sybase ASE host using the same user that was used for installing Cohesity Windows Agent.
- Review and understand the limitations.

Supported Versions

For more information on supported SAP Sybase ASE versions, see SAP Sybase ASE versions.

Port Requirements

You must open certain ports in the firewall to allow the Cohesity cluster to transmit and receive data. Ensure that you open the following port in the firewall to allow data transfer between the SAP Sybase ASE host and the Cohesity cluster:

Port	Source	Target	Direction	Network Protocol
50051	SAP Sybase ASE	Cohesity cluster	Bidirectional	TCP/IP
11113	SAP Sybase ASE	Cohesity cluster	Unidirectional	TCP/IP

Considerations

Review and understand the following requirements before you protect your single-node SAP Sybase ASE deployment.

- Backup or restore for SAP Sybase ASE HA cluster deployments are not supported.
- Ensure that the database is not in use (active connections) during recovery.
- Once Cohesity protects the Sybase ASE database, it is not recommended to use the SAP backup and/or restore tools outside of Cohesity, including native dump or load tools provided by SAP Sybase ASE.
- The same database should not be part of multiple Protection Groups.
- Kerberos authentication is currently not supported.
- Backup or restore not supported if the SAP Sybase ASE server is listening on IPV6.
- If the hole is introduced in the backup timeline, then point-in-time recovery (PITR) is possible only after the subsequent successful full or incremental backup is completed.
- Cohesity supports backup and restore of SAP Sybase ASE databases running on dual-stack (IPv4 and IPv6) mode, single-stack (only IPv6) mode, and single-stack (only IPv4) mode.

Prerequisites

- The user account used to install the Cohesity Windows Agent and Cohesity ASE connector must have administrator privileges
- Local System Account must not be used for installing and running Cohesity Windows agent and connector.
- Enable incremental backups on the SAP Sybase ASE host using the following database command:

```
use <dbname> master ..sp_dboption <dbname>, "allow incremental dumps", false
```

> **Note:** After enabling incremental backups, you have to first initiate a full backup through the Cohesity Dashboard.

- Only IPv4 is supported and hence ensure that the SAP Sybase ASE server listens on IPv4 address.
- In a multitenant environment that has a hybrid extender-based configuration, if the data source does not use a secure gRPC client, then ensure that you execute the following commands on each hybrid extender before registering the SAP Sybase ASE source on the Cohesity cluster:

```
$bifrost.sh stop
$bifrost.sh start --magneto_grpc_timeout_msecs=300000 --tcp_proxy_
connector_ports=445:445,2049:2049,80:80,443:443,11117:11117,11113:11113
$sudo firewall-cmd --add-port=11113/tcp --zone=public --permanent &&
sudo firewall-cmd --reload
```

- Set the following variables in the system environment variables:

Environment Variable	Description
SYBASE	The path of the Sybase installation directory.
SYBASE_ASE	The subdirectory of the Adaptive Server component.
SAP_JRE8_64	The location of the Java Runtime Environment.
SAP_JCONNECT	The Sybase JDBC driver.

- To securely connect to the Cohesity cluster while reading or writing data, you can use secure gRPC.

  Note:
  - It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
  - Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
  - Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

  To use secure gRPC, perform the following:

  1. Generate the certificate config file on the Cohesity node using the following command:

```
~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_
name <clustername> -output_file /home/cohesity/<clustername>.cfg
```

2. Copy `/home/cohesity/<clustername>.cfg` from the Cohesity node to all SAP ASE nodes at the same location. For example, `/opt/certs/<clustername>.cfg`.

   > **Note:** Ensure that the administrator account user used to install the Cohesity Windows agent and SAP ASE Connector has read access to the generated certificate config file.

3. While registering the SAP Sybase ASE source, provide the path of the certificate config file in the **Cohesity SSL Settings** field.

- Increase the SAP Sybase ASE backup server memory to 1GB. For more information, see Utility Commands Reference in SAP Documentation.

Download and Install the Cohesity Windows Agent

The Cohesity Windows Agent is available as an executable installer. Download this installer and install it accordingly on the SAP Sybase ASE hosts to be protected. For detailed instructions, see Download and Install the Windows Agent.

> **Note:** Make sure that you do not select any add-ons.

Download and Install the SAP Sybase ASE Connector Installer

Cohesity Platform provides a SAP Sybase ASE Connector for backup and restore of enterprise SAP Sybase ASE deployments.

Download the SAP Sybase ASE Connector

Cohesity Platform provides a SAP Sybase ASE Connector to backup and restore SAP Sybase ASE databases. The SAP Sybase ASE Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for SAP Sybase ASE Connector, and copy the datasource agent to the SAP Sybase ASE host.

Install the Cohesity SAP Sybase ASE Connector

Before you install the Cohesity SAP Sybase ASE connector, ensure that you are logged as the same user with which you installed the Cohesity Windows agent.

To install the Cohesity SAP Sybase ASE connector:

1. Log on to the SAP Sybase ASE node where you have copied the Cohesity SAP Sybase ASE connector installer executable, `Cohesity_SAP_ASE_connector_<version>_Win_`

`x64_Installer.exe`.

2. Double-click the installer and follow the instructions to Install the Cohesity SAP Sybase ASE connector.

Upgrade the SAP Sybase ASE Connector

To upgrade the SAP Sybase ASE Connector:

1. Log on to the SAP Sybase ASE node where you have copied the Cohesity SAP Sybase ASE connector installer executable, `Cohesity_SAP_ASE_connector_<version>_Win_x64_Installer.exe`. Ensure that you are logged in as the same user with which you installed the Cohesity Windows agent.
2. Double-click on the Cohesity SAP Sybase ASE Connector installer, and when prompted, select **Upgrade** and follow the instructions.

Uninstall the SAP Sybase ASE Connector

To uninstall the Cohesity SAP Sybase ASE connector:

1. Log on to the SAP Sybase ASE host.
2. Select the **Start** menu and look for the SAP Sybase ASE Connector in the list shown.
3. Press and hold (or right-click) on the SAP Sybase ASE Connector, then select **Uninstall**.

Register and Manage the SAP Sybase ASE on Windows Source

To protect your SAP Sybase ASE databases, you need to register your SAP Sybase ASE deployment as a source on the Cohesity Platform. After registering your SAP Sybase ASE deployment as a source, you can:

- Update the SAP Sybase ASE source configuration.
- Refresh the source details.
- Unregister the SAP Sybase ASE source from the Cohesity Platform.

Register SAP Sybase ASE as a Source

Before you register SAP Sybase ASE as a source, ensure that you:

- Review the considerations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP Sybase ASE on Windows Protection.

To register SAP Sybase ASE as a source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
   a. From the **Source Type** drop-down menu, select **SAP ASE**.
   b. From the **Host OS Type** drop-down menu, select **Windows**.
   c. In the **Hostname/IP Addresses** field, enter the hostname or IP address of the node(s) you have identified to run the SAP Sybase ASE connector. For more information, see Plan and Prepare for SAP Sybase ASE on Windows Protection.

      > **Note:** Cohesity supports backup and restore of SAP Sybase ASE databases running on dual-stack (IPv4 and IPv6) mode, single-stack (only IPv6) mode, and single-stack (only IPv4) mode.

   d. In the **Datasource Agent Installation Path** field, enter the directory path on the SAP Sybase ASE node, where you have installed the SAP Sybase ASE connector. For more information, see Install the Cohesity SAP Sybase ASE Connector.
   e. In the **Authentication Settings**, configure any one of the following authentication methods based on your requirements:

Authentication Method	Fields	Description	
Password	Select if you want to use password-based authentication.		
	Username	The SAP ASE backup or restore user name.	
	Password	The SAP ASE backup or restore user password.	

Authentication Method	Fields	Description	
SecureStoreKey	Select if you want to use key-based authentication.		

	SecureStoreKey	The secure store key for key-based authentication.
		Note: Ensure that the key is generated using the same user with which Cohesity agent service is running. For more details, refer to SAP jConnect for JDBC Programmers Reference.
		The following is the sample command to create a key:
		`aseuserstore set <KEY> <HOST>:<PORT> <USERNAME> <PASSWORD>`
		Sample key:
		`> aseuserstore list`
		`KEY: SYSKEY`
		`CON_ENV : COH6XTXXT2:5000`
		`USERNAME : sa`
		`PASSWORD : ***`	

f. Optional. Enable **SAP ASE Client SSL Setting** and provide the absolute SSL certificate path on the SSL-enables SAP ASE cluster to enable a secure connection between the SAP ASE client and the SAP ASE Server.

> **Note:** Ensure that the Cohesity Agent user has access to the SSL Certificate path.

g. Enable **Cohesity SSL Settings** to enable a secure connection to the Cohesity cluster while reading or writing data. Enter the path to the SSL certificate used for secure data transfer between the Cohesity cluster and the SAP ASE cluster.

Note:

- It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
- Ensure that you generate and copy the certificate file from the Cohesity node to all the SAP ASE nodes. For more information, see Prerequisites.
- Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
- Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

h. Under **Source Settings**, enter the details of the SAP Sybase ASE source:

Field	Mandatory?	Description
Datasource Name	Yes	A unique name to identify your SAP Sybase ASE host.

SAP ASE Server Name	Yes	The SAP Sybase ASE server name. The server name can be found with the following SQL command. Connect to the database and execute the following commands:
		`>SELECT @@SERVERNAME`
		`>go`

SAP ASE Server Hostname/IP	No	The hostname or IP address where the SAP ASE server is listening.
SAP ASE Server Port	No	If you are not using the default port, specify the port on which the SAP Sybase ASE server is listening. The default port is 4091.

i. Click **Register**.

Update the SAP Sybase ASE Source Configuration

After registering your SAP Sybase ASE as a source on the Cohesity Platform, you can update the SAP Sybase ASE source configuration.

> **Important:** You cannot modify the source name.

To update the SAP Sybase ASE source configuration:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase ASE source and select **Edit**.
3. On the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Refresh the SAP Sybase ASE Source Configuration

You can refresh the SAP Sybase ASE Source and re-verify the source with the available source registration parameters.

To refresh the SAP Sybase ASE Source Configuration:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase ASE Source and select **Refresh**.

Unregister the SAP Sybase ASE Source

If you plan to stop taking the backup of your SAP Sybase ASE databases, you can unregister the SAP Sybase ASE source from the Cohesity Platform.

To unregister the SAP Sybase ASE source:

1. In the Cohesity Dashboard, select **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP Sybase ASE source and select **Unregister**.

Backup SAP Sybase ASE on Windows Database

Cohesity Platform provides a simple workflow to back up your SAP Sybase ASE databases. After registering the SAP Sybase ASE as a source on the Cohesity Platform, you can start configuring your backups.

Backup Process

Cohesity Platform provides a SAP Sybase ASE connector for backup and restore of SAP Sybase ASE deployments. The SAP Sybase ASE connector leverages the native backup capabilities provided by SAP Sybase ASE. Cohesity SAP Sybase ASE connector uses the JDBC interface to connect the SAP Sybase ASE Server. Ensure that the JDBC driver is installed for the SAP Sybase ASE database.

After registering the SAP Sybase ASE source on the Cohesity Platform, you can create a Protection Group and specify the objects such as databases you want to backup.

Protect SAP Sybase ASE on Windows Deployments

In Cohesity Platform, SAP Sybase ASE deployments are protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects such as databases that you want to backup.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP Sybase ASE deployments:

1. In the Cohesity Dashboard, do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Universal Data Adapter**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the SAP Sybase ASE source and then select **Protect**.
2. From the **Source** drop-down list, select the registered SAP Sybase ASE source or click **Register Source** to register a SAP Sybase ASE source. For more information, see Register SAP Sybase ASE as a Source.
3. In the **Objects** section, add the object that you plan to protect. The object should be a database name. To back up a database, enter the name of the database. For example, `H00`. Click the plus icon (+) to enter multiple objects that you plan to back up. Then, click **Save Selection**.
4. In the **Protection Group** section, do one of the following:
   - Select a **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy.
6. In the **Backup Settings** section, enable the following configuration based on your

requirements:

Options	Mandatory?	Description
Convert Incremental Backup To Full Backup On failure	No	Enable if you want to convert the incremental backup run to full backup in case of an error.
Enable Incremental dumps for databases	No	Enable incremental/differential backups on the SAP Sybase ASE host. **Note:** After you have enabled incremental b ackups, ensure that you first initiate a full ba ckup through Cohesity Dashboard.

7. Optional. Enable **Advanced Options** to configure advanced settings:

Field	Mandatory?	Default	Description
Parallel Objects	No	6	The number of databases for parallel data movement.
IO Thread Count	No	0	The number of threads that can be used for processing parallel I/O operations. The default value is calculated by the SAP ASE connector.
RPC Timeout In Seconds	No	120	The RPC request timeout between the SAP ASE source and the Cohesity cluster.
Hang Timeout in Seconds	No	1200	Define the duration in seconds, after which if there is no data movement, the protection run is marked as hung and killed by the SAP ASE Connector.
Max IO Write Size To Cohesity	No	4194304 bytes	The maximum size of the data packet that can be sent to the Cohesity cluster.
Environment Variables	No		The custom environment variables used by SAP ASE connector. **Expected value**: `KEY=VALUE` **Valid Environment Variable**: `ASE_BLOCK_SIZE`

8. In the **Settings** section, do the following:
    1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
    2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

       Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
    3. In the **Mounts** field, enter the number of VIPs you want to use while writing data to the Cohesity storage.

       > **Note:** Cohesity recommends that the number of mounts is equal to the number of nodes or Hybrid Extenders.
9. After you have completed the **Settings**, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
10. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the Cohesity cluster. By default, maximum scaling is 16. Available concurrency is distributed across multiple databases to be backed up or restored.
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs

Additional Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the SAP Sybase ASE on Windows Backup Configuration

In the Cohesity Platform, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase ASE.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
6. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to

execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase ASE.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SAP Sybase ASE.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Protection Group Only
   - Delete Protection Group and Snapshots

Recover SAP Sybase ASE on Windows Database

You can restore to the same SAP Sybase ASE deployment or to a different SAP Sybase ASE deployment. The Cohesity SAP Sybase ASE connector leverages the restore tool provided by SAP Sybase ASE. You can create a Recovery configuration and specify the object such as a database to recover.

Recover Sybase ASE on Windows Databases

To restore SAP Sybase ASE databases:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character ***.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery. You can also select a point-in-time snapshot or specific data backup that you want to recover.
6. Click **Next: Recover Options**.
7. In the **Recover As** section, based on your requirements, select any one of the following options:

Options	Description
Replace Original	To restore to the same SAP Sybase ASE database on the same SAP Sybase ASE host.

Options	Description
New Object	To restore to a different SAP Sybase ASE host or to the same SAP Sybase host as different objects. Select the **Target** SAP Sybase ASE host to which the data needs to be restored

8. In the **Restore Settings** section, enable **Create Database** if you want the restore job to create a database schema on the target cluster. For the database creation during restoration, the details provided in the following table are honored:

	Database already exists	Overwrite	Create Database	Action
1	False	-	-	Create
2	True	Yes	Enabled	Drop and create
3	True	Yes	Disabled	Load
4	True	no	-	Error

> **Note:** For options (1) and (2), schema from the source database will be replayed as it is; Hence "disk init" will use the same location as the source database to create a physical device

9. Optional. Enable **Advanced Options** to configure advanced settings:

Field	Mandatory?	Default	Description
Parallel Objects	No	6	The number of databases for parallel data movement.
IO Thread Count	No	0	The number of threads that can be used for processing parallel I/O operations. The default value is calculated by the SAP ASE connector.
RPC Timeout In Seconds	No	120	The RPC request timeout between the SAP ASE source and the Cohesity cluster.

Field	Mandatory?	Default	Description
Hang Timeout in Seconds	No	1200	Define the duration in seconds, after which if there is no data movement, the protection run is marked as hung and killed by the SAP ASE Connector.
Max IO Write Size To Cohesity	No	4194304 bytes	The maximum size of the data packet that can be sent to the Cohesity cluster.
Environment Variables	No		The custom environment variables used by SAP ASE connector. **Expected value**: `KEY=VALUE` **Valid Environment Variable**: `ASE_BLOCK_SIZE`

10. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	The maximum number of restore streams that is created to exchange data with the cluster.
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

11. Click **Recover**.

Restore SAP Sybase ASE System Databases

The Cohesity SAP Sybase ASE connector enables you to restore system databases such as `sybsystemdb` and `master` database.

Restore Sybsystemdb Database

Before you restore the `sybsystemdb` database, ensure that you perform the following on the SAP ASE host:

1. Execute the following SQL command:

```
sp_configure "enable xact coordination",0
```

2. Stop and start the SAP ASE service.

To restore the `sybsystemdb` database, follow the steps in the SAP Sybase ASE Databases restore procedure, and ensure that you disable the **Create Database** option in the **Restore Settings** section.

Restore Master Database

Before you restore the master database, ensure that you perform the following on the SAP ASE host:

1. Start the Adaptive Server in single-user mode, for detailed instructions, see How to Start Adaptive Server in Single-User Mode in SAP documentation.
2. Restart the SAP ASE server.

To restore the master database, follow the steps in the SAP Sybase ASE Databases restore procedure, and ensure that you disable the **Create Database** option in the **Restore Settings** section.

Cancel a Recovery Run

When you cancel a recovery run, the SAP Sybase ASE connector cancels the job by terminating the recovery processes.

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click the action icon (...) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
    - **Continue Recovery** if you want to continue the recovery task.
    - **Cancel Recovery** if you want to cancel the recovery task.

Troubleshooting Sybase ASE on Windows Protection

The following table lists the logs that you can use to troubleshoot failures during the data protection workflows:

Component	Path to the log file
Cohesity SAP Sybase ASE Connector logs	`C:\ProgramData\Cohesity\uda`
Cohesity Windows Agent service logs	`C:\ProgramData\Cohesity`

> **Note:** The `C:\ProgramData` folder is a hidden folder in Windows.

Known Issues and Workaround

The following table lists the known issues and workarounds:

Known Issue	Workaround
After the Cohesity Windows Agent is restarted, in certain cases a few stale DBs may continue to run during a recovery job run.	Manually drop the database, recreate and retry the restore.
PITR restore fails with the following error: `error "[13]:Cannot find data backup run for PITR 1623883236"`	If the point in time selected for restore is before the first full backup, then you may see this error. Please select a point later than the end of the first full backup.
In case of SAP Sybase ASE version mismatch , database fails to come online after a successful restore.	Try to online the database manually, `online database <dbname>`

SAP on Oracle for Linux

Cohesity natively integrates with SAP on Oracle to provide backup and recovery for SAP on Oracle deployments. Cohesity provides an Oracle Connector that enables you to do the following:

- Backs up data files, control files, and online redo log files of the database.
- Backs up offline redo log files.
- Restores data files, control files, and online and offline redo log files.
- Recovers database files and restores profiles and log files.
- Manages the database instance, space, and segments (that is, tables and indexes).
- Performs database administration tasks such as statistics update, check database system, adapt next extents, clean up logs and DBA tables.
- Functions as a help tool to monitor the database during a backup.
- Functions as an internal help tool started by BRBACKUP, BRARCHIVE, and BRRESTORE.

Core Components

The following are the core components for protecting SAP on Oracle databases:

- Oracle Database
- SAP on Oracle Connector

- Cohesity Connector Agent
- Cohesity View

As illustrated, you can integrate the Cohesity with your SAP on Oracle deployments using the SAP on Oracle Connector and Cohesity Connector Agent.

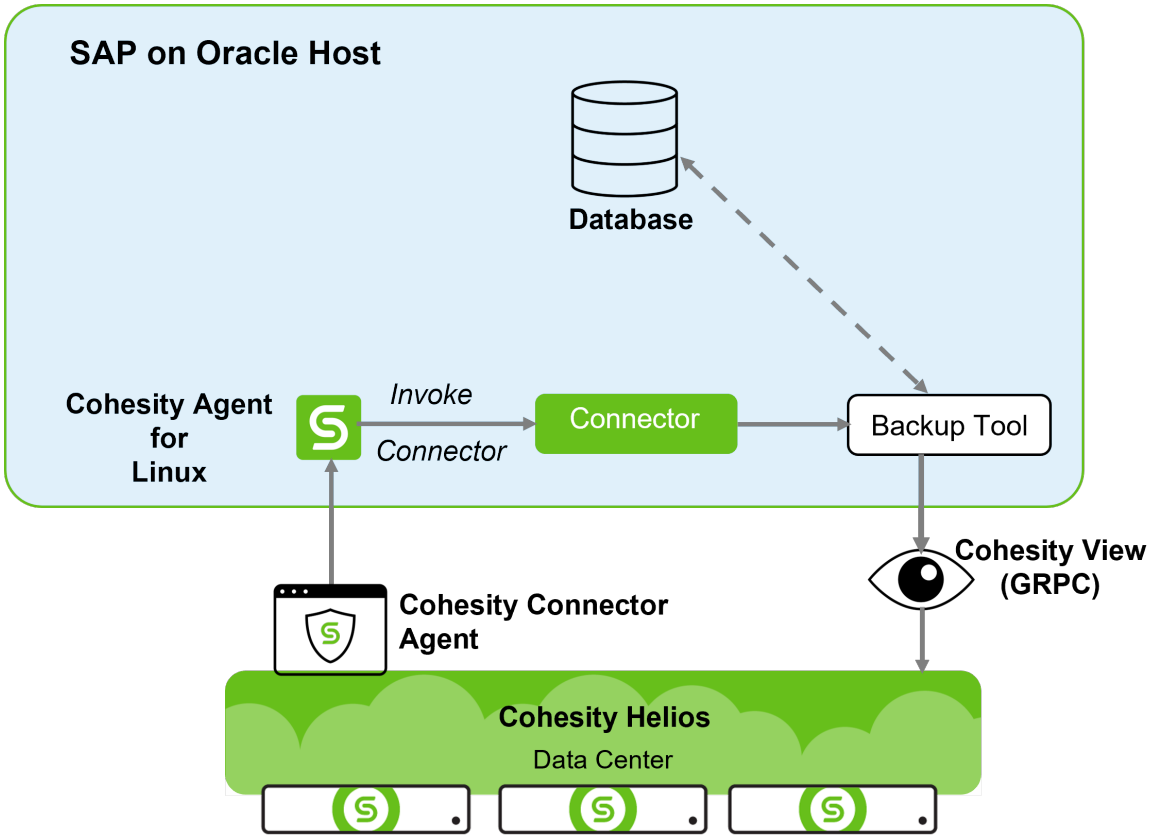

Plan and Prepare for SAP on Oracle for Linux Protection

Before you register your SAP on Oracle deployment as a source with Cohesity and protect SAP on Oracle databases, ensure you meet the prerequisites in the install section and understand the limitations.

Supported Versions

For more information on supported SAP on Oracle versions, see Supported Software versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	SAP on Oracle Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP on Oracle Local Host	SAP on Oracle Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
SAP on Oracle Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
SAP on Oracle Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Considerations

Review and understand the following before you protect your Oracle database:

- SAP on Oracle installations performed using the oracle user are not supported because they require a BRTools configuration as specified in SAP Note 1598594. However, Cohesity only supports the BRTools configuration outlined in SAP Note 113747.
- Automatic database creation is not supported. You need to create a database before running either backup or restore.
- Alternate restore to a new database on a different SAP on Oracle system with a different SID option is not supported.
- SAP on Oracle Connector agent does not support auto-discovery and entity hierarchy; hence you cannot browse or search and select objects for backup.
- Offline backup of the database is not supported.
- Indexing is not supported.
- You cannot protect multiple hosts using a single Protection Group.

- You cannot backup multiple databases in a single Protection Group.
- You cannot restore multiple objects, search, and recover an object.
- The RPM-based installer is not currently supported.
- SAP Oracle ASM deployment is not supported.
- Ensure that the SAP on Oracle DB instance uses an SPFILE to start up as recommended by SAP. For more information, see https://me.sap.com/notes/1598594.
- It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.

Download and install the Cohesity Linux Agent

The Cohesity Linux Agent is available as script-based and RPM-based installers, based on your requirements download the installer and install it accordingly on the SAP on Oracle database. For detailed instructions, see Download and Install the Linux Agent.

> **Note:** Ensure you meet the following prerequisites below before you install the Cohesity Linux agent.

Download and Install the SAP on Oracle Connector

Cohesity provides a SAP on Oracle Connector to backup and restore SAP on Oracle databases. The SAP on Oracle Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for SAP on Oracle Connector, and copy the datasource agent to the SAP on Oracle host.

Install the SAP on Oracle Connector

On the SAP on Oracle host where you have copied the SAP on Oracle Connector (`scp cohesity_secure_connector_service_6.6.0b_ent_sap_oracle_java_installer`), perform the following to install the SAP on Oracle Connector:

Before You Install

- Before you install the SAP on Oracle Connector, ensure the following prerequisites:

  Ensure the variables $SAPDATA_HOME and $ORACLE_HOME are set for the user performing the installation.

  > **Note:** Cohesity recommends using the "root" user account to install the SAP on Oracle Connector.

  Example:

```
$ export ORACLE_HOME=</oracle/SID/121>
$ export SAPDATA_HOME=</oracle/SID>
```

- During the SAP on Oracle Connector installation, you need permission to write to /usr/sap/<SID>/ directory as it is the default installation location. If you do not have write access, please specify the custom installation directory using "-I" or "--install-dir" while running the installer.

  Example:

  ```
  ./cohesity_secure_connector_service_sap_oracle_java_installer -- --
  install -s SID -w uda -S ora<sid> -I /home/ora<sid>/uda_install/
  ```

- The BRTOOLS are pre-installed on the SAP on Oracle System. The Cohesity SAP on Oracle connector integrates with BRTOOLS for backup and restore. SAP recommends BRTOOLS version 7.40 Patch 46 or greater. For more information, see https://launchpad.support.sap.com/#/notes/3060252.
- Ensure that the BRTOOLS' ownership and permissions follow the SAP note: https://me.sap.com/notes/113747.

  For example:

  ```
  -rwsrwxr-- 1 orasid sapsys 10022600 Aug 23 2012 brarchive
  -rwsrwxr-- 1 orasid sapsys 10251536 Aug 23 2012 brbackup
  -rwsrwxr-- 1 orasid sapsys 12179560 Aug 23 2012 brconnect
  -rwxr-xr-x 1 sidadm sapsys 10708840 Aug 23 2012 brrecover
  -rwxr-xr-x 1 sidadm sapsys 4140576 Aug 23 2012 brrestore
  -rwxr-xr-x 1 sidadm sapsys 12778384 Aug 23 2012 brspace
  -rwxr-xr-x 1 sidadm sapsys 4711664 Aug 23 2012 brtools
  ```

- Use "bash" shell to install the SAP on Oracle Connector and to run backups and restore workflows of SAP on Oracle databases.
- Ensure that the NTP time is in sync between the client and the Cohesity cluster. Perform the following on the client:
  - Verify if NTP is enabled and synchronized in the client using the following command:

    ```
    timedatectl
    ```

  - If the NTP time is not synchronized, use the following command to synchronize the NTP time

    ```
    timedatectl set-ntp true
    ```

To install the SAP on Oracle Connector:

1. Log on to the SAP on Oracle host as a root user where you have copied the SAP on Oracle script-based installer, `cohesity_secure_connector_service_6.6.0b_ent_`

sap_oracle_java_installer.

> **Note:** During the SAP on Oracle connector installation, by default, the script installer extracts the files to the /tmp directory. If you plan to extract the files to a different directory, you can define the directory path in the TMPDIR environment variable as follows: export TMPDIR=<directory_path>.

2. Run the following command to install the SAP on Oracle Connector (cohesity_secure_connector_service_6.6.0b_ent_sap_oracle_java_installer):

```
./cohesity_secure_connector_service_6.6.0b_ent_sap_oracle_java_
installer -- --install -I <install_directory> -s <SID> -w uda
```

Where:

1. **-i <install_directory>** - The directory on the SAP on Oracle host on which you want to install the SAP on Oracle Connector.

   By default, the SAP on Oracle Connector is installed in the /usr/sap/<SID>/cohesity directory.

   You can use -i <install_directory> if you want to change the default installation location from /usr/sap/SID/cohesity/.

2. **-w <workflow>** - The type of workflow you intend to install the SAP on Oracle Connector. For UDA workflow, specify **uda**.

3. **-s <SID>** - SAP System Identification (SID) is a unique three-character code that is available for every R/3 installation (SAP system) that consists of a database server and several application servers.

   If you have different permissions (for the owner user and group) in the above-mentioned files, it is necessary to run backup and restore operations with the sudo user. Cohesity recommends adding the following in the sudoers file:

   1. Add the following in the sudoers file:

      ```
      ora<sid> ALL=(ALL) NOPASSWD:SETENV: /bin/chown, /bin/java,
      /bin/chmod
      ```

   2. Edit following line in ${INSTALL_DIR}/backint.sh to:

      ```
      JAVA="sudo -E java"
      ```

Refer to SAP recommendations for the default owner and permissions of init<DBSID>.ora, init<DBSID>.sap, $SAPDATA_HOME: https://launchpad.support.sap.com/#/notes/113747.

You can also run the following command to know more details about the options used in the above command.

```
$ ./cohesity_secure_connector_service_6.6.0b_ent_sap_oracle_java_
installer -- --help
```

3. After installation, if the user does not have access to `$SAPEXE (/usr/sap/<SID>/SYS/exe/run)`, create a soft link to backint.sh in `$SAPEXE` directory using the following command.

   Example:

   ```
   ln -sf '/home/ora<sid>/uda_install/backint.sh' '${SAPEXE}/backint'
   chown -h "ora<sid>":"sapsys" "${SAPEXE_DIR}/backint"
   ```

4. Optional. Create a softlink to uda_scripts manually, if the installer (`cohesity_secure_connector_service_6.6.0b_ent_sap_oracle_java_installer`) fails to create the soft link using the following command:

   ```
   chmod -R 755 /opt/cohesity ln -sf /home/ora<sid>/uda_install/uda_
   scripts /opt/cohesity/agent/uda_scripts
   ```

The following is a sample output of the successful installation:

```
------------------------------------------------------------------------
--
               --- S Y S T E M     I N F O R M A T I O N ---

***** DATE AND TIME INFORMATION *****
Fri May  7 04:42:03 EDT 2021

***** HOSTNAME INFORMATION *****
   Static hostname: saiakhil_sapnetweaver
         Icon name: computer-vm
           Chassis: vm
        Machine ID: 466bc8302ebc4873a6c8541d0854e0fd
           Boot ID: 747d1d4256424a0186df466e2ecc81b5
    Virtualization: vmware
  Operating System: Oracle Linux Server 7.7
       CPE OS Name: cpe:/o:oracle:linux:7:7:server
            Kernel: Linux 4.1.12-124.34.1.el7uek.x86_64
      Architecture: x86-64

***** ORACLE DATABASE INFORMATION *****
Warning: empty oracle database version received from lsinventory.

***** FILE SYSTEM DISK SPACE USAGE *****
Filesystem           Size  Used Avail Use% Mounted on
devtmpfs             7.7G     0  7.7G   0% /dev
tmpfs                7.8G     0  7.8G   0% /dev/shm
tmpfs                7.8G  759M  7.0G  10% /run
```

```
tmpfs                 7.8G     0  7.8G   0% /sys/fs/cgroup
/dev/sda1              97G   90G  7.0G  93% /
/dev/mapper/ol-u01     70G   11G   60G  16% /u01
/dev/mapper/ol-home    51G   11G   40G  22% /home
tmpfs                 1.6G   36K  1.6G   1% /run/user/54321
tmpfs                 1.6G     0  1.6G   0% /run/user/0

***** FREE AND USED MEMORY *****
              total        used        free      shared  buff/cache   avail-
able
Mem:       16150324     2782008      175700     4831032    13192616
9518368
Swap:      17825788      108720    17717068

***** SYSTEM UPTIME AND LOAD *****
 04:42:04 up 86 days,  3:52,  4 users,  load average: 0.19, 0.10, 0.06

***** CURRENTLY LOGGED-IN USERS *****
oracle   :0           2021-02-10 00:22 (:0)
oracle   pts/0        2021-02-10 00:22 (:0)
oracle   pts/1        2021-02-10 00:23 (:0)
root     pts/2        2021-05-07 04:27 (10.16.254.126)

***** TOP 5 MEMORY-CONSUMING PROCESSES *****
%MEM %CPU COMMAND
 5.3  0.0 ora_dbw1_sal
 5.2  0.0 ora_dbw0_sal
 3.8  0.0 ora_cjq0_sal
 1.8  0.0 ora_mmon_sal
 1.8  0.0 ora_smon_sal
--------------------------------------------------------------------------------
--
==> Installing Cohesity java backint SAP on Oracle Connector
--------------------------------------------------------------------------------
--
Are you sure the provided user 'saladm' is correct? Enter y[es]/n[o]:
Continuing installation...
--------------------------------------------------------------------------------
--
Found java executable in path

openjdk version "1.8.0_232"
OpenJDK Runtime Environment (build 1.8.0_232-b09)
OpenJDK 64-Bit Server VM (build 25.232-b09, mixed mode)
```

```
Java Runtime Environment check ................................ [passed]
Environment variables check ................................... [passed]
SAP profile files check ....................................... [passed]
Service user OS roles privileges check ........................ [passed]
------------------------------------------------------------------------------
--
Architecture Name     :
Operating System Name :
Service User          : saladm
Install Directory     : /usr/sap/SID/cohesity
Install Log Directory : /tmp/cohesity_oracle_java_backint_installer_
1620376923_99985.log
SAP Database SID      : SID
Cluster IPs           :
View Name             :
Sudo Command Prefix   :
------------------------------------------------------------------------------
--
==> Creating the required directories.
Creating install directory '/usr/sap/SID/cohesity' if not present.
Creating scripts directory '/usr/sap/SID/cohesity/scripts' if not present.
Creating lib directory '/usr/sap/SID/cohesity/lib' if not present.
Creating profile directory '/usr/sap/SID/cohesity/profile' if not present.
------------------------------------------------------------------------------
--
==> Copying the required files.
Copying 'backint.sh' to install directory '/usr/sap/SID/cohesity/'
Copying 'backint.jar' to install directory '/usr/sap/SID/cohesity/'
Copying 'backup_job.sh' to scripts directory '/us-
r/sap/SID/cohesity/scripts/'
Copying 'purge_job.sh' to scripts directory '/usr/sap/SID/cohesity/scripts/'
Copying 'libsbt_6_and_7_linux-x86_64.so' to lib directory '/us-
r/sap/SID/cohesity/lib/'
Copying 'initSID.utl' to profile directory '/usr/sap/SID/cohesity/profile/'
------------------------------------------------------------------------------
--
==> Customizing the required files.
Creating symbolic link from '/usr/sap/SID/cohesity/backint.sh' to '/us-
r/sap/SID/SYS/exe/run/backint'
Renaming '/usr/sap/SID/cohesity/profile/initSID.utl' to
'/usr/sap/SID/cohesity/profile/initSID.utl'
Creating symbolic link from '/oracle/SID/sapbackup' to '/or-
acle/SID/121/sapbackup'
Creating symbolic link from '/usr/sap/SID/cohesity/profile/initSID.utl' to
```

```
'/oracle/SID/sapprof/initSID.utl'
------------------------------------------------------------------------------
--
99408 99406
Creating uda_scripts directory '/usr/sap/SID/cohesity/uda_scripts' if not
present.
Copying scripts to uda_scripts directory /usr/sap/SID/cohesity/uda_scripts
Creating soft link to /opt/cohesity/agent/uda_scripts
Please provide the Control node(leave empty for default: 10.2.204.19):
Please provide the Source name(leave empty for default: saiakhil_
sapnetweaver): Please provide the Source UUID(leave empty for default:
b71f7f1d56fab8b60fe136a06f86a8e5): ----------------------------------------
-----------------------------------
==> Finding group of service user: saladm
==> Group name from /etc/group : sapsys
==> Group of service user: sapsys
------------------------------------------------------------------------------
--
==> Cohesity java backint SAP on Oracle Connector is installed successfully.
------------------------------------------------------------------------------
--
```

After the SAP on Oracle Connector is installed, the agent is available at:

```
/usr/sap/<SID>/cohesity
```

Uninstall the SAP on Oracle Connector

To uninstall the SAP on Oracle Connector:

1. Log on to the SAP on Oracle host where you have installed the SAP on Oracle Connector.
2. Run the following command to uninstall the SAP on Oracle Connector:

   ```
   ./cohesity_secure_connector_service_6.6.0b_ent_sap_oracle_java_
   installer
   -- -u -I <install_directory> -s <sid> -w uda
   ```

The following is a sample output of the successful uninstallation:

```
------------------------------------------------------------------------------
--
               --- S Y S T E M     I N F O R M A T I O N ---

***** DATE AND TIME INFORMATION *****
Fri May  7 04:37:54 EDT 2021

***** HOSTNAME INFORMATION *****
```

```
   Static hostname: saiakhil_sapnetweaver
         Icon name: computer-vm
           Chassis: vm
        Machine ID: 466bc8302ebc4873a6c8541d0854e0fd
           Boot ID: 747d1d4256424a0186df466e2ecc81b5
    Virtualization: vmware
  Operating System: Oracle Linux Server 7.7
       CPE OS Name: cpe:/o:oracle:linux:7:7:server
            Kernel: Linux 4.1.12-124.34.1.el7uek.x86_64
      Architecture: x86-64

***** ORACLE DATABASE INFORMATION *****
./cohesity_oracle_java_backint_install.sh: line 274: /OPatch/opatch: No such
file or directory
Warning: empty oracle database version received from lsinventory.

***** FILE SYSTEM DISK SPACE USAGE *****
Filesystem           Size  Used Avail Use% Mounted on
devtmpfs             7.7G     0  7.7G   0% /dev
tmpfs                7.8G     0  7.8G   0% /dev/shm
tmpfs                7.8G  759M  7.0G  10% /run
tmpfs                7.8G     0  7.8G   0% /sys/fs/cgroup
/dev/sda1             97G   90G  7.3G  93% /
/dev/mapper/ol-u01    70G   11G   60G  16% /u01
/dev/mapper/ol-home   51G   11G   40G  22% /home
tmpfs                1.6G   36K  1.6G   1% /run/user/54321
tmpfs                1.6G     0  1.6G   0% /run/user/0

***** FREE AND USED MEMORY *****
              total        used        free      shared  buff/cache   avail-
able
Mem:       16150324     2286248      273972     4834132    13590104
9524980
Swap:      17825788      108728    17717060

***** SYSTEM UPTIME AND LOAD *****
 04:37:54 up 86 days,  3:48,  4 users,  load average: 0.05, 0.07, 0.05

***** CURRENTLY LOGGED-IN USERS *****
oracle   :0           2021-02-10 00:22 (:0)
oracle   pts/0        2021-02-10 00:22 (:0)
oracle   pts/1        2021-02-10 00:23 (:0)
root     pts/2        2021-05-07 04:27 (10.16.254.126)
```

```
***** TOP 5 MEMORY-CONSUMING PROCESSES *****
%MEM %CPU COMMAND
 5.3  0.0 ora_dbw1_sal
 5.2  0.0 ora_dbw0_sal
 3.8  0.0 ora_cjq0_sal
 1.8  0.0 ora_mmon_sal
 1.8  0.0 ora_smon_sal
------------------------------------------------------------------------
--
==> Uninstalling Cohesity java backint SAP on Oracle Connector
------------------------------------------------------------------------
--
Are you sure you want to uninstall? Enter y[es]/n[o]: Continuing unin-
stallation...
------------------------------------------------------------------------
--
Removed symbolic link for '/usr/sap/SID/SYS/exe/run/backint'
Removed symbolic link for '/oracle/SID/sapprof/initSID.utl'
Preserved initSID.sap into '/oracle/SID/sapprof/initSID.sap'
Removed install directory '/usr/sap/SID/cohesity' and its contents.
------------------------------------------------------------------------
--
==> Cohesity java backint SAP on Oracle Connector is uninstalled suc-
cessfully.
------------------------------------------------------------------------
--
```

Upgrade the SAP on Oracle Connector

Before upgrading the connector, ensure the following:

- Version 6.x uses Bash scripts and hence does not require Java.
- Version 7.x onwards uses Java and hence requires Java runtime.

To upgrade the SAP on Oracle Connector, run the following command on the SAP on Oracle hosts:

```
./cohesity_secure_connector_service_6.6.0b_ent_sap_oracle_java_installer
-- -n -I <install_directory> -s <sid> -w uda
```

After you have upgraded the SAP on Oracle Secure Connector, log on to the SAP on Oracle host as root user and create a soft link to uda_scripts using the following command:

```
chmod -R 755 /opt/cohesity ln -sf
/home/ora<sid>/uda_install/uda_scripts
/opt/cohesity/agent/uda_scripts
```

Register and Manage the SAP on Oracle for Linux Source

To protect your SAP on Oracle databases, you need to register your SAP on Oracle deployment as a source on Cohesity. After registering your SAP on Oracle deployment as a source, you can:

- Update the username and password of the database user.
- Refresh the source details.
- Unregister the SAP on Oracle source from Cohesity.

Register SAP on Oracle as a Source

Before you register SAP on Oracle as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP on Oracle for Linux Protection.

To register SAP on Oracle as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
    1. From the **Source Type** drop-down menu, select **SAP Oracle**.
    2. From the **Host OS Type** drop-down menu, select **Linux**.
    3. In the **Hostname / IP Addresses** field, enter the hostname or IP address of the SAP on the Oracle host you have identified to run the secure connector service for SAP Oracle. For more information, see Plan and Prepare for SAP on Oracle for Linux Protection.
    4. In the **Datasource Agent Installation Path**, enter the directory path on the SAP on Oracle host, where you have installed the secure connector service for SAP Oracle. For more information, see Download and Install the SAP on Oracle Connector.
    5. In the **APP Authentication** section, provide the SAP on Oracle database's Username and Password.

        > **Note:** The Password field is optional.

    6. Enable **Custom Options**, and perform the following:
        1. Ignore the **Mounts View** option. It is not applicable for SAP on Oracle source registration.

2. Ignore the **Source Registration Options**. It is not applicable to SAP on Oracle registration.

7. Click **Register**.

Update the SAP on Oracle Source Configuration

After registering your SAP on Oracle as a source on Cohesity, you can:

- Update the username and password of the database user.
- Update the Script Directory Path only if you have uninstalled and reinstalled the secure connector service for SAP Oracle in a different directory on the SAP on Oracle host.

To update the SAP on Oracle source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP on Oracle source and select **Edit**.
3. On the **Update Universal Data** Adapter page, update the required configurations.
4. Click **Update**.

Refresh the SAP on Oracle Source Configuration

You can refresh the SAP on Oracle Source and re-verify the source with the available source registration parameters.

To refresh the SAP on Oracle Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP on Oracle Source and select **Refresh**.

Unregister the SAP on Oracle Source

If you plan to stop taking the backup of your SAP on Oracle databases, you can unregister the SAP on Oracle source from the Cohesity Cluster.

To unregister the SAP on Oracle source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP on Oracle source and select **Unregister**.

Backup SAP on Oracle for Linux Database

Cohesity provides a simple workflow to back up your SAP on Oracle databases. After registering the SAP on Oracle as a source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides a SAP on Oracle Connector for backup and restore of enterprise SAP on Oracle deployments. The SAP on Oracle Connector leverages the backup capabilities provided by SAP.

After registering the SAP on Oracle source on Cohesity, you can create a Protection Group and specify the objects such as databases and tables that you want to backup. In the Protection Group, you need to define the parameters that enable you to provide additional parameters required by the backup statements provided by SAP on Oracle.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your SAP on Oracle backup configuration:

Backup Parameters	Mandatory?	Description
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during backup. For example: `--use-secure-grpc=true` > **Note:** It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during backups. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`
`--sbt-certificate-file=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during backups. For example: `--sbt-certificate-file=/tmp/cert.cfg` Use the following command to generate the certificate: `~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name <clustername> -output_file /home/cohesity/<clustername>.cfg`

Protect SAP on Oracle for Linux Deployments

SAP on Oracle deployments is protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects such as databases and tables that you want to back up.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP on Oracle deployments:

1. Do one of the following:
    1. Select **Data Protection** > **Protection**. Click Protect on the top-right of the page and then select **Universal Data Adapter**.
    2. Select **Data Protection** > **Source**. Click the actions menu (⋮) next to the SAP on Oracle source and then select **Protect**.
2. From the **Source** drop-down list, select the registered SAP on Oracle source or click **Register Source** to register an SAP on Oracle source. For more information, see Register SAP on Oracle as a Source.
3. In the **Objects** section, add the object that you plan to protect. The object can be a database or a table. To back up a database, enter the name of the database.
4. In the **Protection Group** section, do one of the following:
    1. Select a **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
    2. Select **Existing Group** to use the settings of an existing Protection Group. From the Group list, select an existing Protection Group and click Protect.
5. In the **Settings** section, do the following:
    1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
    2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

       Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.

3. In the **Custom Options** section, you need not add any backup parameters. If you want to use secure gRPC communication during backup, enter the respective Backup Parameters.

   For advanced log backup use cases, you can provide the BRARCHIVE command options in the **Log Backup** field. By default, the BRARCHIVE command option for -o parameter is `-cds | -copy_delete_save`.

   ```
   -o <brarchive_command_option>
   ```

4. In the **Mounts** field. Select the number of mounts for the SAP on the Oracle host.

   > **Note:** Cohesity recommends that the number of mounts is equal to the number of nodes or Hybrid Extenders.

5. After you have completed the **Settings** configuration, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click **Edit** on the right. For more information, see Additional Settings below.

6. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable Pause Future Run only if you do not want to schedule future runs of the Protection Group.
End Time	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts before this date runs until completion even if it completes after this date.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	The maximum number of restore streams that are created to exchange data with the cluster. By default, maximum scaling is set to 16. > **Note:** Cohesity recommends that concurrency should be set to the number of nodes on the Cohesity cluster.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	Optional. Enter a brief description of the Protection Group.

Manage the SAP on Oracle for Linux Backup Configuration

In the Cohesity Platform, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for SAP on Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection

Group.

5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for SAP on Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups, local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.

- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details will not be displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for SAP on Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    1. Delete Protection Group Only
    2. Delete Protection Group and Snapshots

Recover SAP on Oracle for Linux Database

You can restore the SAP on Oracle host or object(s) such as databases in SAP on Oracle deployment. You can restore to the same SAP on the Oracle host or a different host. The SAP on Oracle Connector leverages the restore capabilities provided by SAP.

You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

Prerequisites

Ensure that:

- The Oracle Database state should be `nomount` state when performing alternate or cross restores.
- The Protection Group associated with the object that you plan to restore is paused or in not running state.
- If you perform an alternate restore, ensure the SAP on Oracle versions on the source and the target host are the same.

Recovery Parameters

In the Recovery configuration, you can define the following script parameters for your SAP on Oracle recovery:

Recovery Parameters	Mandatory?	Description
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during restore. For example: `--use-secure-grpc=true`
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during restore. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`
`--sbt-certificate-file=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during restore. For example: `--sbt-certificate-file=/tmp/cert.cfg`

Recover SAP on Oracle for Linux Databases

To restore SAP on Oracle deployments:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery.
6. Accept the latest Recovery Point or click the change icon and select a different recovery point.
    1. **Snapshots**: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
    2. **Point in Time**: Use the slide control to adjust the date and time and click Save.
7. Click **Next: Recover Options**.
8. In the **Recover As** section, select one of the following.

- **Replace Original**: To restore to the same SAP on Oracle host.
- **New Object**: To restore to a different SAP on Oracle host, do the following:
  - Select the **Target** SAP on the Oracle host to which the data needs to be restored.

    Note: When restoring to a different SAP system (different SID), ensure that the SID is explicitly updated using the **Rename** option. Otherwise, the restore may be treated as a same-SID restore and fail.
  - Enable **Overwrite existing object with the same name** toggle option to overwrite the existing data.

9. In the **Recovery Options** section, based on your requirements, select any one of the following options:
   - **Mounts**: Select the number of mounts for the SAP on the Oracle host.
   - **Concurrency**: The maximum number of restore streams that are created to exchange data with the cluster. By default, maximum scaling is set to 1.

     > **Note:** Cohesity recommends that concurrency should be set to the number of nodes on the Cohesity cluster.

   - **Custom Options**: Specify the parameters based on your requirements. See Sample Recovery Use Cases for examples.

     > **Note:** if you want to use it for secure gRPC communication during restore, ensure that you use the appropriate recovery parameter.

   - **Rename**: Provide the new name for the database, if you want to restore the database with a new name.

     Important: If you are recovering to a different SID, you must explicitly specify the target SID name in the **Rename** field. This ensures that the recover operation is treated as a cross-SID recovery.
   - **Task Name**: Change the default name of the recovery task.
10. Click **Recover**.

    > **Note:** Post recovery, the timestamp displayed on the files are the same as the recovery job time.

Sample Recovery Use Cases

Use Case	Instructions with Sample Parameters
Overwrite Point In Time Restore (PITR) to the same SAP on Oracle database.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, does not specify any parameters in the **Custom Options** field.
Restore specific data backup from the snapshot to the same SAP on Oracle database.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, does not specify any parameters in the **Custom Options** field.
Point In Time Restore (PITR) to a different SAP on Oracle host with the same database name.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown. 2. Under the **Recovery Options** section, do not specify any parameters in the **Custom Options** field. 3. Specify the existing database name in the **Rename** field.
Restore specific data backup from the snapshot to the different SAP on Oracle host with the same database name.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown. 2. Under the **Recovery** Options section, do not specify any parameters in the **Custom Options** field. 3. Specify the existing database name in the **Rename** field.

Use Case	Instructions with Sample Parameters

Point In Time Restore (PITR) to a different SAP on Oracle host with different database name.	Perform the following:
	1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown.
	2. Under the **Recovery Options** section, do not specify any parameters in the **Custom Options** field.
	Note:
	If the source database does not have the `SAPDATA_HOME` with the format `/oracle/SID`, then provide the source database `SAPDATA_HOME` path using the following parameter in the **Custom Options** field.
	`-h | --source-sapdata-home`
	For example: If the source database `SAPDATA_HOME` path is `/usr/sap/oracle/SID` then you pass the following arugment in the **Custom Options** field:
	`-h /usr/sap/oracle/SID`
	3. Specify the existing database name in the **Rename** field.

Use Case	Instructions with Sample Parameters

Restore specific data backup from the snapshot to the different SAP on Oracle host with different database name.	Perform the following:
	1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown.
	2. Under the **Recovery** Options section, do not specify any parameters in the **Custom Options** field.
	Note:
	If the source database does not have the `SAPDATA_HOME` with the format `/oracle/SID`, then provide the source database `SAPDATA_HOME` path using the following parameter in the **Custom Options** field.
	`-h | --source-sapdata-home`
	For example: If the source database `SAPDATA_HOME` path is `/usr/sap/oracle/SID` then you pass the following arugment in the **Custom Options** field:
	`-h /usr/sap/oracle/SID`
	3. Specify the existing database name in the **Rename** field.

Use Case	Instructions with Sample Parameters
Restore the SAP on Oracle instance to its original host, but assign it a new SID. Ensure that the original database is not overwritten during this process.	For example, consider an instance named `XR1` on `dbhost1` that has been successfully backed up. If you want to restore it as a new instance on the same host, follow these steps before starting the restore process: 1. Ensure the new instance already exists on the original host or create it manually. 2. Since Cohesity protects SAP on Oracle at the instance level, install the SAP on Oracle connector for the new instance (SID). For details, see Install the SAP on Oracle Connector. 3. Register the new SAP on Oracle instance as a source. For more information, see Register SAP on Oracle as a Source. To restore the backup to the new SID on the original host: 1. Recover the `XR1` instance to a new object. • In the **Recover As** section, select **New Object** and choose the newly registered source from the **Target** dropdown. • Enable the **Overwrite existing object with the same name** option. 2. In the **Recovery Options** section, leave the **Custom Options** field empty. 3. Specify the new SID name in the **Rename** field.

Cancel a Recovery Run

When you cancel a recovery run, the SAP on Oracle Connector leverages the Cancel-Job capabilities provided by SAP on Oracle.

To cancel a recovery run:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click the action menu (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
    1. Click **Continue Recovery** if you want to continue a recovery run.
    2. Click **Cancel Recovery** if you want to cancel the recovery run.

SAP on Oracle for Microsoft Windows

Cohesity natively integrates with SAP on Oracle to provide backup and recovery for SAP on Oracle deployments running on the Microsoft Windows platform. Using Cohesity Connector Agent, you can backup and restore the SAP on Oracle databases. The Cohesity Connector Agent provides a simple workflow to back up and restore your SAP on Oracle deployment.

Cohesity provides an SAP on Oracle connector that enables you to perform a full backup, incremental (differential) backup, log backup, restore and point in time restore at the database level.

Core Components

The following are the core components for protecting SAP on Oracle databases:

- Oracle Database
- Cohesity SAP on Oracle Connector
- Cohesity Connector Agent
- Cohesity View

As illustrated, you can integrate Cohesity with your SAP on Oracle deployments using the SAP on Oracle connector and Cohesity Connector Agent.

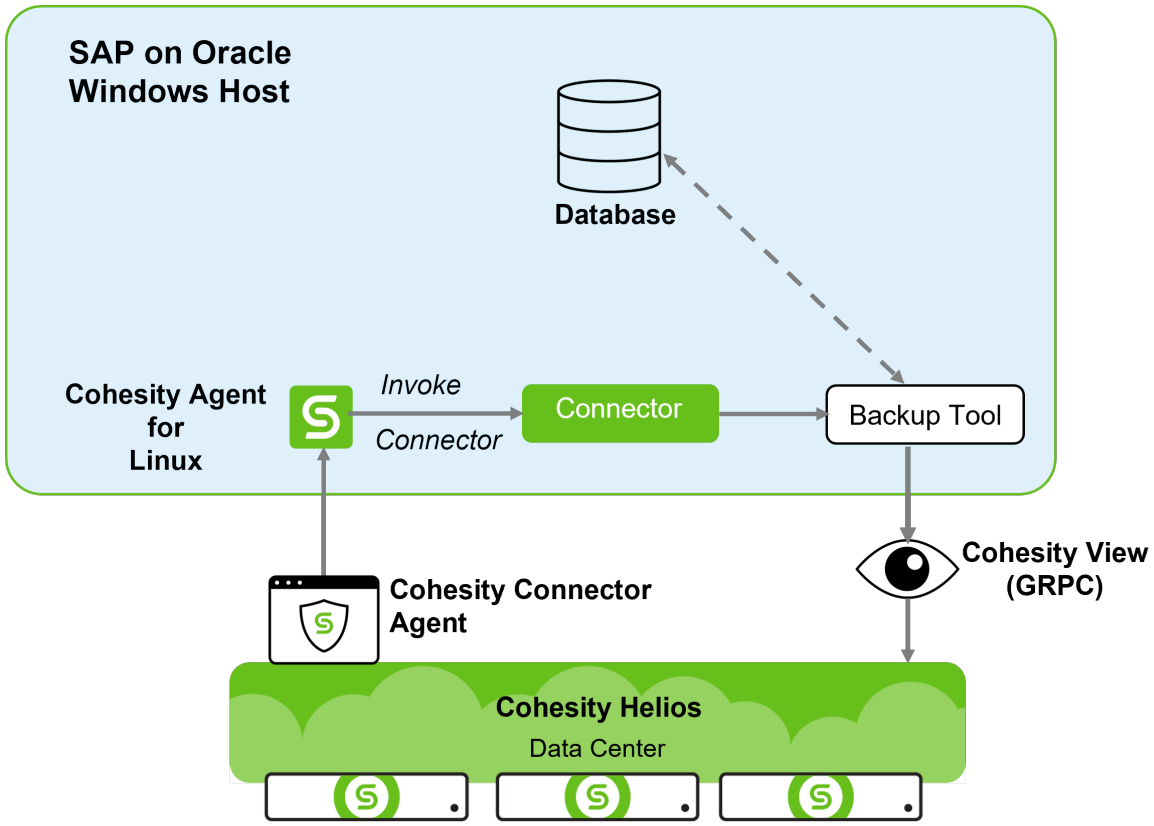

Plan and Prepare for SAP on Oracle for Windows Protection

Before you register your SAP on Oracle deployment as a source with Cohesity, ensure the following prerequisites:

- The user account of the SAP on Oracle instance must have database admin privileges.
- Download and install the Cohesity Windows Agent on the SAP on Oracle host.
- Download and install the Cohesity SAP on the Oracle Connector on the SAP on Oracle host using the same user account that you used to install the Cohesity Windows Agent.

Supported Versions

For more information on supported SAP on Oracle versions, see Supported Software versions.

Considerations

Review and understand the following before you protect your Oracle database:

- Automatic database creation is not supported. You need to create a database before running either backup or restore.
- In a protection run, you cannot back up multiple objects.
- In a recovery run, you cannot restore multiple objects.
- Cross-platform (Windows to Linux or vice versa) backup and restore are not supported.
- SAP on Oracle Connector agent does not support auto-discovery and entity hierarchy; hence you cannot browse or search and select objects for backup.
- Indexing is not supported.
- Offline backup of the database is not supported.
- You cannot protect multiple hosts using a single Protection Group.
- You cannot backup multiple databases in a single Protection Group.
- You cannot restore multiple objects, search and recover an object.
- You cannot cancel a running restore operation.
- SAP Oracle ASM deployment is not supported.
- Ensure that the SAP on Oracle DB instance uses an SPFILE to start up as recommended by SAP. For more information, see https://me.sap.com/notes/1598594.
- It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.

Allowlist the SAP on Oracle Windows Host IP address

To provide effective authorization and prevent rogue access to the Cohesity cluster, the SAP on Oracle Windows host IP address must be whitelisted in the Cohesity view configuration.

To allowlist the SAP on Oracle Windows host IP address:

1. Navigate to **File Services** > **Views**.
2. Click the **Global Settings** tab.
3. Under **Global Subnet Allowlist**, click the **+Add** button.
4. In the **Add Allowlist** dialog box, perform the following:
    1. In the **Server** field, enter the SAP on the Oracle host IP address.
    2. Select **Read/Write** for **SMB Permissions**.
    3. Select **None** for **NFS Squash**.
    4. In the **Description** field, enter a brief description.
    5. Click **Add**.

Download and Install the Cohesity Windows Agent

The Cohesity Windows Agent is available as an executable installer. Download this installer and install it accordingly on the SAP on Oracle host to be protected. For detailed instructions, see Download and Install the Windows Agent.

> **Note:** Ensure that you:
> - Install the Cohesity Windows Agent using the `<SID>adm` user account.
> - Select **Oracle SBT Library** as an add-on component in the Cohesity Windows Agent installer.

Download and Install the SAP on Oracle Connector

Cohesity provides a SAP on Oracle Connector to backup and restore SAP on Oracle databases. The SAP on Oracle Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for SAP on Oracle Connector, and copy the datasource agent to the SAP on Oracle host.

Install the Cohesity SAP on Oracle Connector

On the SAP on Oracle host where you have copied the SAP on Oracle Connector (`Cohesity_SAP_Oracle_connector_<version>_Installer.exe`), perform the following to install the SAP on Oracle Connector.

Prerequisites

Before you install the SAP on Oracle Connector, ensure that:

- The environment variables SAPDATA_HOME and ORACLE_HOME are set for the user performing the installation.
- You have Installed java runtime environment (JRE) version 1.8 (64-bit) or above on the SAP on Oracle host.
- The Cohesity cluster and the SAP on Oracle source are on the same NTP server.
- The BRTOOLS are pre-installed on the SAP on Oracle System. The Cohesity SAP on Oracle Connector integrates with BRTOOLS for backup and restore. SAP recommends BRTOOLS version 7.40 Patch 46 or greater. For more information, see more information, refer https://launchpad.support.sap.com/#/notes/3060252.
- Ensure that the NTP time is in sync between the client and the Cohesity cluster.

To install the SAP on Oracle Connector:

1. Log on to the SAP on the Oracle node where you have copied the Cohesity SAP on Oracle Connector installer executable, `Cohesity_SAP_Oracle_connector_<version>_Installer.exe`. Ensure that you are logged in as the same user with which you installed the Cohesity Windows agent.
2. Double-click on the Cohesity SAP on the Oracle Connector installer and follow the instructions.

Silently Install the SAP on Oracle Connector

After you have downloaded the SAP on Oracle Connector, you can perform a silent installation of the Cohesity SAP on Oracle Connector.

To perform the silent installation:

1. Log on to the SAP on the Oracle node where you have copied the Cohesity SAP on Oracle Connector installer executable, `Cohesity_SAP_Oracle_connector_<version>_Installer.exe`. Ensure that you are logged in as the same user with which you installed the Cohesity Windows agent.
2. Using the command-line interface, navigate to the directory where you have copied the Cohesity SAP on Oracle Connector installer executable, `Cohesity_SAP_Oracle_connector_<version>_Installer.exe`.
3. Use the following command for the silent installation:

   ```
   <InstallerExeFile> [/log="<absolute-path-to-log-file>"]
   [/dir="<absolute-path-to-install-location>"] [/silent, /verysilent]
   [/suppressmsgboxes] [/norestart]
   ```

   Where:

   - `/log="<absolute-path-to-log-file>"` - Records the events during installation and helps to debug the errors, if any, during the installation process.

- `/silent` - Suppresses the installation wizard and background window. Instead, it only displays the installation progress window.
- `/verysilent` - Suppresses the installation progress window.
- `/suppressmsgboxes` - Suppresses any message boxes if `/verysilent` is specified.
- `/norestart` - Prevents the setup from restarting the system following a successful installation, or after a Preparing to Install failure that requests a restart.
- `/dir="<absolute-path-to-install-location>"` - Specify the directory path where the SAP Oracle Connector for Windows should be installed.

  > **Note:** The installer creates the necessary folders if they do not exist. If this parameter is not specified, the SAP Oracle Connector is installed in the default directory path: `"%SystemDrive%\Program Files\Cohesity\SAP Oracle"`.

Sample:

```
Cohesity_SAP_Oracle_connector_<version>_Installer.exe
/log="C:\tmp\install_logs.txt" /verysilent /suppressmsgboxes /norestart
```

Upgrade the SAP on Oracle Connector

Before upgrading the connector, ensure the following:

- Version 6.x uses Bash scripts and hence does not require Java.
- Version 7.x onwards uses Java and hence requires Java runtime.

Before you upgrade the SAP on Oracle Connector, ensure that you pause the existing protection runs.

To upgrade the SAP on Oracle Connector:

1. Log on to the SAP on the Oracle node where you have copied the Cohesity SAP on Oracle Connector installer executable, `Cohesity_SAP_Oracle_connector_<version>_Installer.exe`. Ensure that you are logged in as the same user with which you installed the Cohesity Windows agent.
2. Double-click on the Cohesity SAP on the Oracle Connector installer, and when prompted, select **Upgrade** and follow the instructions.

Silently Upgrade the SAP on Oracle Connector

Before you upgrade the SAP on Oracle Connector, ensure that you pause the existing protection runs.

To silently upgrade the SAP on Oracle Connector, use the following command:

```
<InstallerExeFile> [/log="<absolute-path-to-log-file>"] [/dir="<absolute-path-to-install-location>"] [/silent, /verysilent] [/suppressmsgboxes] [/norestart]
```

1. Where:
   - `/log="<absolute-path-to-log-file>"` - Records the events during the upgrade and helps to debug the errors, if any, during the upgrade process.
   - `/silent` - Suppresses the upgrade wizard and background window. Instead, it only displays the upgrade progress window.
   - `/verysilent` - Suppresses the upgrade progress window.
   - `/suppressmsgboxes` - Suppresses any message boxes if `/verysilent` is specified.
   - `/norestart` - Prevents the setup from restarting the system following a successful upgrade or after a Preparing to upgrade failure that requests a restart.
   - `/dir="<absolute-path-to-install-location>"` - Specify the directory path where the SAP Oracle Connector is installed.

   Sample:

   ```
   Cohesity_SAP_Oracle_connector_<version>_Installer.exe /log="C:\tmp\upgrade_logs.txt" /verysilent /suppressmsgboxes /norestart
   ```

Uninstall the SAP on Oracle Connector

To uninstall the Cohesity SAP on Oracle Connector:

1. Log on to the SAP on the Oracle host.
2. Select the **Start** menu and look for the **SAP on Oracle Connector** in the list shown.
3. Press and hold (or right-click) on the SAP on Oracle Connector, then select **Uninstall**.

Silently Uninstall the SAP on Oracle Connector

To perform silent uninstallation:

1. Log on to the SAP on the Oracle host.
2. Using the command-line interface, navigate to the directory where you have installed the Cohesity SAP on Oracle Connector, For example, `C:\Program Files\Cohesity\SAP Oracle\Installer\`.

> **Note:** The details of the uninstaller executable are available in the registry value "HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\{CCEC1E9A-895C-4038-8CC9-31BBD31B9CA7}_is1\UninstallString". The default value is "C:\Program Files\Cohesity\SAP Oracle\Installer\unins000.exe", assuming C: is the default operating system installation volume and "C:\Program Files\Cohesity\SAP Oracle" is the default SAP Oracle Connector installation path.

3. Use the following command for the silent uninstallation:

```
<UninstallerExeFile>  [/log="<absolute-path-to-log-file>"] [/silent, /verysilent] [/suppressmsgboxes] [/norestart] [/preservesettings]
```

Where:

- /log="<absolute-path-to-log-file>" - Records the events during the uninstallation process and helps to debug the errors, if any, during the uninstallation process.
- /silent - Suppresses the uninstallation wizard and background window. Instead, it only displays the uninstallation progress window.
- /verysilent - Suppresses the uninstallation progress window.
- /suppressmsgboxes - Suppresses any message boxes if /verysilent is specified.
- /norestart - Prevents the setup from restarting the system following a successful uninstallation.
- [/preservesettings] - Preserve the registry settings when uninstalling the SAP Oracle Connector.

> **Note:** Uninstalling the SAP Oracle Connector clears registration settings from the registry. You can preserve the registration settings in the registry by using the /preservesettings flag when uninstalling from the command line. Preserving registry settings are not supported when uninstalling from the GUI.

Sample:

```
unins000.exe /log="C:\tmp\uninstall_logs.txt" /verysilent /suppressmsgboxes /norestart
```

Register and Manage the SAP on Oracle for Windows Source

To protect your SAP on Oracle databases, you need to register your SAP on Oracle deployment as a source on Cohesity. After registering your SAP on Oracle deployment as a source, you can:

- Update the username and password of the database user, if you connect to the database using password-based authentication.
- Refresh the source details.
- Unregister the SAP on Oracle source from Cohesity.

Register SAP on Oracle as a Source

Before you register SAP on Oracle as a source, ensure that you:

- Review the requirements and limitations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for SAP on Oracle for Windows Protection.

To register SAP on Oracle as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the **Register** on the top-right of the page and then select **Universal Data Adapter**.
3. In the **Register Universal Data Adapter** page, perform the following:
    1. From the **Source Type** drop-down menu, select **SAP Oracle**.
    2. From the **Host OS Type** drop-down menu, select **Windows**.
    3. In the **Hostname / IP Addresses** field, enter the hostname or IP address of the SAP on the Oracle host you have identified to run the SAP on Oracle Connector. For more information, see Plan and Prepare for SAP on Oracle for Windows Protection.
    4. In the **Datasource Agent Installation Path**, enter the directory path on the SAP on Oracle host, where you have installed the SAP on Oracle Connector. For more information, see Install the Cohesity SAP on Oracle Connector.
    5. In the **APP Authentication** section, provide the SAP on the Oracle database's **Username** and **Password**.

        > **Note:** The **Password** field is optional.

    6. Enable **Custom Options**, and perform the following:
        1. Ignore the **Mounts View** option. It is not applicable for SAP on Oracle source registration.

2. In the **Source Registration Options** field, enter the following optional source registration parameters in a space-separated format based on your requirement:

   ```
   --source-name=<unique_name_of_your_choice>
   ```

   If you do not define the source name using the parameter, the default hostname is used as the source name.

7. Click **Register**.

Update the SAP on Oracle Source Configuration

After registering your SAP on Oracle as a source on Cohesity, you can:

- Update the username and password of the database user.
- Update the **Datasource Agent Installation Path** only if you have uninstalled and reinstalled the SAP on Oracle Connector in a different directory on the SAP on Oracle host.

To update the SAP on Oracle source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP on Oracle source and select **Edit**.
3. On the **Update Universal Data** Adapter page, update the required configurations.
4. Click **Update**.

Refresh the SAP on Oracle Source Configuration

You can refresh the SAP on Oracle Source and re-verify the source with the available source registration parameters.

To refresh the SAP on Oracle Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP on Oracle Source and select **Refresh**.

Unregister the SAP on Oracle Source

If you plan to stop taking the backup of your SAP on Oracle databases, you can unregister the SAP on Oracle source from Cohesity.

To unregister the SAP on Oracle source:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the SAP on Oracle source and select **Unregister**.

Backup SAP on Oracle for Windows Database

Cohesity provides a simple workflow to back up your SAP on Oracle databases. After registering the SAP on Oracle as a source on Cohesity, you can start configuring your backups.

Backup Process

Cohesity provides a SAP on Oracle Connector for backup and restore of enterprise SAP on Oracle deployments. The SAP on Oracle Connector leverages the backup capabilities provided by SAP.

After registering the SAP on Oracle source on Cohesity, you can create a Protection Group and specify the objects such as databases and tables that you want to backup. In the Protection Group, you need to define the script parameters that enable you to provide additional parameters required by the backup statements provided by SAP on Oracle.

> **Note:** Cohesity recommends that you configure periodic full backups.

Backup Parameters

In the Protection Group, you can define the following parameters for your SAP on Oracle backup configuration:

Backup Parameters	Mandatory?	Description
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during backup. For example: `--use-secure-grpc=true` > **Note:** It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during backups. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`

Backup Parameters	Mandatory?	Description
`--sbt-certificate-file=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during backups. For example: `--sbt-certificate-file=/tmp/cert.cfg` Use the following command to generate the certificate: `~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name <clustername> -output_file /home/cohesity/<clustername>.cfg`

Troubleshooting Backup Error

Issue: Log backup fails with the following error:

```
BR0013W No offline redolog files found for processing
```

Problem: The issue occurs when the Oracle database is opened after recovery with `ALTER DATABASE OPEN RESETLOGS`, and the "`log sequence number`" is reset to `1`. However, since the BRARCHIVE log arch<SID>.log in the directory `$ORACLE_HOME/saparch` still contains entries for successful archive log backups with higher numbers. BRARCHIVE does not recognize the new logs as the logs to be saved.

Workaround: To mitigate this issue, refer to the following SAP recommendation:

- https://me.sap.com/notes/490976 (refer to step 4. Database was reset with RESET LOGS)
- https://me.sap.com/notes/10170

Protect SAP on Oracle SAP on Oracle for Windows Deployments

SAP on Oracle deployments is protected using a Protection Group and a Protection Policy. A Protection Group enables you to specify the objects such as databases and tables that you want to backup.

The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select which standard or custom policy to use when creating a Protection Group. For more information, see About Policies and Protection Groups.

To protect SAP on Oracle deployments:

1. Do one of the following:
    1. Select **Data Protection** > **Protection**. Click Protect on the top-right of the page and then select **Universal Data Adapter**.
    2. Select **Data Protection** > **Source**. Click the actions menu (⋮) next to the SAP on Oracle source and then select **Protect**.
2. From the **Source** drop-down list, select the registered SAP on Oracle source or click **Register Source** to register an SAP on Oracle source. For more information, see Register SAP on Oracle as a Source.
3. In the **Objects** section, enter the name of the database that you plan to protect.
4. In the **Protection Group** section, do one of the following:
    1. Select a **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name from the Protection Group.
    2. Select **Existing Group** to use the settings of an existing Protection Group. From the Group list, select an existing Protection Group and click Protect.
5. In the **Settings** section, do the following:
    1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
    2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

       Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
    3. In the **Custom Options** section, you need not add any backup parameters. By default, the BRARCHIVE command option for `-o` parameter is `-cds | -copy_delete_save`. That is, in the case of log backups, it does not delete offline redo log files from the hosts that have been saved once in the Cohesity cluster. You can override the default BRARCHIVE command option by providing the BRARCHIVE command options in the **Log Backup** field as follows:

       ```
       -o=<brarchive_command_option>
       ```

       > **Note:** If you want to use secure gRPC communication during backup, enter the respective backup parameter.
    4. In the **Mounts/VIPs** field. Select the number of mounts or VIPs for the SAP on Oracle host.

> **Note:** Cohesity recommends that the number of mounts is equal to the number of nodes or Hybrid Extenders.

5. After you have completed the **Settings** configuration, if you need to change any additional settings on the **New Universal Data Adapter Protection Group** page, scroll down and click **Edit** on the right. For more information, see Additional Settings below.
6. Click **Protect**.

Additional Settings

Additional Settings	Description
Pause Future Run	Enable Pause Future Run only if you do not want to schedule future runs of the Protection Group.
End Time	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts before this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Define the concurrency or number of channels to use for exchanging data with the cluster. For SAP on Oracle data protection on Windows, the concurrency depends on the SAP on Oracle license. • For Standard editions, only a single channel is supported. • For Enterprise editions, more than one channel is supported.

Additional Settings	Description
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	Optional. Enter a brief description of the Protection Group.

Manage the SAP on Oracle for Windows Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups associated with the Protection Policy. You

can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for SAP on Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for SAP on Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a protection run:

1. Navigate to **Data Protection** > **Protection**.
2. Click the action menu (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups, local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details will not be displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection** > **Protection**.
2. On the Protection page, select the Protection Group that you had created for SAP on Oracle.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    1. **Delete Protection Group Only**
    2. **Delete Protection Group and Snapshots**

Recover SAP on Oracle for Windows Database

You can restore the SAP on Oracle host or object(s) such as databases in SAP on Oracle deployment. You can restore to the same SAP on the Oracle host or a different host. The SAP Oracle Connector leverages the restore capabilities provided by SAP.

You can create a Recovery configuration and specify the objects such as databases backed up using the Protection Group.

Prerequisites

Ensure that:

- The Protection Group associated with the object that you plan to restore is paused or in not running state.
- If you perform an alternate restore, ensure the SAP on Oracle versions on the source and the target host are the same.

Recovery Parameters

In the Recovery configuration, you can define the following script parameters for your SAP on Oracle recovery:

Recovery Parameters	Mandatory?	Description
`--use-secure-grpc=<true/false>`	Optional	Set the parameter to true, if you want to use secure gRPC communication during restore. For example: `--use-secure-grpc=true`
`--sap-certificate-path=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during restore. For example: `--sap-certificate-path=/tmp/certs/sap_cert1837348343.cfg`
`--sbt-certificate-file=<cert_file_path>`	Optional	Specify the path of the self-signed certificate if you want to use it for secure gRPC communication during restore. For example: `--sbt-certificate-file=/tmp/cert.cfg`

Recover SAP on Oracle for Windows Database

To restore SAP on Oracle deployments:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Universal Data Adapter**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. In the **Object** field, enter the name of the database or table that you had backed up using the selected Protection Group. By default, the latest snapshot is selected for recovery.
6. Accept the latest Recovery Point or click the change icon and select a different recovery point.

1. **Snapshots**: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
2. **Point in Time**: Use the slide control to adjust the date and time and click Save.

7. Click **Next: Recover Options**.
8. In the **Recover As** section, select one of the following:

Options	Description
Replace Original	To restore to the same SAP on Oracle host.
New Object	To restore to a different SAP on Oracle host. Select the **Target** SAP on the Oracle host to which the data needs to be restored. Enable **Overwrite existing object with the same name** toggle option to overwrite the existing data.

9. In the **Recovery Options** section, based on your requirements, select any one of the following options:

Options	Description
Mounts	Select the number of mounts for the SAP on the Oracle host.
Concurrency	Define the concurrency or number of channels to use for exchanging data with the cluster. For SAP on Oracle data protection on Windows, the concurrency depends on the SAP on Oracle license. • For Standard editions, only a single channel is supported. • For Enterprise editions, more than one channel is supported.
Custom Options	Specify the recovery parameters based on your requirements. For examples, see Sample Recovery Use Cases. **Note:** if you want to use it for secure gRPC communication during restore, ensure that you use the appropriate recovery parameter.
Rename	Provide the new name for the object, if you want to restore the object with a new name.
Task Name	Change the default name of the recovery task.

10. Click **Recover**.

Sample Recovery Use Cases

Use Case	Instructions
Overwrite Point In Time Restore (PITR) to the same SAP on Oracle database.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, does not specify any parameters in the **Custom Options** field.
Restore specific data backup from the snapshot to the same SAP on Oracle database.	In the **Recover As** section, select **Replace Original**. Under the **Recovery Options** section, does not specify any parameters in the **Custom Options** field.
Point In Time Restore (PITR) to a different SAP on Oracle host with the same database name.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown. 2. Under the **Recovery Options** section, do not specify any parameters in the **Custom Options** field. 3. Specify the existing database name in the **Rename** field.
Restore specific data backup from the snapshot to the different SAP on Oracle host with the same database name.	Perform the following: 1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown. 2. Under the **Recovery** Options section, do not specify any parameters in the **Custom Options** field. 3. Specify the existing database name in the **Rename** field.

Use Case	Instructions

Point In Time Restore (PITR) to a different SAP on Oracle host with different database name.	Perform the following:
	1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown.
	2. Under the **Recovery Options** section, do not specify any parameters in the **Custom Options** field.
	> **Note:**
	If the source database does not have the `SAPDATA_HOME` with the format `/oracle/SID`, then provide the source database `SAPDATA_HOME` path using the following parameter in the **Custom Options** field.
	`-h | --source-sapdata-home`
	For example: If the source database `SAPDATA_HOME` path is `/usr/sap/oracle/SID` then you pass the following arugment in the **Custom Options** field:
	`-h /usr/sap/oracle/SID`
	3. Specify the existing database name in the **Rename** field.

Use Case	Instructions

Restore specific data backup from the snapshot to the different SAP on Oracle host with different database name.	Perform the following:
	1. In the **Recover As** section, select **New Object** and choose the target source name in the **Target** dropdown.
	2. Under the **Recovery** Options section, do not specify any parameters in the **Custom Options** field.
	> **Note:**
	If the source database does not have the `SAPDATA_HOME` with the format `/oracle/SID`, then provide the source database `SAPDATA_HOME` path using the following parameter in the **Custom Options** field.
	`-h | --source-sapdata-home`
	For example: If the source database `SAPDATA_HOME` path is `/usr/sap/oracle/SID` then you pass the following arugment in the **Custom Options** field:
	`-h /usr/sap/oracle/SID`
	3. Specify the existing database name in the **Rename** field.

Cancel a Recovery Run

When you cancel a recovery run, the SAP on Oracle Connector leverages the Cancel-Job capabilities provided by SAP on Oracle.

To cancel a recovery run:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click the action menu (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog box, select any one of the following:
    1. Click **Continue Recovery** if you want to continue a recovery run.
    2. Click **Cancel Recovery** if you want to cancel the recovery run.

PostgreSQL

Cohesity Platform natively integrates with PostgreSQL to provide backup and recovery for PostgreSQL deployments. Using Cohesity Connector Agent, you can backup and restore the PostgreSQL databases.

Cohesity Platform provides a PostgreSQL connector that enables you to perform a full backup, differential backup, log backup, restore and point in time restore at the cluster level.

Core Components

The following are the core components for protecting PostgreSQL databases:

- PostgreSQL Installation
- Cohesity PostgreSQL connector
- Cohesity Connector Agent

As illustrated, you can integrate the Cohesity Platform with your PostgreSQL deployments using the Cohesity PostgreSQL connector and Connector Agent.

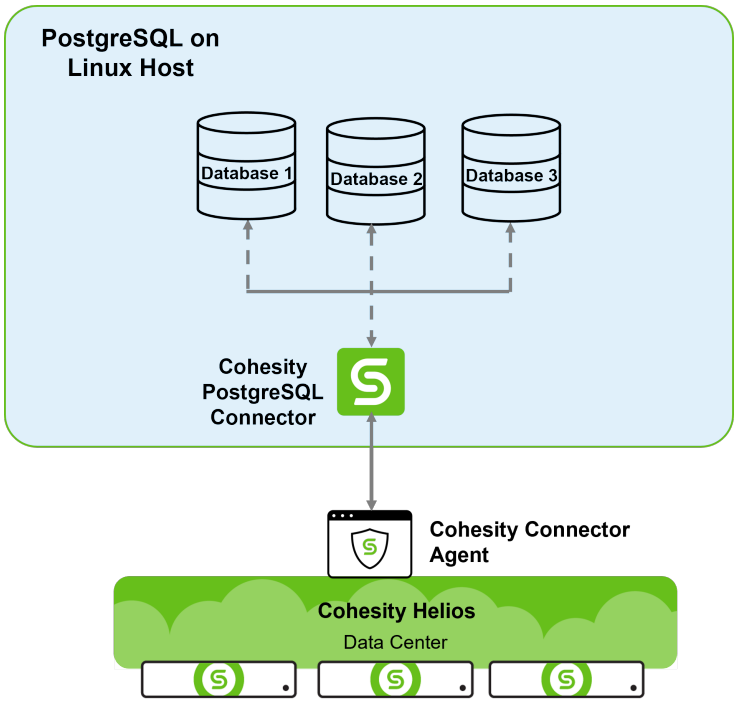

Plan and Prepare for PostgreSQL Protection

Before registering your PostgreSQL deployment as a source with Cohesity Platform, ensure the following prerequisites:

- The user account of the PostgreSQL instance must have database admin privileges.
- Cohesity supports Transparent Data Encryption (TDE).
- Download and install the Cohesity Linux Agent RPM on the PostgreSQL host.
- Download and install the Cohesity PostgreSQL Connector RPM on the PostgreSQL host

- Cohesity provides a script to migrate existing PostgreSQL source to the new workflow. After running the script, the instances appear under **Data Protection** > **Sources** > **PostgreSQL**. For more information, see Migration script for New User Experience.
- Review and understand the limitations.

Supported Versions

For more information on supported PostgreSQL versions, see PostgreSQL versions.

Port Requirements

Outgoing Traffic

Source	Destination	Port	Protocol	Type of Traffic
Cohesity cluster	PostgreSQL Node	50051	TCP/IP	Control path for all workflows

Bidirectional Traffic

Source	Destination	Port	Protocol	Type of Traffic
PostgreSQL Local Host	PostgreSQL Local Host	59999	TCP/IP	Required only if local-to-local communication (used for self-monitoring and debugging purposes) is blocked by the firewall.

Incoming Traffic

Source	Destination	Port	Protocol	Type of Traffic
PostgreSQL Node	Cohesity cluster	11113	TCP/IP	Backup and Recovery
PostgreSQL Node	Cohesity cluster	11117	TCP/IP	Backup and Recovery with secure gRPC

Prerequisites

Ensure that:

- You use a root user or non-root user to install the Cohesity Linux agent and Cohesity PostgreSQL Connector. If you choose to use a non-root user, you must add the `sudo` privileges for the non-root user using the syntax given below:

  ```
  <user_name> ALL=(ALL) NOPASSWD:SETENV: /path/to/command1,
  /path/to/command2, /path/to/command3
  ```

  The following is the `sudo` commands for PostgreSQL Connector:

```
cohesityagent ALL=NOPASSWD:SETENV: /bin/chmod, /bin/chown, /bin/mkdir,
/bin/rm, /bin/psql, /usr/bin/ps, /usr/sbin/runuser, /bin/java,
/usr/bin/netstat, /bin/test
```

> **Note:**
> - If the **cohesityagent** user entry is already present in **/etc/sudoers**, append the list of commands to the existing entry.
> - The actual paths for each command might vary in each system. For more information, see Download and Install the Linux Agent.

- Java version 1.8 or higher is installed on all PostgreSQL nodes and available in the path for the Cohesity agent, the root user, and the postgres user.
- Java binary is available in the `$PATH` environment variable.
- The default `postgres` database is present on the PostgreSQL instance.
- To securely connect to the Cohesity cluster while reading or writing data, you can use secure gRPC.

> **Note:**
> - It is **mandatory** to connect to the Cohesity cluster securely using secure gRPC when writing or reading data.
> - Using non-secure gRPC for data transfer between Cohesity and database hosts will result in unsecured communication.
> - Enabling secure gRPC will increase data transfer time between Cohesity and database hosts compared to non-secure communication.

To use secure gRPC, perform the following:

1. Generate the certificate config file on the Cohesity node using the following command:

   ```
   ~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_
   name <clustername> -output_file /home/cohesity/<clustername>.cfg
   ```

2. Copy `/home/cohesity/<clustername>.cfg` from the Cohesity node to all PostgreSQL nodes at the same location. For example, `/opt/certs/<clustername>.cfg`.

> **Note:** Ensure that the `root` or `non-root` user used to install the Cohesity Linux agent and PostgreSQL Connector has read access to the generated certificate config file.

3. While registering the PostgreSQL source, provide the path of the certificate config file in the **Cohesity SSL Settings** field.

- If you are restoring to the same or different cluster, ensure that the port you want to use for the restored instance is not in use.
- In case of Cohesity Cloud Edition deployment with any PostgreSQL setup, ensure that the Cohesity cluster name resolves to the IP addresses of all nodes through DNS.

Considerations

Review and understand the following limitations and known issues before you protect your PostgreSQL deployment:

- Only physical backup is supported.
- Backup or restore is performed at the cluster level. Object-level backup or restore is not supported.
- Restore across different PostgreSQL versions is not supported.
- For optimum performance in log backups, it is recommended to set the WAL size to 1GB.
- If the hole is introduced in the backup timeline, then point-in-time recovery (PITR) is possible only after the subsequent successful full or incremental backup is completed.
- Cohesity supports backup and restore of PostgreSQL databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.
- Cohesity supports granular-level database permissions which allows the non-superuser to register the PostgreSQL source and perform the backups. See Configure Privileges for Non-Admin Users.
- FIPS mode is not currently supported for PostgreSQL backups.

Considerations for HA Cluster Deployments

Review the following additional considerations when protecting HA cluster:

- Register the source by specifying all Patroni member node names or IP addresses in the hostname field during source registration.
- Backups run only on the current Patroni leader node. Standby and replica nodes are not backed up.
- Immediately trigger a full backup after a failover, switchover, or any event that changes the Patroni leader node.

- Restores are attempted only on the first Patroni member node defined in the source registration's hostname field. To change the restore target, update the source registration and reorder the node list.
- SSL keys must not be stored in the PostgreSQL data directory for Patroni deployments. Because Cohesity backs up only the active leader node, SSL keys stored in the data directory may not be available across all nodes after restore. Store SSL certificates in an external directory and specify the custom SSL certificate path during source registration.
- Restores are not Patroni-aware and always result in a standalone PostgreSQL instance. Additional manual steps are required to rebuild the Patroni cluster after restore. For more information, see Recover a Patroni-Managed PostgreSQL Cluster.
- Backups from a Patroni-managed (replicated) source can only be restored to a new location.
- Point-in-time recovery (PITR) from archived backups is not supported for Patroni deployments.
- Instant restore is not supported for Patroni-managed PostgreSQL deployments.
- In the case of the HA cluster, the backup is executed from the active node only.
- In case of a failover in the HA cluster, the next backup run will be a full backup.

Consideration for Transparent Data Encryption (TDE)

- Cohesity does not encrypt or decrypt PostgreSQL data or WAL files; encryption is handled by PostgreSQL.
- If you enable TDE, ensure that a master encryption key is created using the Cohesity REST API before registering the PostgreSQL source. The Cohesity user invoking the API must have appropriate Users and Groups to create and manage encryption keys.

  > **Note:** If your organization rotates the PostgreSQL master wrapping key (for example, every 30 or 90 days), ensure that the key required to decrypt the backup is available on the restore target. Cohesity recovers encrypted data files as-is. The recovery job can complete successfully, but PostgreSQL will not start if the required master wrapping key is missing or inaccessible.

- Cohesity does not create, manage, back up, or recover encryption keys (including the master wrapping key).
- After a recovery, PostgreSQL can start only if the required master wrapping key is available on the restore target. If the key is missing or inaccessible, the restore job may complete successfully, but PostgreSQL will not start.
- You cannot restore a non-TDE snapshot to a PostgreSQL instance that is already running with TDE enabled on that port. The recovery will fail.

> **Note:** Cohesity backs up and restores EDB PostgreSQL TDE data as-is, and PostgreSQL can start after restore only if the required unwrap configuration and master wrapping key are available on the restore target.

Download and Install the Cohesity Linux Agent

The Cohesity Linux Agent is available as RPM-based installers, based on your requirements download the installer and install it accordingly on the PostgreSQL hosts to be protected. For detailed instructions, see Download and Install the Linux Agent.

Download and install the PostgreSQL Connector

Cohesity Platform provides a PostgreSQL Connector for backup and restore of enterprise PostgreSQL deployments.

Download the PostgreSQL Connector

Cohesity Platform provides a PostgreSQL Connector to backup and restore PostgreSQL deployments. The PostgreSQL Connector is available as a datasource agent on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective datasource agent for PostgreSQL Connector, and copy the datasource agent to the PostgreSQL database host.

Install the PostgreSQL Connector

On the PostgreSQL database host where you have copied the PostgreSQL Connector, perform the following to install the PostgreSQL Connector.

To install the Cohesity PostgreSQL connector:

1. Log on to the PostgreSQL Linux machine using the user with sudo privileges.
2. Navigate to the directory where you have copied the RPM-based installer.
3. Install the Cohesity PostgreSQL connector using the following command:

   Using `yum`:

   ```
   # sudo yum -y install cohesity-postgres-connector-<version>.rpm
   ```

   Using `rpm -ivh`:

   ```
   # sudo rpm -ivh cohesity-postgres-connector-<version>.rpm
   ```

Uninstall the PostgreSQL Connector

To uninstall the Cohesity PostgreSQL connector, run the following command on PostgreSQL host:

Using `yum`:

```
# sudo yum -y remove cohesity-postgres-connector-<version>
```

Using `rpm`:

```
# sudo rpm -e cohesity-postgres-connector-<version>.rpm
```

Upgrade the PostgreSQL Connector

To upgrade the Cohesity PostgreSQL connector, run the following command on PostgreSQL host:

```
sudo rpm -U cohesity-postgres-connector-<version>.rpm
```

Migration script for New User Experience

Cohesity provides a new improved user experience that streamlines the process of registration, protection, and recovery of PostgreSQL database. To migrate the existing PostgreSQL database to the new workflow, Cohesity provides a migration script. However, the source and protection jobs for the existing PostgreSQL will continue to function as they are.

After the migration script is runned successfully, Cohesity displays the PostgreSQL instances listed separately under the **DataProtection** > **Sources** > **PostgreSQL** section.

Ensure that:

- Migration script file should have execute permission on the PostgreSQL server.
- PostgreSQL server can reach the Cohesity cluster on the required network ports listed in this document.
- Cohesity admin or equivalent user should perform the task.

To migrate the PostgreSQL database to new workflow:

1. Log on to the PostgreSQL Linux machine.
2. Navigate to `/opt/cohesity/postgreSQL/scripts`.
3. Run the following script files using `chmod +x` command:
    a. `migrate_scource.sh` - to migrate the source configurations.
    b. `migrate_protection.sh` - to migrate the protection jobs.
    For example: `chmod +x migrate_source.sh`
    To view the help options, run the command as `chmod +x migrate_source.sh --help`

Note:

- Original backup jobs are paused and not deleted.
- No roll back script provided. If the migration script fails, check the logs to resolve the issue and re-run the migration script.

  The log file for **Source** is located: `registration_$(date +%Y%m%d_%H%M%S).log`

  The log file for **Protection Groups** is located: `create_protection_groups_$(date +%Y%m%d_%H%M%S).log`
- Verify the new source in the UI: Navigate to **Data Protection** > **Sources**.

Configure Privileges for Non-Admin Users

Previously, the PostgreSQL connectors required admin or superuser privileges for backups due to PostgreSQL's design for physical backups using CHECKPOINT. See, CHECKPOINT in PostgreSQL for more details.

Cohesity now supports granular-level database permission. From PostgreSQL version 15 onwards, non-admin users can perform backup operation by applying the required configurations and grants. This can be configured manually and script based.

Ensure that:

- The instance must be running PostgreSQL version 15 or later.
- The PostgreSQL instance is up-and-running on the specified port.
- The user executing the script must have PostgreSQL superuser credentials.

To configure the privileges using script:

1. Navigate to Datasource Agent Installation path.
2. Run the *configure-pg-backup.sh* script file with the following parameters on the PostgreSQL host machine where the database is installed:

   ```
   <installation_dir>/configure-pg-backup.sh --host <host> --port <port> --db <database> \
   ```

   ```
   --superuser <user> --superuser-pass <password> \
   ```

   ```
   --db-user <backup_user> [--restart true|false] [--pgctl /path/to/pg_ctl]
   ```

   For example:

   ```
   ../configure-pg-backup.sh --host 10.14.20.60 --port 6004  \
   ```

   ```
   --superuser 'postgres' --superuser-pass xxxxxx \
   ```

```
--db-user user1 --restart true --pgctl /usr/pgsql-15/bin/pg_ctl
```

Refer to the table below for script parameters

3. You need to restart PostgreSQL database to reflect the changes. Run the following command:

```
pg_ctl restart -D <datadirectory>
```

Script parameters	Parameter descriptions	Default values
--host	Database host	127.0.0.1
--port	Database port	5432
--db	Database to connect	postgre
--superuser	PostgreSQL superuser (required)	N/A
--superuser-pass	Superuser password (required)	N/A
--db-user	Database user to whom privileges to be granted (required)	N/A
--restart	Restart PostgreSQL after applying settings (true or false)	false
--pgctl	Explicit path to pg_ctl binary (if not in PATH)	N/A
-h, --help	Display help and exit	N/A

To configure the privileges manually:

1. Log on to the PostgreSQL database as a super user and execute the following command to grant the required roles to the backup users:

```
GRANT pg_read_all_settings, pg_checkpoint TO <dbuser>;
```

2. Grant the execute permissions on the SQL functions using the following command:

```
GRANT EXECUTE ON FUNCTION pg_switch_wal() TO <dbuser>;
GRANT EXECUTE ON FUNCTION pg_backup_start(text, boolean) TO <dbuser>;
GRANT EXECUTE ON FUNCTION pg_backup_stop(boolean) TO <dbuser>;
```

3. Edit and update the archive parameter in postgresql.cong file with the following parameters:

```
archive_mode = on
wal_level = replica
```

```
archive_command = '<installation_dir>/stream-log.sh %p %f
<installation_dir>/archive_config_<system_id>'
```

4. Restart PostgreSQL database using the following command. The restart is required to reflect the changes to archive_mode and wal_level:

```
pg_ctl restart -D <data_directory>
```

Sample output:

```
==> Checking server version...
[18-11-2025 11:11:40 IST configure-pg-backup.sh:154] Info: server_version_
num=150015
[18-11-2025 11:11:40 IST configure-pg-backup.sh:192] Info: Generated
archive_command: /opt/cohesity/postgres/scripts/stream-log.sh %p %f /op-
t/cohesity/postgres/scripts/archive_config_7573935469751098619
==> Verifying required predefined roles exist...
==> Ensuring target db user exists ...
==> Setting WAL archiving parameters...
==> Reloading config...
==> Granting roles...
==> Granting EXECUTE on backup/WAL functions ...
[18-11-2025 11:11:41 IST configure-pg-backup.sh:141] Info: Using /usr/pgsql-
15/bin/pg_ctl restart -D /data/pgdata/testsuper1
waiting for server to shut down.... done
server stopped
waiting for server to start....2025-11-18 11:11:41.133 IST [2948688] LOG:
redirecting log output to logging collector process
2025-11-18 11:11:41.133 IST [2948688] HINT: Future log output will appear in
directory "log".
done
server started
[18-11-2025 11:11:41 IST configure-pg-backup.sh:143] Info: pg_ctl restart OK
```

Register and Manage the PostgreSQL Source

To protect your PostgreSQL databases, you need to register your PostgreSQL deployment as a source on theCohesity Platform. After registering your PostgreSQL deployment as a source, you can:

- Update the PostgreSQL source configuration.
- Refresh the source details.
- Unregister the PostgreSQL source from the Cohesity Platform.

Register PostgreSQL as a Source

Before you register PostgreSQL as a source, ensure that you:

- Review the considerations.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for PostgreSQL Protection.

To register PostgreSQL as a source:

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** at the top-right of the page. Click **Databases** > **PostgreSQL**.
3. In the **Source Details** section of the **Register PostgreSQL** dialog, perform the following:
   a. Specify a unique **Source Name** to identify the PostgreSQL deployment.
   b. Specify the Host details such as IP Addresses and Port number.

      > **Note:** For HA cluster deployments, specify all Patroni member node names or IP addresses in the hostname field, separated by commas. The first node in the list is the default target for restore operations. To change the restore target node, update the source configuration and reorder the hostname list so that the preferred restore target appears first.

   c. In the **Datasource Agent Installation Path**, enter the directory path on the PostgreSQL Linux machine, where you have installed the PostgreSQL connector scripts.

      For example: `/opt/cohesity/ postgres/scripts/`.

      For more information, see Install the PostgreSQL Connector.
   d. In the **Authentication Method** section, configure anyone of the following based on your requirements and then click **Proceed**.

Authentication Method	Fields	Description
User Credentials	Select if you want to use password-based authentication. **Note:** For more information on the required setting for client authentication, see Client Authentication in PostgreSQL documentation.	
	Username	The PostgreSQL database username. **Note:** Ensure that the PostgreSQL database user has superuser privileges or equivalent. Refer to Plan and Prepare for PostgreSQL Protection section to grant super privileges.
	Password	The PostgreSQL database user's password. **Note:** If the database is configured to use authentication methods that do not require passwords such as ident, trust, cert, or similar–the **Password** field can be blank.
Kerberos	Select if you want to use Kerberos-based authentication.	
	Kerberos Principal	The Kerberos principal that is required to connect to the database.
	Keytab File Path	The absolute path of the keytab file for the given Kerberos principal. **Note:** The path must be accessible by the Cohesity Agent user.

Authentication Method	Fields	Description
Certification	Select if you want to use certification-based authentication	
	Username	The PostgreSQL database user with superuser or other user with the equivalent privileges.
	SSL Certificate Directory	Specify the absolute path of the directory where the SSL certificates such as client.crt, client.key, or root.crt are located. **Note:** The client certificate should be in X.509 format (V1 or V3).

4. In the **Optional Configuration** of the **Register PostgreSQL** dialog, specify the following parameters:
   a. In the **Host Configurations**, specify the Hostname or IP address of PostgreSQL in **PostgreSQL Hostname/IP** field.

      > **Note:** This is required in case of private network. If not specified, the system uses localhost by default. Ensure that the PostgreSQL server is reachable on the IP address or hostname provided for the Source, then specify the IP address or hostname on which the PostgreSQL server is listening.

   b. Specify the absolute directory location where PostgreSQL executable files (binaries) are installed on your PostgreSQL host machine in PostgreSQL binaries path.

      For example: `/usr/pgsql-15/bin/`

   c. In **Advanced Settings**, specify the **Environment Variables**.

      Environment variables override defaults and allow fine-tuning of backup performance, storage usage, and cleanup. They influence paths, compression, snapshot size, and cleanup actions.

   d. Enable **Check Database Connectivity** if you want to verify the PostgreSQL connection as a part of source verification.

      Default: Disabled

e. Enable **Archive Mode** to restart the PostgreSQL during the first backup. Archive mode is a PostgreSQL setting that enables archiving Write-Ahead Logs (WAL) files. This will restart your database as part of the first full backup and may lead to downtime, so plan accordingly before enabling this option.

   Default: Disable

5. Click **Register**.

Update the PostgreSQL Source Configuration

After registering your PostgreSQL as a source on Cohesity, you can update the PostgreSQL source configuration.

> **Important:** You cannot modify the source name.

To update the PostgreSQL source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the PostgreSQL source and select **Edit**.
3. On the **Update Universal Data Adapter** page, update the required configurations.
4. Click **Update**.

Unregister the PostgreSQL Source

If you plan to stop protecting your PostgreSQL databases, you can unregister the PostgreSQL source from the Cohesity Platform.

To unregister the PostgreSQL source:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the PostgreSQL source and select **Unregister**.

Backup PostgreSQL Database

Cohesity Platform provides a simple workflow to back up your PostgreSQL deployment. After registering the PostgreSQL as a source on Cohesity DataPlatform, you can start configuring your backups.

Backup Process

Cohesity Platform provides a PostgreSQL connector for backup and restore of PostgreSQL deployments. Cohesity PostgreSQL connector uses the JDBC interface to connect the PostgreSQL Server. After registering the PostgreSQL source on the Cohesity Platform, you can create a Protection Group with the required parameters and start protecting PostgreSQL deployment.

Backup Behavior for HA Cluster Deployments

- For Patroni-managed PostgreSQL deployments, backups are always executed on the current leader node. Standby and replica nodes are not backed up directly.
- After a failover or switchover event that changes the Patroni leader, you must immediately trigger a full backup. The next scheduled backup after a leader change will automatically be a full backup.
- Ensure that the Cohesity Linux Agent and PostgreSQL Connector are installed and running on all Patroni member nodes to support automatic leader detection during backup.

Protect PostgreSQL Deployments

In Cohesity Platform, PostgreSQL deployments are protected using a Protection Group and a Protection Policy. The Protection Group uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. You can select a standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

> **Important:** If you have upgraded your PostgreSQL server version after registering the source, it's important to update the directory path to the PostgreSQL binaries that you set using the `--pg-bin` parameter during registration. Make sure to update the path to the new PostgreSQL binaries path, such as "`/usr/pgsql-14/bin`" if `pgsql-14` is the upgraded version.
> Also, you need to create a new Protection Group to protect the upgraded PostgreSQL instance.

To protect PostgreSQL deployments:

1. In the Cohesity Dashboard, do one of the following:
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Databases** > **PostgreSQL**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the PostgreSQL source and then select **Protect**.
2. From the **Source** drop-down list, select the registered PostgreSQL source or click **Register Source** to register a PostgreSQL source. For more information, see Register PostgreSQL as a Source.
3. Click **Add Objects** and **Add Objects** dialog is displayed. Select the required instance. Cohesity displays all the database in the selected instance. Select the instance and click **Continue**.

> **Note:** Cohesity supports instance level and NOT database level protection.

4. Click **Protection Group**, to specify a unique name.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Manage Protection Policy.

   Note: If you re initialize an EnterpriseDB PostgreSQL instance to enable Transparent Data Encryption (TDE), the encrypted instance is treated as a new instance. Existing protection groups created for the unencrypted instance cannot be reused. You must create a new protection group to protect the encrypted instance.
6. Click **More Options** and perform the following:
   1. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
   2. In the **Start Time** field, define when the protection run should start. This option is only available if the selected policy is set to Backup Daily. Enter the hour and minutes, or use the up and down arrows on your keyboard.

      Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
   3. In the **SLA**, configure the settings which defines the expected level of service for backup and recovery operations.

      By default, 2 hours for Full Backup and 1 Hour for incremental backup is defined.
7. After you have completed the **Settings**, if you need to change any additional settings on the **New Database Protection Group** page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
8. Click **Protect**.

Additional Settings

Additional Settings	Description
Convert to Full Backup on Node Failure	Enable this option if you want to convert the incremental backup to full backup in case of an error.
Pause Future Run	Enable **Pause Future Run** only if you do not want to schedule future runs of the Protection Group.

Additional Settings	Description
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the cluster. By default, maximum scaling is set to 16.
Alerts	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.

Additional Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity Platform supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Description	**Optional**. Enter a brief description about the Protection Group.

Manage the PostgreSQL Backup Configuration

In the Cohesity Platform, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Object(s) to the Protection Group

You can add new databases or tables to the existing Protection Group for backup. The newly added object is backed up from the next protection run, and it will be a full backup for the newly added object.

To add object(s) to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Objects** section.
5. In the **Objects** section, Click the plus icon (**+**) to enter multiple objects that you plan to back up. click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired, then a full backup is triggered.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Cancel Protection Run

To cancel a recovery run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click the action icon (⋮) next to the protection run and click **Cancel Run**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for PostgreSQL.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - **Delete Snapshots**: If you enable **Delete Snapshots** and then click **Delete Group**, protection group along with the snapshot will be deleted.
   - **Delete Group**
5. Click **Delete Group**.

Recover PostgreSQL Database

You can restore the protected PostgreSQL data to the same PostgreSQL deployment or to a different PostgreSQL deployment.

Instant Restore

Instant restore allows you to start a PostgreSQL instance instantly after the restore. Instant restore mounts the Cohesity View as a PostgreSQL data directory and immediately starts a PostgreSQL instance pointing to the same data directory.

For EDB PostgreSQL with TDE enabled, Cohesity restores encrypted files as-is. To start PostgreSQL after the restore, ensure that the master wrapping key is available on the restore target. If the master wrapping key is not available, the restore can complete but PostgreSQL will not start.

The pulse logs show the instance restore's status and whether the PostgreSQL instance is ready to use. The following is a sample of the pulse log:

```
Instant restore is done, it's ready for use. Please cancel the restore job
when done using the instantly restored data
```

Note:

- For external tablespaces, soflinks are created from the mounted Cohesity View (pointing to the tablespace data) to the expected tablespace location on the local filesystem.
- The instant restore run is shown in the **Running** state, until you explicitly cancel it from the Cohesity UI.
- To shut down the PostgreSQL instance that started using instant restore, you must cancel the instant restore run from Cohesity UI. Once you cancel the instant restore run, the PostgreSQL instance is stopped, and the Cohesity View is unmounted from the cluster.
- Instant restore does not support point-in-time recovery (PITR).

Recover using Instant Restore to Original Location

Instant restore to original location allows you to instantly restore the PostgreSQL instance to the existing location by selecting the protection group, choosing the new source, and configuring recovery options.

To restore PostgreSQL:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Databases** > **PostgreSQL**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character ***.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. Click **Next: Recover Options**.
6. In the **Recover To** section, to restore to the **Original Location** and type **YES** to perform instant restore to the same PostgreSQL node.

> **Note:** If you enable **Start PostgreSQL Server After Restore** for an encrypted (TDE-enabled) EDB instance, PostgreSQL must have access to the master wrapping key at the restore target to start successfully. Recovery fails if you attempt to restore a non encrypted snapshot to an encrypted PostgreSQL instance.

7. Configure the following parameters for **Recovery Options**:
   - Enable **Start PostgreSQL Server After Restore**, to start the PostgreSQL server after the data restore is complete. By default, **Start PostgreSQL Server After Restore** is disabled.

     > **Note:** Ensure that the port you have specified during the PostgreSQL source registration is not in use.

   - Disable **Permissions for Restored Files** option to specify the permissions to the restored files or directories. The permission options is displayed. Select the required permissions.

     By default, the **Permissions for Restored Files** option is enabled which applies the original permissions to the restored files or directories as it was in the source.
   - In the **PostgreSQL Server CLI** option, enter the options required to start the PostgreSQL instance. For more information on the options, see pg_ctl command reference in the PostgreSQL documentation.

     > **Note:** Ensure that you do not pass the **-D <datadir>** option.

   - For the **Mount/VIPs**, **Concurrency**, and **Task Name** fields, refer to the table below.

Options	Description
Mounts/VIPs	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	Number of parallel streams to use for reading data from Cohesity storage.
Task Name	Change the default name of the recovery task.

8. Click **Recover**.

Recover using Instant Restore to a New Location

Instant restore to new location allows you to instantly restore the PostgreSQL instance to a new location by selecting the protection group, choosing a new source, and configuring recovery options.

When restoring an EDB PostgreSQL instance with TDE to a different location, ensure that the target system is configured with the same key store access and wrap/unwrap command configuration as the source. Cohesity recovers encrypted data files as-is and does not modify `postgresql.conf` or key management settings. PostgreSQL must be able to access the unwrap command and master wrapping key at startup.

To restore PostgreSQL:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Databases** > **PostgreSQL**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character ***.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. Click **Next: Recover Options**.
6. In the **Recover To** section, click **New Location** option and select the appropriate **Source** from the drop-down menu.
7. Configure the following parameters for **Recovery Options**:
   - Enable **Start PostgreSQL Server After Restore**, to start the PostgreSQl server after the data restore is completed. By default, **Start PostgreSQL Server After Restore** is disabled.

     > **Note:** Ensure that the port you have specified during the PostgreSQL source registration is not in use.

   - In the **PostgreSQL Service User**, specify the user having access to PostgreSQL with super-user privileges. The default user is "`postgres`".

     > **Note:** Cohesity now supports granting admin privileges to other non-admin users. For more information, see Plan and Prepare for PostgreSQL Protection section.

- Disable **Permissions for Restored Files** option to specify the permissions you want to set for the files or directories after the restoration. Select the appropriate privileges.

  By default, **Permissions for Restored Files** option is enabled to apply original permissions to the restored files or directories as it was in the source.
- In the **PostgreSQL Server CLI** options field, enter the options required to start the PostgreSQL instance. For more information on the options, see pg_ctl command reference in the PostgreSQL documentation.

  > **Note:** Ensure that you do not pass the **-D <datadir>** option.
- For the **Mount/VIPs**, **Concurrency**, and **Task Name** fields, refer to the table below.

Options	Description
Mounts/VIPs	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	Number of parallel streams to use for reading data from Cohesity storage.
Task Name	Change the default name of the recovery task.

8. Click **Recover**.

Regular Restore

Regular restore allows you to restore PostgreSQL instance using latest backups and applying write-ahead-logging (WAL) files during recovery. This recovery method allows you to specify mapping of the PostgreSQL tablespaces to an alternate directory.

Note:

- For versions prior to 7.1.2_u4, redirected restore is not supported for external tablespaces. All tablespaces are restored to the same location as the source. Ensure that the same location as the source is available on the cluster where you are restoring the PostgreSQL data. If the same locations are unavailable on the target cluster, you can create softlinks pointing to the actual storage location.

  For example:

  ```
  ln -s <actual-storage-location> <tablespace-location-same-as-source>
  ```

For EDB PostgreSQL with TDE enabled, Cohesity recovers encrypted files as-is. To start PostgreSQL after the recovery, ensure that the master wrapping key is available on the restore target. If the master wrapping key is not available, the recovery can complete but PostgreSQL will not start.

Recover using Regular Restore to the Original Location

Perform a regular restore of the PostgreSQL instance to its original location by selecting the protection group, confirming restoration, and configuring recovery options.

To restore PostgreSQL:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Databases** > **PostgreSQL**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character ***.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. Click **Next: Recover Options**.
6. In the **Recover To** section, click **Original Location** and type **YES** to perform instant restoration to same PostgreSQL node.
7. Configure the following parameters for **Recovery Options**:
   - Enable **Start PostgreSQL Server After Restore**, to start the PostgreSQL server after the restore is complete. By default, **Start PostgreSQL Server After Restore** is disabled.

> **Note:** Ensure that the port you have specified during the PostgreSQL source registration is not in use.
> If you enable **Start PostgreSQL Server After Restore** for an encrypted (TDE-enabled) EDB instance, PostgreSQL must have access to the master wrapping key at the restore target to start successfully. Recovery fails if you attempt to restore a non encrypted snapshot to an encrypted PostgreSQL instance.

- Disable **Permissions for Restored Files** option to specify the permissions to the restored files or directories. The permission options is displayed. Select the required permissions.

  By default, the **Permissions for Restored Files** option is enabled which applies the original permissions to the restored files or directories as it was in the source.
- In the **PostgreSQL Server CLI** options field, enter the options required to start the PostgreSQL instance. For more information on the options, see pg_ctl command reference in the PostgreSQL documentation.

  > **Note:** Ensure that you do not pass the **-D <datadir>** option.
- For the **Mount/VIPs**, **Concurrency**, and **Task Name** fields, refer to the table below.

Options	Description
Mounts/VIPs	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	Number of parallel streams to use for reading data from Cohesity storage.
Task Name	Change the default name of the recovery task.

8. Click **Recover**.

Recover using Regular Restore to a New Location

Perform a regular restore of PostgreSQL to a new location by selecting the protection group, specifying new data/log paths, configuring overwrite and recovery options.

> **Note:**
> Using this recovery method you can:
>
> - Map the directory to an alternate directory by specifying the parameters in the **Tablespace Directories** field.
> - Recover PostgreSQL using the latest backup snapshots and WAL file using **Roll Forward Database** option.

To restore PostgreSQL:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** at the top right of the page and then select **Databases** > **PostgreSQL**.
3. Search for the Protection Group that contains the snapshots to recover from by entering characters of the Protection Group name, or by specifying the wildcard character *.
4. From the list of search results, select the Protection Group that contains the snapshots to recover.
5. Click **Next: Recover Options**.
6. In the **Recover To** section, recovering to a New Location, configure the following fields:
    a. select the appropriate **Source** from the drop-down menu.
    b. Specify the new path on the target cluster where you want to restore the new path in the **New Data Path** field.
    c. Specify the absolute path of the directory for PostgreSQL server logs in **New Log Path** (optional) field.
    d. Enable **Overwrite existing database(s)** option and type **YES**, if you want to overwrite the existing instance with the same name.
7. In the **Recover Options**, configure the following fields:
    a. Enable **Start PostgreSQL Server After Restore**, if you want to start the PostgreSQL server after the data restore is complete. By default, **Start PostgreSQL Server After Restore** is disabled.

> **Note:**
> Select this option:
>
> - Ensure that the port you have specified during the PostgreSQL source registration is not in use.
> - This option must be selected when using **Roll Forward Database** option.
> - When restoring an EDB PostgreSQL instance with TDE to a different location, ensure that the target system is configured with the same key store access and wrap/unwrap command configuration as the source. Cohesity restores encrypted data files as-is and does not modify postgresql.conf or key management settings. PostgreSQL must be able to access the unwrap command and master wrapping key at startup.

b. In the **PostgreSQL Service User**, specify the user having access to PostgreSQL with super-user privileges. The default user is "`postgres`".

> **Note:** Cohesity supports granting admin privileges to other non-admin users. For more information, see Configure Privileges for Non-Admin Users.

c. Disable **Permissions for Restored Files** to specify the permissions you want to set for the files or directories after the restore. Select the appropriate privileges.

   By default, **Permissions for Restored Files** field is enabled to apply original permissions to the restored files or directories as it was in the source.

d. Enable **Roll Forward Database** to recover the instance using the most recent snapshot backup along with write-ahead-logging (WAL) files that are backed up to Cohesity. Select one of the options from the following:

> **Note:** By default, **No Rollforward** option is selected.

   i. **Automatic Rollforward**: Cohesity automatically applies latest snapshot along with the WAL files during recovery. Specifying the path of the WAL files is optional.

ii. **Manual Rollforward**: Cohesity only copies the WAL files; you must apply them manually to achieve point-in-time recovery. Specify the path where the WAL files need to be copied with read-write permissions.

For more information, see Roll Forward Recovery Mode section.

e. In the **Tablespace Directory** field, Cohesity auto-populates the directories. You need to specify the path for each tablespace along with their names.

For example: /tmp/ts1. The /tmp/ts1 is the absolute alternate path of the directory where you want tablespace to be restored.

> **Note:**
> - If path is not specified, PostgreSQL tablespace directory path from backup instance will be used.
> - If the directories are not auto-populated, refer to the Directories are not auto-populated for Tablespace during recovery in the troubleshooting section.

f. In the **PostgreSQL Server CLI** options field, enter the options required to start the PostgreSQL instance. For more information on the options, see pg_ctl command reference in the PostgreSQL documentation.

> **Note:** Ensure that you do not pass the **-D <datadir>** option.

g. For the **Mount/VIPs**, **Concurrency**, and **Task Name** fields, refer to the table below.

Options	Description
Mounts/VIPs	Enter the number of VIPs you want to use while reading data from the Cohesity storage.
Concurrency	Number of parallel streams to use for reading data from Cohesity storage.
Task Name	Change the default name of the recovery task.

8. Click **Recover**.

Roll Forward Database

The **Roll Forward Database** feature allows to recover PostgreSQL instance using the most recent snapshot backup along with write-ahead-logging (WAL) files that are backed up to

Cohesity but not part of a log backup. This feature is supported only when recovering to the latest snapshot backup and applying all WAL files that are available after it.

To enable this option during recovery, select **Regular Restore** to a **New Location** option.

> **Note:** Default is **No RollForward** where recovery is performed without applying the WAL files.

Cohesity supports two roll forward recovery modes:

1. **Automatic RollForward**: Cohesity automatically applies these logs during recovery. Select this method for quick recovery.
2. **Manual RollForward**: Cohesity only copies WAL files to the specified path. Specify the path where the WAL files need to be copied with read-write permissions.

Refer the following table to understand the differences between **Automatic RollForward** and **Manual RollForward** options.

UI Field	Automatic RollForward	Manual RollForward	UI Location
Start PostgreSQL Server After Restore	This option must be **enabled**. Enabling this option will re-start the database, automatically. If not enabled, manual restart is required.	This option must be disabled. If enabled, restoration fails. After the recovery, the database is left in stopped state. You need to start the database manually. For more information to start the database, refer PostgreSQL: Documentation: 15: 19.3. Starting the Database Server.	**Recovery Options**
WAL Files Path	**Optional**: Cohesity applies all the WAL files.	**Required**: Specify the path where WAL files, with read-write permission, will be copied. Cohesity only copies WAL files; you must apply them manually to achieve point-in-time recovery. However, you need to apply them manually after snapshot restore. If not specified, recovery is performed without applying the backed-up WAL files. For more information, see PostgreSQL document.	

> **Note:** In case of **Automatic RollForward**, Cohesity does not support timestamp-based selection of the WAL files; therefore, all the available WAL files are applied during recovery. Using Instant restore mode of recovery is not supported.

Cancel a Recovery Run

When you cancel a recovery run, the PostgreSQL connector cancels the job by terminating the recovery processes.

To cancel a recovery run:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click the action menu (⋮) next to the recovery run and click **Cancel Run**.
3. In the **Cancel Recovery** dialog, select any one of the following:
    1. Click **Continue Recovery** if you want to continue a recovery run.
    2. Click **Cancel Recovery** if you want to cancel the recovery run.

Recover HA Cluster Deployment

This topic describes the recovery workflows for PostgreSQL high availability (HA) cluster deployments. The procedures and examples use Patroni as a sample configuration; however, they can be used as a reference for any PostgreSQL HA deployment. Adjust the commands and configuration paths as needed for your specific HA framework.

Depending on your recovery requirements, you can:

- Perform a full disaster recovery to rebuild the entire cluster.
- Restore an individual node as a standalone PostgreSQL instance.
- Convert a standalone instance back into an HA-managed cluster.

Recover a Patroni-Managed PostgreSQL Cluster (Disaster Recovery)

When recovering a Patroni-managed PostgreSQL cluster, Cohesity restores data to a single node (the first node defined in the source registration). After the restore, you must manually rebuild the remaining standby and replica nodes to re-form the Patroni cluster.

> **Note:** While this section uses Patroni as a reference, the recovery procedures described here are applicable to any PostgreSQL HA cluster deployment. Adjust the commands and configuration paths as needed for your specific setup.

The overall disaster recovery workflow consists of three phases:

- Pre-Restore
- Restore
- Post-Restore

Pre-Restore

Before restoring the Patroni cluster, perform the following steps:

1. Stop Patroni on all nodes:

   ```
   systemctl stop patroni
   ```

2. Remove the old cluster state from the Distributed Configuration Store (DCS):

   ```
   patronictl -c /etc/patroni.yml remove <cluster-name>
   ```

   > **Note:** Replace `<cluster-name>` with the name of your Patroni cluster. The Patroni configuration file path may vary based on your installation.

3. Identify the restore target node. Cohesity restores to the first node defined in the source registration's hostname field. To restore to a different node, update the source registration and move the preferred node to the first position in the hostname list before initiating the restore.

Restore (Single Identified Node)

Recover the PostgreSQL data using a Regular Restore to the original location. Follow the standard restore procedure described in the Recover using Regular Restore to the Original Location topic.

> **Note:** Only PITR (point-in-time recovery) restores are supported for Patroni-managed PostgreSQL deployments.

Post-Restore (Recover Standby and Replica Nodes)

After the restore completes successfully on the identified node, perform the following steps to rebuild the Patroni cluster:

1. Stop the PostgreSQL instance that was started by Cohesity during recovery:

   ```
   pg_ctl stop -D <data-directory> -m fast
   ```

   Example:

   ```
   pg_ctl stop -D /var/lib/pgsql/16/data -m fast
   ```

2. Start Patroni on the restored node only:

```
systemctl start patroni
```

3. Verify that the restored node has assumed the Leader role:

```
patronictl -c /etc/patroni.yml list
```

Expected output shows the restored node as Leader in running state.

4. Rebuild each standby or replica node:

> **Note:** Perform on all nodes that were NOT restored.

1. Stop Patroni (confirm it is stopped):

```
systemctl stop patroni
```

2. Remove the old PostgreSQL data directory contents:

```
rm -rf <data-directory>/*
```

> **Important:** This command permanently deletes all data in the specified directory. Verify the correct path before executing.

3. Start Patroni to trigger automatic replica rebuild:

```
systemctl start patroni
```

> **Note:** Wait at least 10 seconds before starting Patroni on the next standby node. It is recommended to verify the node has joined the cluster using 'patronictl list' before proceeding to the next node.

5. Verify the complete cluster is operational:

```
patronictl -c /etc/patroni.yml list
```

The output should show all member nodes with their respective roles (Leader, Replica) and a **running** state.

Restore a Patroni Node as Standalone PostgreSQL

You can restore a backup from a Patroni-managed source to a standalone PostgreSQL instance. This is useful when you need to recover data without rebuilding the full Patroni cluster.

Pre-Restore

1. Remove the target node from the Patroni source registration:

   Update the PostgreSQL Source Configuration to remove the node you plan to restore as standalone.

2. Register the node as a standalone PostgreSQL source:

   Follow the standard source registration procedure in Register PostgreSQL as a Source.

3. Stop Patroni on the target node only:

```
systemctl stop patroni
```

Restore

Restore to the alternate host by selecting the standalone source registered in the pre-restore step. Follow the standard procedure in the Recover using Regular Restore to a New Location topic.

> **Note:** Only PITR (point-in-time recovery) restores are supported for Patroni-managed PostgreSQL deployments.

Convert a Standalone PostgreSQL Instance Back to Patroni

After restoring a Patroni node as a standalone PostgreSQL instance, you can convert it back to a Patroni-managed cluster using the following procedure:

1. Stop Patroni on all cluster nodes:

```
systemctl stop patroni
```

2. Remove the old cluster state from the DCS:

```
patronictl -c /etc/patroni.yml remove <cluster-name>
```

3. Stop the restored standalone PostgreSQL database:

```
pg_ctl stop -D <data-directory> -m fast
```

4. Start Patroni on the restored node only:

```
systemctl start patroni
```

5. Verify the restored node is the Leader:

```
patronictl -c /etc/patroni.yml list
```

6. Rebuild each standby/replica node:

   Perform the following on all nodes that were NOT restored:

1. Stop Patroni:

   ```
   systemctl stop patroni
   ```

2. Remove old data directory contents:

   ```
   rm -rf <data-directory>/*
   ```

3. Start Patroni:

   ```
   systemctl start patroni
   ```

   > **Note:** Wait at least 10 seconds before starting Patroni on the next node.

7. Verify the complete cluster:

   ```
   patronictl -c /etc/patroni.yml list
   ```

8. Update the source registration:
   1. Remove the standalone PostgreSQL source that was registered during the pre-restore steps.
   2. Add the restored node back as a hostname under the original Patroni source registration. See Update the PostgreSQL Source Configuration.

Troubleshooting PostgreSQL Protection

The following table lists the logs that you can use to troubleshoot failures during the data protection workflows:

Component	Path to the log file
Cohesity PostgreSQL Connector logs	`/var/log/cohesity/uda/`
Cohesity Linux Agent service logs	`/var/log/cohesity/`

Directories are not Auto-populated for Tablespace during Recovery

When recoverying a PostgreSQL instance using the new workflow, the directories for tablespaces may not auto-populate. In such cases, manually specify the mapping between directories and tablespaces using the following format:

1. In the **Tablespace Directory** field, specify the tablespace along with the path in the following format:

   ```
   <TablespaceName>:/<AbsolutePath>
   ```

For multiple tablespace entries, specify comma separated values.

For example, `Ts1:/tmp/ts1,Ts2:/mytmp/ts2`

2. Click ☑.

In the above example,

- `Ts1` and `Ts2` are tablespaces.
- `/tmp/ts1` is the absolute path of the directory where you want to restore `Ts1`.
- `/mytmp/ts2` is the absolute path of the directory where you want to restore `Ts2`.

Chapter 14: NAS

Cohesity supports backing up multiple NAS volumes in a single Protection Group. Protection Groups can run repetitively on a schedule.

To learn more about the relationship between Protection Groups and Policies, see Manage Protection Policy. To learn more about managing Protection Groups, see Manage Protection Groups.

For example, you could create a Protection Group with a schedule that backs up specific NAS volumes every hour as illustrated in the following figure.

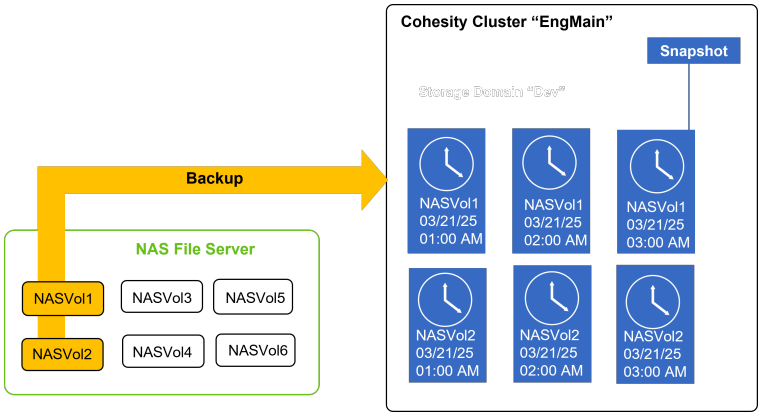

During a Protection Group Run, Cohesity Snapshots of the NAS volumes are saved on the Cohesity cluster with a date and time stamp.

For instructions on how to use the Cohesity Dashboard to back up NAS volumes, see Protect NAS.

Plan and Prepare for NAS Protection

Review and understand the following before you perform the backup and recovery of NAS:

Elastifile

Considerations

- Recovering an NFS 4.1 volume to NFSv3 is not supported.
- For application containers, you must add Cohesity client IPs for all the nodes with read/write access onto the container's root share.
- CohesityCohesity Cloud Edition must have access to the Elastifile load balancer VIP.
- For Elastifile, the Incremental forever backup uses Cohesity built-in Change File Tracking (CFT) to locate the changed file/folder.
- If multiple shares support both SMB and NFS in one Protection Group, you can select all of them to back up using only SMB or NFS.
- Cohesity supports the backup of the symbolic link reparse point, **IO_REPARSE_TAG_SYMLINK** in an SMB share.
- Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.

Generic NAS

Considerations

- Generic NAS mount-based backup can be used to back up any NAS device.
- For Generic NAS, Cohesity scans through the live file system during backup. And, the open or locked files may be excluded during backup.
- Custom Pre/Post scripts can be used to take snapshot-based backups.
- Cohesity supports the backup of the symbolic link reparse point, **IO_REPARSE_TAG_SYMLINK** in an SMB share.
- Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.
- Encryption authentication with NFS v4.1 is not supported.

> **Note:**
> - NFS export permissions may be required on the source to allow the Cohesity IPs to be added to the Nutanix export rules for Generic NAS share.
> - Generic NAS protection supports the backup of regular files, directories, symlinks (only in NFS), and hardlinks.
> - Kerberos authentication with NFS v4.1 is supported.

IBM Spectrum Scale (GPFS)

Considerations

- Cohesity only supports fileset-based NFS backups that use simple user-based authentication.
- Unlinked filesets are not displayed in the cluster source tree.
- Junction Path can be used to create a fileset in a fileset, but Cohesity excludes these during backups as they can have cyclic dependencies.
- Cohesity supports only the backup and restore of POSIX attributes for GPFS. The backup of ACLs is not supported for GPFS.
- For IBM Spectrum Scale, Incremental forever backup uses Cohesity built-in Change File Tracking (CFT) to locate the changed file or folder.
- Cohesity supports NFS-based backup and restore of GPFS cluster.
- Cohesity supports the backup of the symbolic link reparse point, **IO_REPARSE_TAG_SYMLINK** in an SMB share.
- Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.

Prerequisites

- The NFS protocol must be enabled on the GPFS cluster.
- GUI services must be installed on at least one of the GPFS nodes as REST APIs are enabled only if the GUI services are enabled that are used for registration and other orchestrations.
- Cohesity uses the IP addresses of Cluster Export Services(CES) to connect the GPFS cluster using the NFS protocol. Therefore, at least one IP address of Cluster Export Services(CES) must be reachable from the Cohesity cluster nodes. The CES IP addresses that are not reachable must be added to the **Deny IPs** field while registering the GPFS cluster with the Cohesity cluster. Else, the backup of the GPFS cluster will fail.
- If AD is not configured make sure the file level authentication is set up properly.

Isilon

For information about the prerequisites and considerations for Isilon, see Dell EMC Isilon.

NetApp

For information about the prerequisites and considerations for NetApp ONTAP, see NetApp ONTAP.

Nutanix Files

For information about the prerequisites and considerations for Nutanix Files, see Nutanix Files.

Pure FlashBlade

Considerations

- Incremental forever backup uses Cohesity built-in Change File Tracking (CFT) to locate the changed file or folder.
- If multiple shares support both SMB and NFS in one protection group, you can select all of them to back up using only SMB or NFS.
- Cohesity supports the backup of the symbolic link reparse point, **IO_REPARSE_TAG_SYMLINK** in an SMB share.
- Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.

Prerequisites

- For Cohesity to protect Pure FlashBlade File Systems:
  - Each File System must have NFS and/or SMB enabled.
  - Each File System must have the Special Directory "Snapshot" enabled.
- To register the Pure Storage Flash Blade with the Cohesity cluster, ensure that you have the default *pureuser* FlashBlade administrator account. If not, create the *pureuser* with full admin permissions.
- For Pure Storage FlashBlade, the recovery of SMB data may fail if the files and folders being recovered have restrictive permissions. Cohesity recommends you either create folders or directories manually in the target location before initiating the recovery or use an SMB user with read/write permissions when registering the Pure Storage FlashBlade source.

Register or Edit NAS

Registering a NAS makes its content eligible to be protected or backed up on the Cohesity cluster.

1. Navigate to **Data Protection > Sources**.
2. Click **Register > NAS**.
3. Select the **NAS Source**.
4. Enter the information for the NAS.

   Isilon (Cluster)

For information on registering and protecting Isilon, see Dell EMC Isilon.

NetApp

For information on registering and protecting NetApp, see NetApp ONTAP.

GPFS

Cohesity supports NFS-based backup and restore of the GPFS cluster.

Prerequisites:

Before you register the GPFS cluster with the Cohesity cluster, ensure that:

- NFS protocol is enabled on the GPFS cluster.
- GUI services are installed on at least one of the GPFS nodes.
- Cohesity uses the IP addresses of Cluster Export Services(CES) to connect the GPFS cluster using NFS protocol. Therefore, at least one IP address of Cluster Export Services(CES) must be reachable from the Cohesity cluster nodes. The CES IP addresses that are not reachable must be added to the **Deny IPs** field while registering the GPFS cluster with the Cohesity cluster. Else, the backup of the GPFS cluster will fail.

Field Name	Description
Host or IP Address	Enter the host or IP address of one of the GPFS nodes with UI access.
Username	Enter the user name used to access the GPFS cluster. A user with GUI role should register the GPFS on Cohesity. Cohesity uses this GUI user to access the REST APIs. If you are providing a domain user account, then you must specify the user name in the following format: `fully_qualified_domain_name\username`. For example, `xydc.local\user1`.
Password	Enter the password for the specified user.

Field Name	Description
Filter IPs	Enable this option to filter the IP addresses of the GPFS cluster. By filtering IP addresses, you can allow or deny the communication of Cohesity cluster to specific IP addresses or subnets of the GPFS cluster. To filter the IP addresses of the GPFS cluster, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the GPFS cluster source through which the communication to Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the GPFS cluster source through which the communication to Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.

If Active Directory is not configured make sure the file-level authentication is set up properly for the backups and recovery tasks to go through.

Generic NAS

Field Name	Description
Mode	NFSv3, NFSv4.1, or SMB
Mount Path	The path for mounting the NAS. Use the following format: • Hostname or IP:/Volume for NFS mode • \\Hostname or IP\Share path for SMB mode. **Note:** Do not edit an existing mount path. To change the mount path perform the following: Register the new mount path Source. Edit any Protection Groupto remove the old mount path and add the new mount path. Unregister the old mount path Source.

Field Name	Description
Skip Mount Point validation during registration	The option to skip the mount point validation during registration. The following will be skipped: • for NFS mount point, mounting the volumes that have been tested whether the mount point actually exists and is reachable from the Cohesity cluster. • for SMB mount point, establishing connection with the sources and consequently querying some basic information of the share (like size).
Description	An optional description for the mount point.
Username	SMB mount credentials that allow the reading and backing up of files. If you are providing a domain user account, then you must specify the user name in the following format: `fully_qualified_domain_name\username`. For example, `xydc.local\user1`.
Password	SMB mount credentials that allow the reading and backing up of files.

SMB DFS Referral

The Cohesity SMB client now follows the Distributed File System (DFS) referral. The Distributed File System has a collection of shares on different machines that serve a file system namespace. If one machine does not have a particular file requested by the client, it is referred to check the other machine.

Users can now backup the entire DFS namespace by simply registering the share with the hostname directly as a Generic NAS SMB mount point.

> **Note:**
> Before the 6.8 version, registering each share as a Generic NAS source and backing it up in a Protection Group was suggested by Cohesity support as a workaround in some cases. After the upgrade, the Cohesity cluster may end up backing up the same data multiple times causing data bloat and an increase in the run time.
>
> If you are using this method, then you should contact Cohesity Support to transition to the SMB DFS Referral method.

Pure Storage FlashBlade

Field Name	Description
Hostname or IP Address	The hostname or IP address of the FlashBlade to register.
API Token	API Token The API token of the Admin user for logging in to the FlashBlade. You can use the FlashBlade CLI to find the API token of the administrator: `pureadmin list --api-token --expose`
SMB Volumes Access	Enable this option to back up SMB volumes. Provide Active Directory user credentials that allow read and write access on the FlashBlade. SMB support is currently a beta feature.
Filter IPs	To filter the IP addresses of the FlashBlade, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the FlashBlade source through which the communication to Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the FlashBlade source through which the communication to Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. **Note:** The "Filter IPs" feature can be found under “Additional Settings” on the “Protection Group” page and under “Recovery Options” on the “Restore” page. You must enable the feature flag **nasFilterIpsEnabled** for the Filter IPs feature to appear.

Elastifile

You can perform NFS-based backup and restore of Elastifile data containers present on Elastifile Cloud File System.

Prerequisites:

Before you register the cloud edition of Elastifile with Cohesity cluster, ensure that:

- Elastifile Cloud File System is of version 3.1.0.31 and above.
- Elastifile Cloud File System is deployed using load balancer configuration.
- The load balancer VIP is accessible from the Cohesity cluster as it is used for data ingestion during backups and restores.
- The Elastifile user is an admin or has Provisioner role.

Field Name	Description
Host or IP Address	Enter the host or IP address of the Elastifile Management interface to register.
Username	Enter the username used to access Elastifile Cloud File System. The user must be an admin or a user with provisioner role. If you are providing a domain user account, then you must specify the user name in the following format: `fully_qualified_domain_name\username`. For example, `xydc.local\user1`.
Password	Enter the password for the specified user.
Filter IPs	Enable this option to filter the IP addresses of the Elastifile Cloud File System. By filtering IP addresses, you can allow or deny the communication of Cohesity cluster to specific IP addresses or subnets of the Elastifile Cloud File System. To filter the IP addresses of the Elastifile Cloud File System, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the Elastifile Cloud File System source through which the communication to Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the Elastifile Cloud File System through which the communication to Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.

You must open the following ports in the firewall to allow the Cohesity cluster to send and receive data and register Elastifile on Cohesity cluster to perform data backup operations.

Port Number	Network Protocol	Operationss
443	TCP	For EMS/API operations
111 2049 644 4040 4045	TCP	For NFS Operations

Port Number	Network Protocol	Operationss
111 2049 644 4040 4045	UDP	For NFS Operations

5. Click **Register**.

   When registration finishes, a new entry appears on the Sources page.

Automatic Denylisting in NAS Sources

When connecting to a NAS source's IP address, network issues such as timeouts can result in backup failures. Cohesity now supports implementation of automatic denylisting of problematic IP addresses associated with a NAS source. This feature aims to prevent NAS protection jobs from failing due to specific IPs or interfaces becoming unreachable or having limited bandwidth, ensuring smoother backup operations.

Protect NAS

NAS Protection Group captures snapshots and these snapshots are required for the following tasks:

- Recover volumes
- Replicate snapshots
- Archive snapshots
- Files and folders recovery (CloudArchive Direct)

> **Note:**
> - Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.
> - NAS backup and restore do not support Junction points and ADS (Alternate Data Stream).

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit NAS
- Create or Edit a Standard Policy
- Create or Edit Storage Domains

Create a Protection Group for NAS Volumes

1. Navigate to **Data Protection > Protection** and do one of the following:
   - To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
   - To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. If creating a new Protection Group, select **Protect > NAS**.
3. In the **New Protection** pop-up, do one of the following:
   - Protect the NAS volumes database by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered NAS source or click **Register Source** to register a NAS source. From the **Objects** list, select the volumes you want to back up.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. Click **Protect**.
   - To protect NAS volumes by specifying all the options, click **More Options**.
4. Click **Add Objects** and from the **Registered Source** drop-down list, select the NAS source that contains the volumes you want to back up. You can also register a new source with the Cohesity cluster so you can protect it.

5. From the **Objects** list, select the volumes to back up. You can select volumes that are of NFS exports, SMB shares or the volumes that are exposed to both NFS and SMB protocols (mixed volumes). By default, all volumes in the source are excluded as indicated by the unselected check box. Select the volumes you want to include in the Protection Group. A selected check box ( ) indicates the volume is selected to be backed up.

   If appropriate for your organization, enable Auto Protect.

   Click to turn on Auto Protect. Turn on this option to protect new volumes that are added to a selected parent SVM. When a parent SVM is selected and this option is enabled, all volumes that are added to the SVM in the future are automatically added to this Protection Group and become protected on the next run. For example, you create a Protection Group that backs up a SVM. The Protection Group runs for two weeks and then a new volume is added to the SVM. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new volume is also backed up even though the new volume has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click . To exclude an object from Auto Protect, click .

   - Indicates an Auto Protected object.
   - Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

   You can optionally filter the volumes displayed.

   - In the Search field, enter part of volume or SVM name.
   - In the **Show All** drop-down, select one of the following options:
     - **Selected objects**—Displays the volumes that are currently selected for protection.
     - **Unprotected**—Displays the volumes that are not protected by any Protection Groups in the Cohesity cluster—not protected by Cohesity.
     - **Protected** —Displays the volumes that are protected by all the Protection Groups in the Cohesity cluster.
   - You can also filter with the following options:
     - **NFS** -Displays NFS volumes only.

       For Generic NAS and NetApp ONTAP sources, **NFSv3** and **NFSv4.1** filters are supported from Cohesity 6.6.0d version onwards. To enable this

feature, contact Cohesity Support.

- **SMB** - Displays SMB volumes only.

- Selecting the check box adjacent to **Volumes** selects all filtered volumes.

The graphical representations are described below.

- If you are creating a new Protection Group and the shield appears next to the volume, this volume is protected by another Protection Group.

- If you are editing an existing Protection Group and the shield appears next to the volume and the volume is selected ( ), it will be protected by the current Protection Group. (It may also be protected by another Protection Group.)

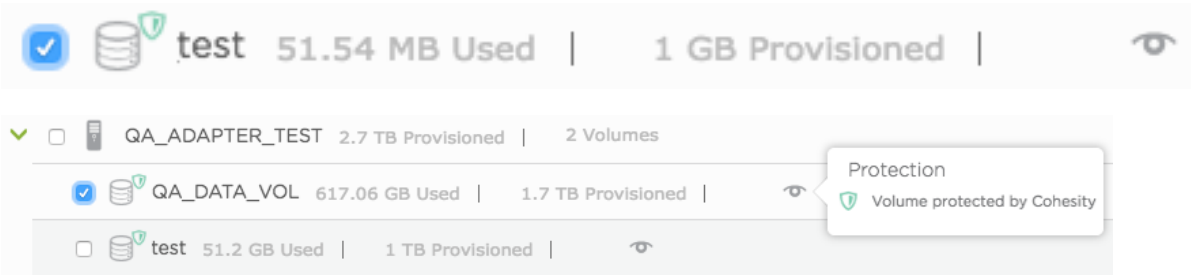

- Click the icon to see details about the object.

6. Select the backup preference mode of the volume you want to backup.
7. Click **Continue** to save and go back to the New Protection page.
8. **Protection Group**: Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes and periods. This field is required.
9. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.
10. **Storage Domain:** Select an existing storage domain or create a new one by selecting **New Storage Domain**.
11. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, Click **Protect**.
12. Optionally, from the **Protection Group** field, you can select an existing Protection Group settings for this Protection Group by enabling **Existing Group**, and then by selecting an existing Protection Group from the **Group** drop-down list.
13. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up

and down arrows on your keyboard. Verify the AM or PM setting.

The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.

14. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
15. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
16. **Pre & Post Scripts:** Toggle **Enable** to run scripts on a Unix server before and/or after a Protection Group runs. If configured, scripts are run every time an object (or entity) is backed by a protection run. For example if there are five mount/shares that are backed up as part of the NAS Protection Group and the Pre Script is configured, then five instances of the Pre Script is executed in parallel. You can distinguish between the scripts that are run on the different objects (entities) using the `COHESITY_BACKUP_ENTITY` environment variable. The Cohesity cluster populates this environment variable with the name of the backed object (or entity).

    > **Note:** The specified Unix server does not need to be attached to the NAS target, it just needs to be a Linux VM or machine that can either invoke APIs on or ssh in to a NAS target.

    Configure the settings for running the Pre and Post scripts.

Hostname or IP Address (Unix only): Specify either the hostname or the IP address of the remote Unix server where the script will run. (Only Unix is currently supported.)

Cluster SSH Public Key: Click **Copy Key to Clipboard** to copy the entire string including the `ssh-rsa` string preceding the key, the public key and the `username@cluster_name` string after the key. Open the `authorized_keys` file (typically located in the `~/.ssh` directory) and paste the entire string to the end of the file.

Username: Specify the username on the remote Unix machine that will execute the script. The username must match the username specified in the `authorized_keys` file.

You can configure a Pre-Script, a Post-Script or both scripts.

Pre-Script

- **Script Full Path:** Specify the full path and name of the script to run. A script matching the specified name must be available in the specified directory. The script can be in any format, such as a shell script or executable. The script must be configured with the permissions to allow execution (`+x`).

  > **Important:** Cohesity does not validate the content of the script and is not responsible for the content of the script or any actions done by the script. You must verify the script actions before a script or executable is run.

- **Script Params**: Optional field for passing one or more parameters to the script. All specified parameter values are passed as a single string to the script.
- **Timeout**: Optional field for specifying how long a script is allowed to execute before the Cohesity cluster stops the execution of the script. By default, the timeout is 15 minutes. If no value is specified, there is no timeout and the script is not stopped. This can cause a protection run to hang if the script never completes. The timeout is applied from the moment the script starts. For example, if the timeout is set to 15 minutes but the script takes 20 minutes to complete, it will be stopped after 15 minutes. Therefore, users should set the timeout based on the expected runtime of the script to avoid premature termination.
- **Continue Backup if script fails**: This toggle applies only to Pre-Scripts and is enabled by default. When this toggle is enabled, the Protection Group backs up the objects even if the Pre-Script fails. If this toggle is disabled and the Pre-Script fails, the backup is canceled but the Post-Script still runs.

Post-Script

- **Script Full Path:** Specify the full path and name of the script to run. A script matching the specified name must be available in the specified directory. The script can be in any format, such as a shell script or executable. The script must be configured with the permissions to allow execution (`+x`).

  > **Important:** Cohesity does not validate the content of the script and is not responsible for the content of the script or any actions done by the script. You must verify the script actions before a script or executable is run.

- **Script Params**: Optional field for passing one or more parameters to the script. All specified parameter values are passed as a single string to the script.
- **Timeout**: Optional field for specifying how long a script is allowed to execute before the Cohesity cluster stops the execution of the script. By default, the timeout is 15 minutes. If no value is specified, there is no timeout and the script is not stopped. This can cause a protection run to hang if the script never completes. The timeout is applied from the moment the script starts. For example, if the timeout is set to 15 minutes but the script takes 20 minutes to complete, it will be stopped after 15 minutes. Therefore, users should set the timeout based on the expected runtime of the script to avoid premature termination.

More about failures, cancellations and logging:

- If the backup is successful but the Pre-Script or Post-Script fails, a warning is issued with the script error. A script failure occurs for any of the following conditions:
  - script is not found in the expected location
  - script cannot be executed
  - script returns a failure return code of non-zero

  A script is successful when the return code is 0.
- If a Post-Script is configured, the script always executes even if the backup fails.
- If a protection run is canceled, all subsequent phases are canceled.
  - For example, if a protection run is canceled while the Pre-Script is running, the Pre-Script is canceled, then the backup task is canceled and then the Post-Script is canceled.
  - For example, if a protection run is canceled while the backup task is running, the backup task is canceled and then the Post-Script is canceled.

- The logging information reported by the script is available from the Protection Group details page. To open the page, log in to the Cohesity Dashboard, select **Data Protection > Protection**, and under **Last Run** click on a date. In the Run Details page, click the green > icon next to the protected object (entity) to view the job details.

Available Environment Variables

The Cohesity cluster populates the following environment variables with values before the Pre-Script or the Post-Script is run. You can use these environment variables in your script.

Variable Name	Available from Script	Description
`COHESITY_JOB_ID`	Pre-Script and Post-Script	Specifies the Cohesity id assigned to the Protection Group.
`COHESITY_JOB_NAME`	Pre-Script and Post-Script	Specifies the name of the Protection Group as displayed in the Cohesity Dashboard.
`COHESITY_BACKUP_TYPE`	Pre-Script and Post-Script	`kRegular` indicates an incremental (CBT) backup. Incremental backups utilizing CBT (if supported) are captured of the target protection objects. The first run of a kRegular schedule captures all the blocks. `kFull` indicates a full (no CBT) backup. A complete backup (all blocks) of the target protection objects are always captured and Change Block Tracking (CBT) is not utilized. `kLog` indicates a database log backup. Capture the database transaction logs to allow rolling back to a specific point in time. `kSystem` indicates a system backup. System backups are used to do bare metal recovery of the system to a specific point in time. (*Applies only to IBM*) `kStorageArraySnapshot` indicates storage array snapshot backup. This backup method creates point-in-time snapshots of data on a storage array, allowing for data protection and recovery without impacting the primary storage's performance.
`COHESITY_BACKUP_ENTITY`	Pre-Script and Post-Script	Specifies the name of the object (or entity) being backed up that triggered the script execution.

Variable Name	Available from Script	Description
COHESITY_BACKUP_HOST_ENTITY	Pre-Script and Post-Script	Specifies the name of the host that contains the object (or entity) being backed up that triggered the script execution.
COHESITY_CLUSTER_IP_ADDRESSES	Pre-Script and Post-Script	Specifies a comma separated list of IP Addresses of the nodes that make up the Cohesity cluster.
COHESITY_JOB_STATUS	Post-Script	SUCCEEDED indicates the protection run ran successfully. ERRORED indicates the protection run reported an error.

17. **Skip Files on Errors:** Toggled on by default. The Protection Group continues to run even if it encounters errors on files, such as permissions errors. If files are skipped, the protection run details page indicates a warning status and provides additional information. If toggled off, the Protection Group stops when it encounters an error.
18. **Encryption**: Toggle on this option to enable in-flight encryption of data traffic between the Cohesity cluster and external NAS source while backing up NFS AND SMB Volumes. This option is disabled by default.

    > **Note:** Encrypting in-flight data will have a minor impact on the performance. To use the encryption functionality for NFS volumes, you must join the Cohesity cluster to the active directory as the Kerberos server or add the Kerberos provider in the Cohesity cluster Access Management section.

19. **File DataLock**: Toggle on this option to preserve the access time of the files and folders of the selected volumes. The access time provides the lock period of the files and folders that are applied when these files and folders are recovered to the Cohesity View. For more information on File DataLock, see NAS File DataLock.

    Select one of the following **DataLock** modes:

    - **Compliance**: In this mode, files and folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files folder. However, a root user can delete the files and folders. Only a user with Data Security role can edit the File DataLock settings in the Compliance mode.

- **Enterprise**: In this mode, files and folders cannot be modified by any user before or after the lock retention period expiration. However, a root user can delete the files and folders before or after the lock retention period. An admin user or a user with a Data Security role can edit the File DataLock settings in the Enterprise mode. This mode is useful for giving storage administrators additional control over the file.

You can further configure the File DataLock settings to:

- Set the lock period for the new files and folders created on the recovered Cohesity View

  Provide the following information to set the lock period for the new files and folders that will be created on the recovered Cohesity View. These settings are not applicable to the existing NAS source files and folders recovered on the Cohesity View.

  - **Default Lock Period**: Provide the time period during which the files should remain locked. When the time period ends, files can be deleted but not changed. By default, the files are locked forever.
  - **Autolock Files**: By enabling this option, the files are automatically locked if they are not modified in the time period you specify. By default, this option is disabled.
  - **Manually Lock Using**: Select when the files are to be locked:
    - **Read-Only**: Select this option to lock files when the file permission is set to read only, for example; by setting the Windows file properties attribute or using the Linux chmod command. By default, the files are locked when the file permission is set to read only.
    - **FutureATime**: Select this option to lock files when the last access time (atime) is set, for example, by using Windows PowerShell commands or the Linux touch command.

      Optionally you can set a Minimum Retention and/or a Maximum Retention period to ensure that a file's last access time cannot be set outside this period. A request to do so results in a permission denied error if the manual locking mode is FutureATime. Otherwise, the request will be silently rejected and the file will be locked for the minimum retention duration.
- Override the locking period of the existing files and folders recovered on the Cohesity View:

  The following option can be applied to all the existing files and folders that will be recovered on the Cohesity View.

- **Override (Lock all files until a specific date)**: Enable this option and select a date to override the locking period of the existing files and folders.

  When you recover the files to the Cohesity View, Cohesity applies the farthest lock period between the override lock period you specify and the source file's lock period.

  For example;

  - The lock period of the source file is 09/09/2023 and the override lock period you specified is 09/08/2022. When you recover the file to the Cohesity View, then the lock period of the source file is applied to the recovered file, which is 09/09/2023.
  - The lock period of the source file is 09/09/2022 and the override lock period you specified is 09/08/2024. When you recover the file to the Cohesity View, then the override lock period you specified is applied to the recovered file, which is 09/08/2024.

  In **Compliance** mode, a user with Data Security privileges can extend the override date into the future, but no user can shorten the date.

  In **Enterprise** mode, a user with admin or Data Security privileges can extend or shorten the override date.

> **Note:**
> - You cannot enable or disable File DataLock once the Protection Group is created.
> - If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed.
> - You can delete Protection Groups and Snapshots having File DataLock enabled.
>
>   File DataLock also preserves the ATimes of the hardlinks and symlinks.

20. **Filter IPs**: Enable this option to filter the IP addresses of the NAS source. By filtering IP addresses, you can restrict the communication of the Cohesity cluster to specific IP addresses or subnets of the NAS source.

    To filter the IP addresses of the NAS source, select one of the following options:

    - **Allow IPs**: Select this option and specify the IP addresses of the NAS source through which the communication to the Cohesity cluster must happen. You can

provide the IP addresses in a comma-separated list or in a CIDR format.

- **Deny IPs**: Select this option and specify the IP addresses of the NAS source through which the communication to the Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.

> **Note:**
> - This option does not apply to the NAS Generic Mount Point Protection Group creation.
> - If you have specified the IP filtering at the source-level also, then the IP filtering specified at the Protection Group-level will take precedence over the IP filtering specified at the source-level.

21. **Exclusions and Inclusions:** Everything is included by default. Toggle on **Exclusions and Inclusions** if you want to exclude or include locations. By creating exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories and therefore minimize the disk space used to store the data.

> **Note:**
> To add an exclusion or inclusion, you must prefix a forward slash ('/') or suffix an asterisk ('*') to the path or to a particular file within the protected object. For example, '**/test**' or '***.txt**'.

Add Inclusion: Click to include a particular path or a particular file within the protected object.

For example, consider four directories - **test**, **test1**, **test2**, and **test3** under the protected object with path `/ifs/TestShare1/Folder1`. The table below lists the input types for an inclusion list of folders and their respective outcome:

Input	Outcome
/	Everything inside the protection path is backed up, if exclusions are not defined.

Input	Outcome
/test	When the inclusion path is not suffixed by '/', Protection Group will include all folders starting from the protection path. In this example, /test will behave as /test* and Protection Group will include '**test**', '**test1**', '**test2**', and '**test3**' from protection path for backup.
/test/	When inclusion path is suffixed by '/' , it will only include specific folder. In this example, /test/ behaves as /test/* and Protection Group will only include '**test**' from protection path for backup.

> **Note:**
> Do not provide the folder name along with the backup path while including files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not include any folder.
>
> If exclusions are defined, they will take precedence over inclusions.

Add Exclusion: Click to exclude a particular path or a particular file within the protected object. You can also specify regular expressions for excluding files in the format `regex:<regex pattern>`

For example, consider there are three directories-**Test**, **Test1**, and **Folder1** and a file, **TestFile.tmp** under the protected object with path **/ifs/TestShare1**.

- **Test1** folder has **TestFile1.tmp**, **TestFile2.tmp** files.
- **Folder1** folder has **TestFile3.tmp**, **TestFile4.tmp files**, and **SubFolder1** which further has a **Test** folder with **File.txt**, **File1.txt** files in it.

The following figure illustrates the directory structure under the protected object with path **/ifs/TestShare1**.

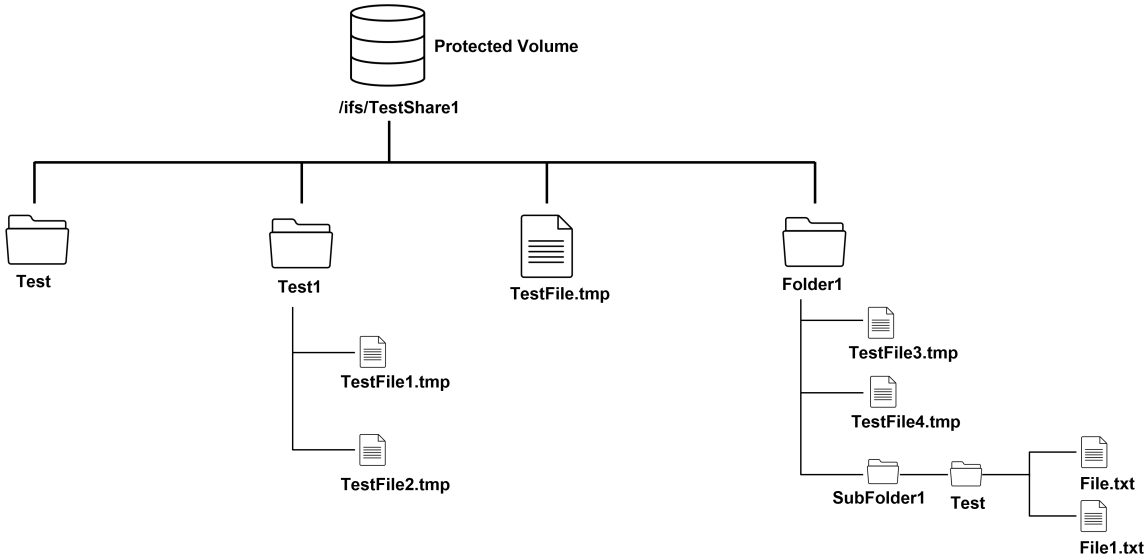

The table below lists the input types for an exclusion list of folders and their respective outcome:

Input	Description	Outcome
/Test	When the exclusion path is not suffixed by a '/', the Protection Group will exclude all files and folders with **Test** as the prefix . For exclusions, **/Test** behaves as **/Test***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes **Test** and **Test1** folders, and **TestFile1.tmp** file from the protection path. **TestFile3.tmp** and **TestFile4.tmp** files in the **Folder1** folder will not be excluded.
/Test/	When the exclusion path is suffixed by '/' , the Protection Group excludes only the files and folders with the name **Test**. For exclusions, **/Test/** behaves as **/Test/***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes the **Test** folder of the parent directory. The **Test** sub-folder in **/Folder1/SubFolder1** will not be excluded.

Input	Description	Outcome
*/Test	Excludes all the files and directories with the name '**Test**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes the **Test** folder of the parent directory and the **Test** sub-folder in **/Folder1/SubFolder1**.
*.tmp	Excludes all the files that with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.
regex:.*\.tmp	Excludes the files and folders with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.
regex:/Folder1/.*/Test/.*\.txt$	Excludes all files with **.txt** extension from all the **Test** folders directory in **/Folder1/**.	Excludes **File.txt** and **File1.txt** files in the **/Folder1/SubFolder/Test** folder.

> **Note:** Do not provide the folder name along with the backup path while excluding files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not exclude any folder.

Cohesity automatically excludes the following NetApp system files:

- .vtoc_internal and .bplusvtoc_internal files
- .copy-offload directory and .tokens file

22. **Indexing**: Toggle on this option to index the data of the Protection Group. Indexing improves files and folders search.

    Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.

    You can add custom indexing rules.

By default, a rule to include the root directory (`/`) is defined and therefore by default the entire file system is indexed. When a directory is indexed, all of the child sub-directories are also indexed. (Indexing is recursive.)

By creating exclusion and inclusion rules, you can limit indexing to a specific set of files and directories and therefore minimizing the disk space used to store the index. For example, you could add two additional rules to the first default rule:

- include `/`
- exclude `/var`
- include `/var/log`

In this example, all files in the `/` root directory are indexed except the files found in the `/var` directory. However, the files in the `/var/log` are indexed because there is an explicit include rule.

23. **Quiet Time:** This option is available only if the selected policy has at least one quiet time period. Toggle the **Stop in-progress runs when quiet time starts** option to pause or abort the runs during quiet time. If this option is toggled off, after a protection run starts, execution continues even when the quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period.
    - **Pause Runs** - Select this option to pause the currently executing protection runs and resume the runs when quiet time ends.
    - **Abort Runs** - Select this option to abort the currently executing protection runs when the quiet time period specified for the Protection Group starts.

    > **Note:** If the quiet time window for multiple protection runs is scheduled at the same time, then when the protection run is resumed, prioritization of the protection run does not work

24. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
    - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
    - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
    - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

    For more information, see Alerts.

25. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
26. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

    > **Note:**
    > The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

27. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
28. **Description:** Optional. Click the plus sign and enter a description for the Protection Group.
29. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

Manage the NAS Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirement.

> **Note:** You cannot update the registered NAS server and the assigned Storage Domain in the Protection Group settings.

Add NAS to the Protection Group

You can add new NAS volumes to the existing Protection Group for backup. The newly added NAS volumes are backed up from the next protection run and it will be a full backup.

Perform the following steps to add NAS volumes to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the edit icon.
5. Select the NAS volume you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added NAS volumes are backed in the next scheduled protection run.

Remove NAS from the Protection Group

From the Edit Protection page, you can remove NAS volumes from the Protection Group that you do not want to back up in the next protection run.

Perform the following steps to remove NAS volumes from the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the edit icon.
5. Unselect the NAS volume you want to remove from the Protection Group.
6. Click **Save Selection** > **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new NAS volumes that are added to the selected parent NAS server. However, you cannot enable auto-protect at the server-level containing both Read-Write volumes and Data-Protect volumes.

Perform the following steps to enable auto protect:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the edit icon.

5. Click the shield icon to enable Auto Protect.
6. Click **Save Selection** > **Save**.

When this option is enabled on the parent NAS server, all the NAS volumes that are added to the parent NAS server in the future are automatically added to the Protection Group and are protected from the next protection run.

For example, you create a Protection Group that backs up a NAS. The Protection Group runs for two weeks and then a new volume is added to the NAS. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new volume is also backed up even though the new volume has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list, and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Excluded objects that are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirement.

Perform the following steps to edit the policy settings:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Policy** field, click the **Edit Policy** icon .

   The **Edit Policy** page appears
5. Edit the settings and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to the Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

Perform the following steps to attach a new protection policy to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, from the **Policy** drop-down, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to back up NAS volumes. It provides you options to back up all the objects, only selected objects, or failed objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- Is configured with replication targets, you can select those targets.
- Is configured with archival targets, you can select those targets.

Perform the following steps to start an on-demand protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. From the **Backup** drop-down list, select any one of the following:
        - **All objects in the Protection Group**.
        - **All objects that failed last run**.
        - **All objects that succeeded last run**.
        - **Selected objects in the protection run**.
    2. In the **Backup Type** drop-down list, select the backup type.
    3. Optional. In the **Replicate** section, select the replication target and the retention duration or click **Add** to add the replication target and the retention duration. Also, select **Replicate only fully successful runs** to replicate the data only from the completed protection runs.
    4. Optional. In the **Archive** section, click **Add** to add the external target to which you want to archive the data and the retention duration. Also, select **Archive only fully successful runs** checkbox to archive the data only from the completed protection runs.
5. Click **Backup Now**.

Pause Protection Runs Instantly

You can instantly pause a protection run of the Protection Group that you have created to back up the NAS sources.

Perform the following steps to pause a currently executing protection run of a Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, click the NAS Protection Group that has been created.
3. Click the Actions menu (⋮) next to the protection run and click **Pause Run**.

   The **Pause Run?** window appears, allowing you to add a note for reference.
4. Click **Confirm**.

   You can see who paused the run and view the note associated with the paused run.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

Perform the following steps to pause the future protection runs:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.
   The **Pause Future Runs** window appears, allowing you to add a note for reference.
4. Click **Confirm**.

   You can see who paused the Protection Group and view the note associated with the paused Protection Group.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered

Perform the following steps to resume paused protection runs:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you paused.
3. Click the Actions menu (⋮) next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Protection Group are not deleted.
- **Delete Protection Group and Snapshots**: Delete the Protection Group and all snapshots, any local snapshots captured by this Protection Group and stored on this cluster are deleted. In addition, any archived snapshots stored on external targets that were captured by this Protection Group are also deleted.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

Perform the following steps to delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for NAS.
3. Click the Actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Run.

If the Protection Group is replicating to a remote cluster, the failover job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the failover job.

Download Error Report for a NAS Protection Group

You can now download the list of all the errors encountered during the protection run for a NAS backup. Once the protection run is complete, you can download the details of the files that are skipped or have encountered errors in a CSV format.

Follow the steps below to download the errors report for a Protection Group.

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. In the **Protection** page, select a protection group that has the warning icon (ICON) in the **Status** column.
3. In the backup run details page, click the **Warning** icon .
4. In the **Warning** page, click **View More** in the **Message** column to view a list of errors for the Protection Group run and then click **Download Error Logs**.

This will download all the error logs in a CSV file format.

Re-sync Protection Groups

Re-sync refers to manually triggering an archive task and archiving the snapshot again, for example, if an archive has failed earlier. The re-sync operation can be performed only if the local copy has not expired.

Follow the steps below to perform the re-sync operation.

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. In the **Protection** page, select the required Protection Group.
3. In the backup run details page, click the action icon (⋮) for the run and click **Edit Run**.
4. In the **Edit Run** page, click **Re-sync** under the **Archival** section.

   > **Note:** This option is displayed if the protection run includes an associated archive run.

5. Click **Update Retention**, provide the retention duration, and click **Save**. Re-sync will be performed and all the objects will be re-archived.

NAS File DataLock

Many organizations now rely on WORM (Write Once Read Many) data storage such as Nutanix Files, NetApp Snaplock, and Isilon Smartlock volumes to retain their files and folders in unmodified form for compliance and regulatory purposes and to protect the files from accidental deletion or modification.

When you backup these WORM-capable volumes to the Cohesity cluster, it is essential to have the ability to easily lock the files and folders of the volumes while protecting themselves from accidental or malicious attempts to delete or modify these files. The backup should also provide a strong guarantee for data non-mutability and non-erasability for a certain period of time.

This capability is achieved using Cohesity's File DataLock feature.

FileDataLock is a Cohesity feature that preserves the WORM attributes of any files and folders created in the Cohesity View. You can apply the File DataLock feature for NAS backup (NFS exports and SMB shares) while creating a NAS Protection Group.

With File DataLock enabled in a NAS Protection Group, Cohesity preserves the future AccessTime (ATime) of the files and folders. The AccessTime provides the lock period for files and folders. You can also override this lock period in the File DataLock settings of the NAS Protection Group. This lock period is applied to these files and folders when they are recovered to the Cohesity View. This prevents the users or administrator from modifying or deleting these locked files until the locking period expires.

You can also set lock periods for the new files and folders created on recovered Cohesity View in the File DataLock settings of the NAS Protection Group.

When you recover File DataLock enabled storage volumes to the Cohesity View, the View details page will display a lock icon for the files and folders having File DataLock.

File DataLock is supported for all types of NAS sources including NetApp 7-mode and is agnostic of the NAS protocol.

For information on how to enable File DataLock, see Protect NAS.

File DataLock Modes

You can use the File DataLock feature in one of the two modes:

- **Compliance mode**: The WORM files are protected at the disk level. This mode implements strict regulatory requirements for data retention and complies with SEC Rule 17a-4. In this mode, files and folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files and folders. However, a root user can delete the files and folders. Only a user with a Data Security role can edit the File DataLock settings in the Compliance mode.
- **Enterprise mode**: The WORM files are protected at the file level. This mode does not implement strict regulatory requirements for data retention. In this mode, files and folders cannot be modified by any user before or after the lock retention period expiration. However, a root user can delete the files and folders during or after the lock retention period. An admin user or a user with the Data Security role can edit the File DataLock settings in the Enterprise mode

> **Note:** An admin user or a user with the Data Security role can edit the File DataLock settings and switch from the Enterprise mode to Compliance mode. But switching from Compliance mode to Enterprise mode is not supported.

Things to remember for File DataLock

- You cannot enable or disable File DataLock once the View is created.
- If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed.
- You can delete Protection Groups and snapshots having File DataLock enabled.

  File DataLock also preserves the ATimes of the hardlinks and symlinks.
- **Recovery**
  - Recovering a File DataLock enabled backed-up volume as a Cohesity View will preserve ATimes and the locking period for the file.
  - Recovering a File DataLock enabled volume to the original location with 'Overwrite existing files/folder' as the recovery option will encounter an error if the files and folders are read-only. However, the ATime and other attributes of other files and folders of the source volumes are preserved.
  - Recovering a File DataLock enabled volume to an alternate NAS recovery will preserve ATimes and other attributes of files and folders of the source volumes.
  - Cloning a recovered View having File DataLock enabled will not inherit the File DataLock properties.
- **Replication**

  When you replicate to a remote Cohesity cluster, the File DataLock settings are also replicated.
- **Archival**

  File DataLock settings are preserved when you recover volume from an archived snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.
- **Cloud Archive Direct**

  File DataLock settings are preserved when you recover from a Cloud ArchiveDirect snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.

NetApp ONTAP

NetApp ONTAP is a hardware and software-based data storage and retrieval system. It responds to network requests from clients and fulfills them by writing data to or retrieving data from its disk array.

You can backup and recovery NetApp ONTAP volumes using the following two backup approaches:

- **Cohesity File-Runner based backup approach**: Cohesity uses the Cohesity's built-in CFT method to protect data. In this method, Cohesity uses its internal File-Runner to discover the delta between two snapshot copies.
- **NetApp SnapDiff-based backup approach**: Cohesity leverages NetApp's SnapMirror cloud functionality to move the data to the Cohesity cluster and NetApp's SnapDiff functionality to get the changed file list.

> **Note:**
> - In this document, NetApp SnapDiff means Cohesity's implementation of SnapDiff using NetApp SnapDiff v3 APIs.

Supplemental terms regarding the Cohesity implementation of the SnapDiff API. Please review before use.

NetApp ONTAP File-Runner Based Backup

In the Cohesity File-Runner based backup approach for NetApp ONTAP, Cohesity uses the Cohesity's built-in CFT method to protect data. In this method, Cohesity uses its internal File-Runner to discover the delta between two snapshot copies.

Plan and Prepare for File-Runner Based NetApp ONTAP Backup

This topic provides information on the supported NetApp ONTAP versions and volumes, minimum permissions, and the prerequisites to register and protect NetApp ONTAP with Cohesity cluster. For more information, see Register NetApp ONTAP.

Before you register your NetApp ONTAP with Cohesity cluster and protect the volumes and data, ensure the following prerequisites:

- The Storage Virtual Machines (SVMs) that you plan to protect must have:
  - an active logical interface attached to the SVM.
  - the NFS and CIFS services configured on the SVM.
- The Cohesity cluster requires access to the snapshot directory on ONTAP. This access is enabled by default. However, it is recommended to verify that it is enabled in your environment.
  1. To check the access status of the snapshot directory, run the following command:

     ```
     volume show -vserver -volume -fields snapdir-access
     ```

     Example output, where volume name is vol1:

```
clus1::> volume show -vserver vs0 -volume vol1 -fields snapdir-
access
```

vserver	volume	snapdir-access
vs0	vol1	true

2. If the snapshot directory access is False, enable it using the following command:

   ```
   volume modify -vserver <SVM_name> -volume <vol_name> -snapdir-
   access true
   ```

   Example output, where volume name is vol1:

   ```
   clus1::> volume modify -vserver vs0 -volume vol1 -snapdir-access
   true
   ```

   Volume modify successful on volume vol1 of Vserver vs0.

For more information about performing these changes through the ONTAP System Manager console, see the NetApp documentation.

- The **Allow Superuser Access** check box is selected in export policy to register SVM with the Cohesity cluster.
- Read or write or full control permissions are added under Share Access Control for the registered SMB user,
  for the volume share being backed up using the SMB protocol.
- The required firewall ports are open to enable communication between Cohesity cluster and NetApp ONTAP.
- If volume export policies are not in use, but export policies are applied at the Qtree level, Cohesity IP addresses must be manually added to the Qtree-level export policies to ensure proper access.
- Review and understand the considerations and best practices.

Support Matrix

Before you register your NetApp ONTAP with Cohesity cluster, ensure that Cohesity cluster supports the NetApp ONTAP versions and volumes you want to protect. For more information, see Supported Software

Supported NetApp ONTAP Volumes

The following table lists the supported NetApp ONTAP versions and volume types for backup:

Volume Type	Volume Sub-Type
Flex Volume*	Normal Flex Volume
	SnapLock Enterprise Volume
	SnapLock Compliance Volume
	Encrypted Volume Storage
Data Protection Volume	SnapMirror Destination Volume
	SnapVault Destination Volume

- Flex Volume support also includes FlexGroup Volume support.
- Cohesity does not support the backup of the following NetApp ONTAP volumes:
  - Flex Volume sub-type of Flex Volume

Supported NFS and SMB versions

The following table lists the supported NFS and SMB versions for backup:

Protocol	Version	Notes
NFS	NFSv3, NFSv4.1	• If NFSv4 volume backup is triggered, the Cohesity cluster performs the backup in NFSv3 mode. • NFS v4.1 Kerberos authentication is supported.
SMB	SMB v2.x, v3	

Supported Security Style of Qtree with Root Volume's Security Style

You can back up Qtrees even if their security style does not match the root-volume security style.

Volume Type	Volume Sub-Type	Qtree	
		Supported	Not Supported*
NTFS	Indicates that the file security style of the qtree is UNIX based.	NTFS	UNIX, MIXED
UNIX	Indicates that the file security style of the qtree is Windows based.	NTFS, UNIX	MIXED
MIXED	Indicates that either the UNIX or the NTFS file security style could be the effective security style for each file or folder within the qtree.	NTFS, UNIX	MIXED

*Cohesity skips the backup of Qtree security styles that are not supported.

Minimum Permissions

Ensure you have the required minimum permissions to register NetApp ONTAP as a source with the Cohesity cluster, and to perform backup and recovery.

Minimum Permissions for Registering NetApp ONTAP

The user account used for registering NetApp ONTAP SVM or NetApp ONTAP Cluster must have the required permissions to communicate with the Cohesity cluster.

The following user-role types are used for File-Runner-based NetApp ONTAP Backup:

- ONTAPI (Legacy)
- HTTP

You can assign a role to the NetApp user which specifies permission for the user. The role type can be either REST or Zephyr Application Programming Interface (ZAPI) permissions role. You can configure both these roles from the NetApp UI.

- **REST role**: Using the REST permission role, you can specify permissions using the NetApp REST APIs. You can use the REST role to specify permissions for both ZAPI and REST APIs.
- **ZAPI ONTAPI role (Legacy)**: Using the ZAPI permission role, you can specify permissions using the NetApp CLI commands or ZAPI calls. It is a legacy role and is used for ZAPI API permissions only.

Configuring Permissions for the REST Role

To configure the permissions using REST API, use the guidelines described in this section.

Considerations:

- To use only REST APIs for backup and recovery, you must create a REST role for configuring the permissions.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

The following tables describe the REST APIs.

For Cluster, Licensing, and Jobs

API	Access Level	Used By	Description
/api/cluster	Read-only	• GET /api/cluster • GET /api/cluster/nodes • GET /api/cluster/jobs/{uuid}	• Discovers cluster identity and version. • Polls job state for various asynchronous operations.
/api/cluster/licensing **Note:** The secondary path for this API is: /licensing	Read-only	GET /api/cluster/licensing/licenses	Inspects cluster licenses.

For SVM and Network

API	Access Level	Used By	Description
/api/svm/svms	Read-only	GET /api/svm/svms	Lists SVMs along with their supported data protocols within the process of entity-hierarchy discovery.
/api/network/ip/interfaces	Read-only	GET /api/network/ip/interfaces	Discovers LIFs, addresses, and services to determine usable NFS/CIFS/management interfaces.

For Protocol and Security Configuration

API	Access Level	Used By	Description
/api/protocols/nfs/services	Read-only	GET /api/protocols/nfs/services/{svm_uuid}	Determines NFS enablement (v3/v4.1) per SVM.
/api/protocols/nfs/kerberos/interfaces	Read-only	GET /api/protocols/nfs/kerberos/interfaces	Maps NFS Kerberos-enabled IPs and hostnames.
/api/protocols/cifs/services	Read-only	GET /api/protocols/cifs/services/{svm_uuid}	Retrieves CIFS server name and AD domain/FQDN.
/api/protocols/cifs/shares	Read-only	GET /api/protocols/cifs/shares	Lists CIFS shares, paths, and ACLs for hierarchy and access checks.
/api/protocols/cifs/users-and-groups/privileges	Read/Write	POST /api/protocols/cifs/users-and-groups/privileges	Adds or updates CIFS user/group privileges when NAS workflows require it.
/api/security/certificates	Read-only	GET /api/security/certificates	Inspects SVM SSL certificate configuration and server-authentication state.

For Volumes, Qtrees, and Export Policies

API	Access Level	Used By	Description
/api/storage/volumes	Read-only	• GET /api/storage/volumes • GET /api/storage/volumes/{uuid}	Discovers NetApp volumes (capacity, junction path, export policy, security style, and related attributes) for entity hierarchy and NAS backup targeting.
/api/storage/qtrees	Read/Write	• GET /api/storage/qtrees • POST /api/storage/qtrees	• Reads qtree security-style and exports policy information. • Creates qtrees as part of NAS export-related operations.
/api/protocols/nfs/export-policies	Read/Write	• GET/POST /api/protocols/nfs/export-policies • GET /api/protocols/nfs/export-policies/{id}/rules	Lists, creates, and modifies NFS export policies and rules that control NAS backup and access behavior.

For Snapshots (NAS backup and recovery)

API	Access Level	Used By	Description
/api/storage/volumes/*/snapshots **Note:** The secondary path for this API is: /*/snapshots	Read/Write	• POST /api/storage/volumes/{uuid}/snapshots • GET /api/storage/volumes/{uuid}/snapshots • GET /api/storage/volumes/{uuid}/snapshots/{snapshot_uuid} • PATCH /api/storage/volumes/{uuid}/snapshots/{snapshot_uuid} • DELETE /api/storage/volumes/{uuid}/snapshots/{snapshot_uuid}	Creates, lists, modifies (rename or change expiry), and deletes volume snapshots for NAS backups and recoveries. In some ONTAP versions, RBAC is purely prefix-based. In these versions, configuring api/storage/volumes with Read/Write access is sufficient to cover snapshots. If the UI allows the more specific /api/storage/volumes/*/snapshots entry, this document assumes it is granted Read/Write access.

For Private CLI Endpoints

API	Access Level	Used By	Description
/api/private/cli/snapshot **Note:** Secondary path: /private/cli/snapshot	Read/Write	GET /api/private/cli/snapshot?snapshot=...&volume=...&fields=instance-uuid	Retrieves CLI snapshot metadata like instance UUID as needed for NAS workflows.

API	Access Level	Used By	Description
/api/private/cli/vserver/export-policy/check-access **Note:** Secondary path: /private/cli/vserver/export-policy/check-access	Read-only	GET /api/private/cli/vserver/export-policy/check-access	Uses the private CLI layer to conduct export-policy access checks.

Considerations

- NAS entity hierarchy discovery is fully supported.
- NAS snapshot based backup and recovery (create, delete, and modify) are supported.
- List-based backup and recovery (create, delete, modify, and list snapshots) are supported.
- NFS/CIFS export policy and privilege management required by NAS workflows are supported.
- SnapDiff, SnapMirror policies and relationships, object store endpoints, cloud targets, VMware store endpoints, cloud targets, VMware storage snapshot workflows, and array snapshot workflows remain disabled at the RBAC-level for this role.

Configuring Permissions Using Legacy Commands

Before registering a NetApp ONTAP Cluster as the source type, ensure you have the following command permissions:

Access-level	Command	Description
All	vserver export-policy	Adds the Cohesity cluster subnet to the export-policy so that Cohesity can mount volumes.
	volume snapshot	Allows fetching, creating, and deleting snapshots for volumes.

Access-level	Command	Description
ReadOnly	vserver cifs	Fetches information about CIFS/SMB shares for volumes .
	cluster identity	Fetches information about the cluster.
	network interface	Fetches information about network interfaces that Cohesity connects to for mounting volumes.
	volume	Fetches information about volumes.
	vserver	Fetches information about SVM.

When registering a NetApp ONTAP SVM as the source type, ensure you have the following command permissions:

Access-level	Command	Description
All	vserver export-policy	Adds the Cohesity subnet to the export-policy so that Cohesity cluster can mount volumes.
	volume snapshot	Allows fetching, creating, and deleting snapshots for volumes.
ReadOnly	vserver cifs	Fetches information about CIFS/SMB shares for volumes.
	network interface	Fetches information about network interfaces to which Cohesity cluster connects for mounting volumes.
	volume	Fetches information about volumes.
	vserver	Fetches information about SVM.

Minimum Permissions for SMB/CIFS Shares Backup and Recovery

To backup NetApp ONTAP SMB/CIFS shares to the Cohesity cluster, the user must have local or domain user credentials that allow at least read access to the SMB share.

To recover the SMB/CIFS shares from the Cohesity cluster, the local or domain user must have full access control on the target where data is being restored.

You can assign the following privileges to the user:

- **SeBackupPrivilege** to the SMB user to be able to back up the SMB data without modifying the ACL on the SMB shares.
- **SeRestorePrivilege** to restore files and directories, override any ACLs, and set a valid user or group SID as the file owner.
- **SeChangeNotifyPrivilege** to bypass traverse checking. Users with this privilege are not required to have traverse (x) permissions to traverse folders, symlinks, or junctions.

Minimum Permissions for NFS Export Backup and Recovery

To backup NFS export, the user must have the read and superuser access permissions on the NFS volume that needs to be backed up and on the parent root volume. Before starting the backup, Cohesity verifies if the user has these permissions, if not, the Cohesity cluster adds a new rule index for the Cohesity node subnet in the export policy attached to the backup volume. You must add the permissions manually for the Cohesity node subnet for the parent root volume.

To restore NFS export, the user needs the read/write and superuser access permissions on the NFS volume that needs to be restored. Before starting the restore, you must add a new rule index for the Cohesity node subnet in the export policy attached to the source volume and parent root volume to give permissions to the Cohesity node subnet/IPs.

Permissions for NTFS Trees on UNIX Security Style Volumes

When you protect NetApp volumes that use a UNIX security style and contain NTFS trees by using File Runner–based backups, the SMB service must have **Read** access to the CIFS *c$* share on the SVM. This access is required because the File Runner workflow accesses the SVM root context during entity discovery.

For restore scenarios that involve full volume recreation or operations that require **Write** access at the SVM root level, the SMB account must have **Full Control** permissions on the *c$* share.

Credentials for NetApp ONTAP Backup with Multiple SVMs

To register NetApp ONTAP with multiple storage virtual machine (SVM), create a custom role with the required permissions and a local user at the SVM level. Assign the custom role to the local user. Use the respective local user account to register multiple SVMs.

Ports Used for Communication

For information on port numbers required, see Manage Firewall Ports.

Best Practices

Cohesity protects NetApp ONTAP volumes by communicating with a NetApp ONTAP cluster or Storage Virtual Machine (SVM).

As a security best practice, you must enable SSL before registering the cluster or SVM as a source on the Cohesity cluster.

To enable SSL, run the following ONTAP CLI commands on your NetApp ONTAP cluster:

1. To verify the settings, run the following command:

```
netapp_cluster::> security ssl show -vserver <cluster/vservername>
```

2. To enable SSL, run the following command:

```
netapp_cluster::> security ssl modify -vserver <cluster/vservername> -
server-enabled true
```

3. After enabling SSL, run the following command to verify the settings:

```
netapp_cluster::> security ssl show -vserver <cluster/vservername>.
```

Considerations

Review and understand the following considerations before you protect your NetApp ONTAP data:

- The NAS storage FQDN must be resolvable. Even if registered by IP, SMB restore workflows depend on the FQDN for repeated SMB share connections. Ensure that the DNS or cluster host mapping resolves the storage FQDN.
- Instant Volume Mount for NetApp ONTAP stub file is not supported.
- You cannot restore the NetApp Data-Protect volume to the original location or to an alternate Data-Protect volume because the Data-Protect volume is a read-only volume.
- To backup NetApp DP volumes with SMB share, the volume needs to be in the junction path on the NetApp.
- Cohesity skips the backup of the following security style Qtrees:
  - Mixed-mode security style Qtrees
  - UNIX security style Qtree when root volume security style is NTFS.
- **Volume recovery**: The restored volume (target volume) should have the same security style as that of the volume being protected.

  When volumes with Qtrees are recovered and the Qtrees have a different export policy than the parent volume, update the export policy on the recovered Qtrees to match the original Qtrees.
- If volume export policies are not in use, but export policies are applied at the Qtree level, Cohesity IP addresses must be manually added to the Qtree-level export policies to ensure proper access.
- **Recovering files/folders with Qtrees**: The security style of the location (target) to which the files or folders are restored should have the same security style as the restored files and folders.

- Cohesity supports the backup of the symbolic link reparse point, **IO_REPARSE_TAG_SYMLINK** in an SMB share.
- Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.

The following are the considerations for NetApp NFSv4.1 volumes:

- Encryption is not supported.
- Kerberos is not supported.
- If the NFS backup preference selected in the Cohesity UI is NFSv4.1, and if the NetApp Vserver has NFSv4.1 disabled, all the volumes belonging to this Vserver will use the NFSv3 protocol for backup.
- Recovering an NFS 4.1 volume to NFSv3 is not supported.
- Recovering an NFS 4.1 volume to the Cohesity View is not supported.

Register and Manage NetApp ONTAP Source

To protect your NetApp ONTAP data, you need to register your NetApp ONTAP as a source on the Cohesity cluster. After registration, you can manage the source on the Cohesity cluster in the following ways:

- Update the NetApp ONTAP source configuration
- Refresh the source details
- Unregister the NetApp ONTAP source from the Cohesity cluster.

Register NetApp ONTAP as a Source

Before you register NetApp ONTAP as a source on the Cohesity cluster, ensure that you:

1. Review the requirements, considerations, and best practices.
2. Perform the prerequisite tasks.

For more information, see Plan and Prepare for File-Runner Based NetApp ONTAP Backup.

> **Note:** Before you register NetApp ONTAP as a source on the Cohesity cluster, ensure that you have reviewed and completed the instructions provided in the Plan and Prepare for File-Runner Based NetApp ONTAP Backup

To register NetApp ONTAP as a source:

> **Note:** To register the NetApp ONTAP cluster for NetApp SnapDiff-based backup, you must obtain the SnapMirror Cloud license from NetApp.

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click **Register** on the top-right of the page and then select **NAS**.
   The Register NAS page appears.
3. From the NAS Source drop-down list, select **NetApp**.
4. Enter the following information, and then click **Register**.

Field Name	Description
NetApp Source Type	The NetApp ONTAP cluster or Vserver/SVM to register. **Note:** Cohesity does not support SnapDiff-based backup if the NetApp source is registered as NetApp ONTAP SVM.
Host or IP Address	The Host or IP address of the NetApp ONTAP cluster or SVM to register. For NetApp SnapDiff-based backup, use the Host or IP address of the NetApp ONTAP cluster to register.
Username	Username used to access the NetApp ONTAP cluster or SVM. Specify a user that has adequate privileges to perform actions on the Source. If you are providing a domain user account, then you must specify the user name in the following format: `fully_qualified_domain_name\username`. For example, `xydc.local\user1`. For more information, see Minimum Permissions. **Note:** The username used to register the NetApp ONTAP cluster and SVM is case sensitive.
Password	Enter the password for the specified user.
Back Up SMB Volumes	Whether NetApp ONTAP will back up SMB volumes. If yes, provide the local or Active Directory user credentials that allow at least read access on the NetApp ONTAP cluster or SVM. **Note:** The user must have full control on the restore target during restore.

Field Name	Description
Filter IPs	Enable this option to filter the IP addresses of the NetApp ONTAP cluster. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the NetApp ONTAP cluster. To filter the IP addresses of the NetApp ONTAP cluster, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the NetApp ONTAP source through which the communication to the Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the NetApp ONTAP source through which the communication to the Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.
Storage Array Snapshots	Enable this option to manage the snapshots of VMware VMs stored in NetApp NFS volumes from the Cohesity cluster. • **Cap number of snapshot per volume**: Enter the maximum number of snapshots that can be stored on each NetApp storage volume. • **Cap storage consumed by snapshot per volume**: Enter the percentage of the NetApp volume's storage you want the snapshots to consume. **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

When registration finishes, a new entry appears on the Sources page. You can now create a Protection Group to backup NetApp ONTAP volumes. For more information, see Backup NetApp ONTAP Using File-Runner .

Update the NetApp ONTAP Source Configuration

You can update the NetApp ONTAP details that you had provided during the registration process with the latest NetApp ONTAP configuration.

To edit the NetApp ONTAP source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the NetApp ONTAP source and select **Edit**.
3. In the Edit NAS page, update the respective configurations.
4. In the **Password** field, enter the password of the NetApp ONTAP user account, and then click **Update**.

Refresh the NetApp ONTAP Source

You can refresh the NetApp ONTAP configuration and fetch the latest changes on the NetApp ONTAP.

To refresh the NetApp ONTAP source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the NetApp ONTAP source and select **Refresh**.

Unregister NetApp ONTAP

If you plan to stop taking backups of the NetApp ONTAP volumes, you can unregister the NetApp ONTAP source from the Cohesity cluster. Before you unregister, you must delete all the Protection Groups associated with the NetApp ONTAP server. For more information, see Delete the Protection Group.

To unregister the NetApp ONTAP server on the Cohesity cluster:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the NetApp ONTAP source and select **Unregister**.
3. Click **Unregister**.

Backup NetApp ONTAP Using File-Runner

After registering your NetApp ONTAP server as a source on the Cohesity cluster, you can create a Protection Group to backup NetApp ONTAP. You can backup one or more NetApp ONTAP NFS or SMB shares in a Protection Group while preserving ACLs on SMB.

Select NetApp ONTAP Volumes for Backup

Cohesity supports the backup of both NetApp ONTAP read-write (RW) and data protection (DP) volumes. NetApp ONTAP DP volumes are read-only volumes which are used as destination volumes for NetApp ONTAP's replication applications such as SnapMirror and SnapVault. These applications replicate data from NetApp ONTAP primary source volumes to NetApp secondary destination DP volumes. The volume can be of NFS export only, SMB share only, or of both NFS and SMB (mixed) type. Based on the volume and protocol you select for backup, you must provide additional information when creating a Protection Group. For more information, see Protect NetApp ONTAP Using File-Runner Based Backup.

> **Note:** The file-runner is supported on the NetApp Storage Virtual Machine (SVM) and Storage Virtual Machine for Disaster Recovery (SVM-DR).

The following table lists the additional information required to perform backup based on the volume type and protocol you choose while creating a Protection Group:

Volume type	Protocol	Required additional information
Data protection	NFS	Snapshot Prefix for DP Volumes
	SMB	• Snapshot Prefix for DP Volumes • User Access Credentials to Back Up SMB Shares
	SMB, NFS	• Backup Preferences For Mixed Mode Volumes • User Access Credentials to Back Up SMB Shares • Snapshot Prefix for DP Volumes
	NFS v3 & NFSv4.1	• Backup Preferences for NFS Volumes • Snapshot Prefix for DP Volumes
	SMB, NFSv3 & NFSv4.1	• Backup Preferences For Mixed Mode Volumes • User Access Credentials to Back Up SMB Shares • Snapshot Prefix for DP Volumes • Backup Preferences for NFS Volumes
Read-write	NFS	None
	SMB	User Access Credentials to Back Up SMB Shares
	SMB, NFS	• Backup Preferences For Mixed Mode Volumes • User Access Credentials to Back Up SMB Shares
	NFS v3 & NFSv4.1	Backup Preferences for NFS Volumes
	SMB, NFSv3 & NFSv4.1	• Backup Preferences For Mixed Mode Volumes • User Access Credentials to Back Up SMB Shares • Backup Preferences for NFS Volumes

Considerations

- You cannot select both RW and DP volumes in the same Protection Group. For example, in a Protection Group, if you select a RW volume, then you cannot select another volume with the DP mode.
- If you are selecting the objects to be backed up at a NetApp ONTAP cluster-level or server-level having both RW volumes and DP volumes, then RW volumes will take precedence over data protection volumes. That is, only RW volumes of the NetApp ONTAP server will be selected for the backup.

- You cannot enable auto-protect at the server-level containing both RW volumes and data protection volumes.
- The snapshot prefix setting is applicable to both DP and RW volumes. However, Cohesity will only implement these settings on DP volumes.
- To backup NetApp DP volumes with NFS/SMB share, the volume needs to be in the junction path on the NetApp.

Snapshot Prefix for DP Volumes

Cohesity cannot create snapshots of the NetApp data protection volume because it is a read-only volume. Therefore, for backing up NetApp data protection volume, Cohesity DataPlatform leverages the existing snapshots present in the DP volume that are replicated from the NetApp ONTAP source volume to the DP volume. When creating a Protection Group, you can select one of the following options and provide the required information to leverage the existing snapshots present in the DP volume:

- **None**: By selecting this option, Cohesity will take the full backup from the oldest snapshot available on the DP volume and incremental backup from the latest snapshots available on the DP volume.
- **Snapshot Prefix**: By selecting this option, Cohesity will take the full and incremental backup from the snapshots that have the prefix name you specify.
  - **Incremental snapshot label**

    The prefix of the snapshot name present in the DP volume from which the Cohesity cluster can take incremental backup. If there are multiple snapshots that match the specified name, then theCohesity cluster will take the backup from the latest snapshot. The create time of this latest snapshot should be later than the create time of the snapshot that was last used in the most recent backup.

    If the snapshot label specified in the Protection Group is not present in the DP volume, then the following error message appears when you run the Protection Group:

    ```
    No snapshot with label: <label_name> found in the volume: <DP_
    volume_name> on vserver <NAS_server_name>. Cohesity discovers
    snapshots for DataProtect volumes and use existing snapshots for
    taking backups. Please verify the labels for the snapshot
    provided.
    ```

  - **Full snapshot label**: The prefix of the snapshot name present in the DP volume from which theCohesity cluster can take the first full backup. If there are multiple snapshots that match the specified name, then the Cohesity cluster will take the backup from the latest snapshot. For example, if you specify the prefix name of snapshot label as weekly and in the DP volume there are three snapshots with labels, weekly1, weekly2, and weekly3 (latest snapshot), then

the Cohesity cluster will take the first full backup from the latest snapshot, which is weekly3.

Considerations

- When creating the SnapMirror policy to replicate the source volume to the DP volume, ensure to provide the snapshot retention period depending on the size of the volume you backup.
- Since Cohesity does not create and manage the snapshots present in the DP volume, you must ensure that the snapshot label of the DP volume you specified does not get deleted within the backup duration you configured in the protection policy.
- When assigning or creating backup policies to a Protection Group for DP volumes, it is recommended to create the schedule of backup policy in-sync or less aggressive than the SnapMirror or SnapVault schedule. This ensures the presence of the snapshot in the DP volume, prior to the incremental backup triggered by Cohesity backup policy schedule.
- When creating snapshots for DP volume backup, ensure not to use the same prefixes that are automatically generated by NetApp SnapMirror.

User Access Credentials to Back Up SMB Shares

To back up SMB shares of any volume type, you must provide the local or Active Directory user credentials that allow at least read access to the NetApp ONTAP SMB share. You can provide these credentials while registering the NetApp ONTAP server with the Cohesity cluster or later when creating the Protection Group. For more information, see, Register and Manage NetApp ONTAP Source.

Backup Preferences For Mixed Mode Volumes

If the volume you plan to backup is exposed to both NFS export and SMB share, then you must choose how you want the Cohesity cluster to backup that volume while creating the Protection Group. Cohesity will backup the ACLs according to the selected backup preference.

If you select NFS as the preference mode, then the SMB access control lists (ACLs) would be lost and only the NFS POSIX ACLs will be backed up.

If you select SMB as the preference mode, then the NFS POSIX ACLs would be lost and only the SMB ACLs will be backed up.

Backup Preferences for NFS Volumes

If you select NFS as the preference mode and the selected volume has:

- NFSv3 enabled, then the SMB ACLs will be lost and only the NFS POSIX ACLs will be backed up.
- NFSv3 and NFSv4.1 enabled, then you can choose to use NFSv4.1 to back up the volume. If you choose NFSv4.1, then the SMB ACLs will be lost, and the NFSv4.1 ACLs

will be backed up. However, if you do not toggle the "Use NFSv4.1 protocol for NFS volumes" option on the Cohesity Dashboard, then only the NFS POSIX ACLs will be backed up.

- NFSv4.1 enabled (NFSv3 disabled), the backup will be performed using NFSv4.1 automatically and the **Use NFSv4.1 protocol for NFS volumes** toggle option will not be available on the Cohesity Dashboard.

> **Note:** NetApp NFSv4.1 volumes are supported from the Cohesity 6.6.0d version onwards. To enable this feature, contact Cohesity Support.

Extended Attributes using SMB2 Protocol on NetApp ONTAP 9.15.1

Cohesity's NAS Data Protection supports the restore of extended attributes of entities using the SMB2 protocol on NetApp. You can now preserve and restore the file content along with any associated extended attributes, which are metadata that provide additional information about the .

Protect NetApp ONTAP Using File-Runner Based Backup

In Cohesity, NetApp ONTAP volumes are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

This topic provides information on how to protect NetApp ONTAP volumes using a Protection Group and a Protection Policy:

- Protect Data Protection Volumes with Only SMB Share or NFS Export
- Protect Data Protection Volumes with Both NFS and SMB Share (Mixed Mode)
- Protect Read-Write volumes with Only SMB Share or NFS Export
- Protect Read-Write Volumes with Both NFS and SMB Share (Mixed Mode)
- Protect Snapshots of VMware VMs stored in NetApp NFS volumes

> **Note:** Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.

Protect Data Protection Volumes with Only SMB Share or NFS Export

You can protect NetApp ONTAP DP volumes with only SMB share or NFS export by creating a Protection Group.

To protect DP volumes with only SMB share or NFS export:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **NAS**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the NAS source that contains the volumes you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it. For more information, see Register and Manage NetApp ONTAP Source.
4. From the **Objects** list, select DP volumes of NFS exports or SMB shares.
5. Click **Continue**.
6. For NetApp volumes enabled for NFSv3 and NFSv4.1 protocol, the Backup Preferences for NFS Volumes field will be enabled. Toggle the **Use NFSv4.1 protocol for NFS volumes** to back up the volume via NFSv4.1 protocol.

   > **Note:** This feature is supported from the Cohesity 6.6.0d version onwards. To enable this feature, contact Cohesity Support.

   To back up SMB shares, provide the user access credentials if you have not provided when registering the NetApp ONTAP source.

   > **Note:** To backup NetApp DP volumes with SMB share, the volume needs to be in the junction path on the NetApp.
7. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
8. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Use an Existing Protection Policy.
9. Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
10. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, click **Protect**.
11. Optionally, from the **Protection Group** field, you can select an existing Protection Group. By selecting an existing Protection Group, the settings configured in the existing Protection Group will be applied to this Protection Group. To select an existing Protection Group, perform the following steps or else, go to step 13:
    a. From the **Protection Group** field, enable **Existing Group**.
    b. Select an existing Protection Group from the **Group** drop-down list.
    c. Click **Protect**.

12. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
13. If you need to change any of the additional settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
14. Click **Protect**.

    Your new Protection Group is now active and running.

Protect Data Protection Volumes with Both SMB Share and NFS Export (Mixed Mode)

You can protect NetApp ONTAP DP volumes that are exposed to both NFS and SMB protocols. If the volume is exposed to both NFS export and SMB share, then you must choose how you want Cohesity to back up the NetApp ONTAP volume. Cohesity will backup the ACLs according to the backup preference mode you select.

If you select NFS as the preference mode and the selected volume has:

- NFSv3 enabled, then the SMB ACLs will be lost and only the NFS POSIX ACLs will be backed up.
- NFSv3 and NFSv4.1 enabled, then you can choose to use NFSv4.1 to back up the volume. If you choose NFSv4.1, then the SMB ACLs will be lost, and the NFSv4.1 ACLs will be backed up. However, if you do not toggle the **Use NFSv4.1 protocol for NFS volumes** option on the Cohesity Dashboard, then only the NFS POSIX ACLs will be backed up.

If you select SMB as the preference mode, then the NFS POSIX ACLs will be lost and only the SMB ACLs will be backed up.

> **Note:** NetApp NFSv4.1 volumes are supported from the Cohesity 6.6.0d version onwards. To enable this feature, contact Cohesity Support.

To protect DP volumes exposed to both NFS and SMB protocols:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **NAS**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the NAS source that contains the volumes you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it. For more information, see Register and Manage NetApp ONTAP Source.
4. From the Object list, select DP volumes that are exposed to both NFS and SMB protocols.
5. Click **Continue**.
6. Select the backup preference mode for the mixed-mode volume you want to backup.

7. For NetApp volumes exposed to the NFSv4.1 protocol, the Backup Preferences for the NFS Volumes field will be enabled. Toggle the **Use NFSv4.1 protocol for NFS volumes** option to use the NFSv4.1 protocol.

   > **Note:** This feature is supported from the Cohesity 6.6.0d version onwards. To enable this feature, contact Cohesity Support.

8. To back up SMB shares, provide the user access credentials if you have not provided when registering the NetApp ONTAP source.

   > **Note:** To backup NetApp DP volumes with SMB share, the volume needs to be in the junction path on the NetApp.

9. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
10. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Use an Existing Protection Policy.
11. Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
12. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, click **Protect**.
13. Optionally, from the **Protection Group** field, you can select an existing Protection Group. By selecting an existing Protection Group, the settings configured in the existing Protection Group will be applied to this Protection Group. To select an existing Protection Group, perform the following steps or else, go to step 13:
    a. From the **Protection Group** field, enable **Existing Group**.
    b. Select an existing Protection Group from the **Group** drop-down list.
    c. Click **Protect**.
14. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
15. If you need to change any of the additional settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
16. Click **Protect**.

    Your new Protection Group is now active and running.

Protect Read-Write volumes with Only SMB Share or NFS Export

You can protect NetApp ONTAP RW volumes with only SMB share or NFS export by creating a Protection Group.

To protect RW volumes with only SMB share or NFS export:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **NAS**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the NAS source that contains the volumes you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it. For more information, see Register and Manage NetApp ONTAP Source.
4. From the Object list, select RW volumes of NFS exports or SMB shares.
5. Click **Continue**.
6. For NetApp volumes enabled for NFSv3 and NFSv4.1 protocol, the Backup Preferences for NFS Volumes field will be enabled. Toggle the Use NFSv4.1 protocol for NFS volumes to back up the volume via NFSv4.1 protocol.

   > **Note:** This feature is supported from the Cohesity 6.6.0d version onwards. To enable this feature, contact Cohesity Support.
7. To back up SMB shares, provide the user access credentials if you have not provided when registering the NetApp ONTAP source.
8. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
9. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Use an Existing Protection Policy.
10. Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
11. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, click **Protect**.
12. Optionally, from the **Protection Group** field, you can select an existing Protection Group. By selecting an existing Protection Group, the settings configured in the existing Protection Group will be applied to this Protection Group. To select an existing Protection Group, perform the following steps or else, go to step 13:
    a. From the **Protection Group** field, enable **Existing Group**.
    b. Select an existing Protection Group from the **Group** drop-down list.
    c. Click **Protect**.

13. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
14. If you need to change any of the additional settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
15. Click **Protect**.

    Your new Protection Group is now active and running.

Protect RW Volumes with Both SMB Share and NFS Export (Mixed mode)

You can protect NetApp ONTAP RW volumes that are exposed to both NFS and SMB protocols. If the volume is exposed to both NFS export and SMB share, then you must choose how you want Cohesity to back up the NetApp ONTAP volume. Cohesity will backup the ACLs according to the backup preference mode you select.

If you select NFS as the preference mode and the selected volume has:

- NFSv3 enabled, then the SMB ACLs will be lost and only the NFS POSIX ACLs will be backed up.
- NFSv3 and NFSv4.1 enabled, then you can choose to use NFSv4.1 to back up the volume. If you choose NFSv4.1, then the SMB ACLs will be lost, and the NFSv4.1 ACLs will be backed up. However, if you do not toggle the **Use NFSv4.1 protocol for NFS volumes** option on the Cohesity Dashboard, then only the NFS POSIX ACLs will be backed up.

If you select SMB as the preference mode, then the NFS POSIX ACLs will be lost and only the SMB ACLs will be backed up.

> **Note:** NetApp NFSv4.1 volumes are supported from the Cohesity 6.6.0d version onwards. To enable this feature, contact Cohesity Support.

To protect RW volumes exposed to both NFS and SMB protocols:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **NAS**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the NAS source that contains the volumes you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it. For more information, see Register and Manage NetApp ONTAP Source.
4. From the Object list, select RW volumes that are exposed to both NFS and SMB protocols.
5. Click **Continue**.
6. Select the backup preference mode for the mixed-mode volume you want to backup.

7. For NetApp volumes exposed to the NFSv4.1 protocol, the Backup Preferences for the NFS Volumes field will be enabled. Toggle the Use NFSv4.1 protocol for NFS volumes option to use the NFSv4.1 protocol.

   > **Note:** This feature is supported from the Cohesity 6.6.0d version onwards. To enable this feature, contact Cohesity Support.

8. To back up SMB shares, provide the user access credentials if you have not provided when registering the NetApp ONTAP source.
9. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
10. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Use an Existing Protection Policy.
11. Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
12. If you want to change or configure any of the additional settings , then select **More Options** and perform the below steps or else, click **Protect**.
13. Optionally, from the **Protection Group** field, you can select an existing Protection Group. By selecting an existing Protection Group, the settings configured in the existing Protection Group will be applied to this Protection Group. To select an existing Protection Group, perform the following steps or else, go to step 13:
    a. From the **Protection Group** field, enable **Existing Group**.
    b. Select an existing Protection Group from the **Group** drop-down list.
    c. Click **Protect**.
14. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
15. If you need to change any of the additional settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
16. Click **Protect**.

    Your new Protection Group is now active and running.

Protect Snapshots of VMware VMs stored in NetApp NFS volumes

The prerequisites for protecting snapshots of VMware VMs stored in NetApp NFS volumes are:

- Register your NetApp ONTAP as a source on Cohesity and enable the Storage Array Snapshots option during registration. For details, see Register and Manage NetApp

ONTAP Source.

- Set how often the Storage Array Snapshots have to be taken and how long they have to be retained in the Protection Policy. For details, see Policy Settings for NetApp ONTAP

Perform these steps to protect snapshots of VMware VMs stored in NetApp NFS volumes by creating a Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **Virtual Machines**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the registered source (the vCenter).
4. From the **Objects** list, select the VMware VMs that are backed up by the registered NAS source.

   You can also register a new source with the Cohesity cluster so you can protect it. For more information, see Register and Manage NetApp ONTAP Source.
5. Click **Continue**.
6. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

   Optionally, select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing **Protection Group**.
7. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. The policy you select or create must have the **Storage Array Snapshot** option enabled. For more information, see Policy Settings for NetApp ONTAP.
8. Selection one of the following options:
   - **Protect**: Select this option if you want to complete the Protection Group creation without changing or configuring any of the additional settings. Your new Protection Group is now active and running.
   - **More Options**: If you want to change or configure any of the additional settings, then select **More Options**. For more information, see Add or Edit a Protection Group for Virtual Servers. After configuring the settings, click **Protect**.

> **Note:** To delete the VMs from a Storage Array Snapshot backup job, select all the VMs in that backup job. The **Delete** or **Delete All Snapshots** buttons are not enabled until all the VMs are selected.

Recover Snapshots of VMware VMs stored in NetApp NFS volumes

To recover Snapshots of VMware VMs stored in NetApp NFS volumes to the original or new location:

1. In the Cohesity Dashboard, select **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then do one of the following:
    1. If you want to recover VMs, select **Virtual Machines** > **VMs**. The Virtual Machines page appears. Continue to Step 3.
    2. If you want to recover virtual disks, select **Virtual Machines** > **VMDK**. Follow the procedure described in Recover Virtual Disks to recover the virtual disk to an original or new location.
3. Search for a protected object (VM) or Protection Group that contains the snapshots to recover from by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *. A list of VMs displays.

    You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more VMs or Protection Groups that contain the snapshots to recover. **Ensure that all the VMs you select are from Storage Array Snapshots backups.**
5. From the selected lists, mouse over the snapshot, and then click the edit icon.

    The Edit recovery point for <VM_name> page appears.

    > **Note:** By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page.

    1. To select a snapshot from the Timeline mode:
        i. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

        ii. Select the snapshot to be recovered by clicking the corresponding blue dot.

        For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar.

        iii. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud (

) icon if available in the **Location** option.

iv. Click **Select Recovery Point** and then go to step 6.

2. To select a snapshot from the **List** mode:
   i. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page.

      By default, the page displays the snapshots available in the last one month.

   ii. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

   iii. Select the snapshots to be recovered.

   iv. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud (

) icon if available in the **Location** option.

      a. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the VM to the original location or select **New Location** to recover the volume to an alternate location as a new storage volume.
8. Select **Select Instant Recovery** as the **Recovery Type**.

   The VMs will be instantly available after recovery in the target location. The data will be moved to the target storage later.

   > **Note:** You should have the VirtualMachine.Config.Rename permission on your user account to use this recovery option.

9. Optional. Edit the default settings under **Recovery options**.
10. Click **Start Recovery**.

    The recovery task is initiated. You can monitor the recovery task from the **Recoveries** page.

Additional Settings

Additional Settings	Description
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.

Additional Settings	Description

Snapshot Prefix for DP Volumes	(*Applies only for DP volumes*) Select one of the following options to backup the snapshots from the DP volume to the Cohesity cluster.
	• **None**: Enable this option if you want the Cohesity cluster to take the full backup from the oldest snapshot available on the DP volume and incremental backup from the latest snapshots available on the DP volume. This option is enabled by default.
	> **Note:** If you have selected None as the Snapshot Backup Preferences option in a Protection Group, then you can edit the Protection Group and change the Snapshot Backup Preferences option to Snapshot Prefix or SnapDiff Backup.
	• **Snapshot Prefix**: Enable this option if you want the Cohesity cluster to take the full and incremental backup from the snapshots that match the prefix name you specify:
	  • In the **Incremental Snapshot Prefix** field, specify the prefix of the snapshot name present in the DP volume from which Cohesity cluster can take incremental backup.
	  • In the **Full Snapshot Prefix** field, specify the prefix of snapshot name present in the DP volume from which Cohesity cluster can take the first full backup.

Pre & Post Scripts	Toggle **Enable** to run scripts on a Unix server before and/or after a Protection Group runs. If configured, scripts are run every time an object (or entity) is backed by a protection run. For example, if there are five mount/shares that are backed up as part of the NAS Protection Group and the Pre Script is configured, then five instances of the Pre Script is executed in parallel. You can distinguish between the scripts that are run on the different objects (entities) using the `COHESITY_BACKUP_ENTITY` environment variable. The Cohesity cluster populates this environment variable with the name of the backed object (or entity).
Skip Files on Errors	Toggled on by default. The Protection Group continues to run even if it encounters errors on files, such as permissions errors. If files are skipped, the protection run details page indicates a warning status and provides additional information. If toggled off, the Protection Group stops when it encounters an error.

Additional Settings	Description
Encryption	**Note:** If you want to encrypt NFS data traffic between the NetApp source and Cohesity cluster while backing up NFS volumes, you must modify the supported encryption types value on the compute object to 24 (0x18). This change is required as NetApp does not support ARCFOUR Encryption. Setting the value to 24 (0x18) configures the object to only use AES Encryption. This is needed only for NFS encryption. Toggle on this option to enable in-flight encryption of data that are backed up to the Cohesity cluster. To use the encryption functionality for NFS volumes, you must join the Cohesity cluster to the active directory as the Kerberos server or add the Kerberos provider in the Cohesity cluster Access Management section. This option is disabled by default. **Note:** Encrypting in-flight data will have a minor impact on the performance.

Additional Settings	Description
File DataLock	Toggle on this option to preserve the AccessTime of the files and folders of the selected volumes. The AccessTime provides the lock period of the files and folders that are applied when these files and folders are recovered to the Cohesity View. For more information on File DataLock, see NAS File DataLock. Select one of the following **DataLock** modes: • **Compliance**: In this mode, files or folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files or folders. However, a root user can delete the files and folders. Only a user with Data Security role can edit the File DataLock settings in the Compliance mode. • **Enterprise**: In this mode, files or folders cannot be modified by any user before or after the lock retention period expiration. However, a root user can delete the files or folders before or after the lock retention period. An admin user or a user with a Data Security role can edit the File DataLock settings in the Enterprise mode. This mode is useful for giving storage administrators additional control over the file. You can further configure the File DataLock settings to: • Set the lock period for the new files and folders created on the recovered Cohesity View Provide the following information to set the lock period for the new files and folders that will be created on the recovered Cohesity View. These settings are not applicable for the existing NAS source files and folders recovered on the Cohesity View. • **Default Lock Period**: Provide the time period during which the files should remain locked. When the time period ends, files can be deleted but not changed. By default, the files are locked forever. • **Autolock Files**: By enabling this option, the files are automatically locked if they are not modified in the time period you specify. By default, this option is disabled. • **Manually Lock Using**: Select when the files are to be locked: • **Read-Only**: Select this option to lock files when the file permission is set to read only, for example; by setting the Windows file properties attribute or using the Linux `chmod` command. By default, the files are locked when the file permission is set to read only. • **FutureATime**: Select this option to lock files when the last access time (atime) is set, for example, by using Windows `PowerShell` commands or the Linux touch command.

Additional Settings	Description
	Optionally you can set a **Minimum Retention** and/or a **Maximum Retention** period to ensure that a file's last access time cannot be set outside this period. A request to do so results in a permission denied error if manual locking mode is FutureATime. Otherwise, the request will be silently rejected and the file will be locked for the minimum retention duration. • Override the locking period of the existing files and folders recovered on the Cohesity View: The following option can be applied to all the existing files and folders that will be recovered on the Cohesity View. • **Override (Lock all files until a specific date)**: Enable this option and select a date to override the locking period of the existing files and folders. When you recover the files to the Cohesity View, Cohesity applies the farthest lock period between the override lock period you specify and the source file's lock period. For example; • The lock period of the source file is 09/09/2023 and the override lock period you specified is 09/08/2022. When you recover the file to the Cohesity View, then the lock period of the source file is applied to the recovered file, which is 09/09/2023. • The lock period of the source file is 09/09/2022 and the override lock period you specified is 09/08/2024. When you recover the file to the Cohesity View, then the override lock period you specified is applied to the recovered file, which is 09/08/2024. In **Compliance** mode, a user with Data Security privileges can extend the override date into the future, but no user can shorten the date. In **Enterprise** mode, a user with admin or Data Security privileges can extend or shorten the override date.

Additional Settings	Description
	Note: • You cannot enable or disable File DataLock once the Protection Group is created. • If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed. • You can delete Protection Groups and snapshots having File DataLock enabled. File DataLock also preserves the ATimes of the hardlinks and symlinks.
Exclusions and Inclusions	Toggle on this option if you want to exclude or include locations. By creating exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories and therefore minimize the disk space used to store the data. All the files and folders are included for by default for protection. For more information, see Exclusions and Inclusions for NAS Protection Group. **Note:** Cohesity does not support inclusion or exclusion of paths in a NetApp Snapdif-based Protection Group.
Indexing	This option is toggled on by default to index the data of the Protection Group. Indexing improves files and folders search. Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.

Additional Settings	Description
Quiet Times	This option is available only if the selected policy has at least one quiet time period. Toggle the **Stop in-progress runs when quiet time starts** option to pause or abort the runs during quiet time. If this option is toggled off, after a protection run starts, execution continues even when the quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period. • **Pause Runs** - Select this option to pause the currently executing protection runs and resume the runs when quiet time ends. • **Abort Runs** - Select this option to abort the currently executing protection runs when the quiet time period specified for the Protection Group starts. > **Note:** If the quiet time window for multiple protection runs is scheduled at the same time, then when the protection run is resumed, prioritization of the protection run does not work. 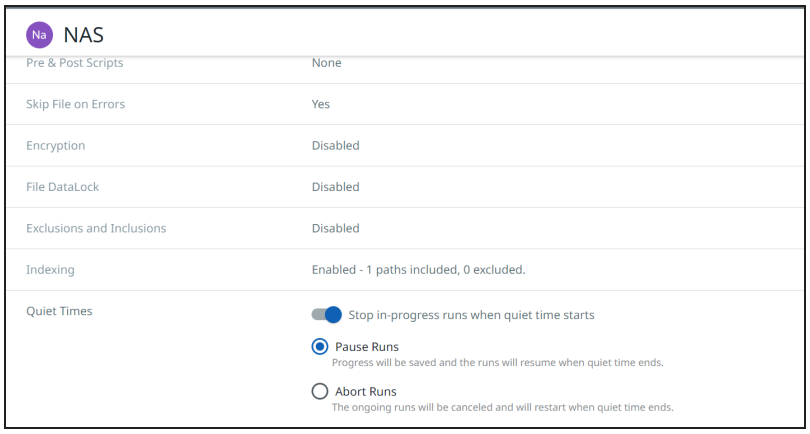

Additional Settings	Description
Alerts	Select one or more of the following **Alerts On** settings if you want alerts to be created for the following triggers: • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • Success: Create an informational alert when the Protection Group completes successfully. Emails are not sent when informational alerts are created. To send emails when alerts are triggered for the Protection Group, in the **Email Recipients** field, add the email addresses of the recipients.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: a. High b. Medium c. Low
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

A Protection Policy is a reusable set of settings. It enables you to define how and when NetApp ONTAP data is protected, replicated, and archived. A standard policy provides various settings that you can configure for a Protection Group.

Create a Protection Policy with the settings listed in the Policy Settings for NetApp ONTAP. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for NetApp ONTAP

The following table lists the applicable policy settings for NetApp ONTAP:

Field	Description
Primary Copy	Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster. Select any external target for performing CloudArchive Direct wherein the primary copy of the data is stored on the external targets. For more information, see CloudArchive Direct.
DataLock	If applicable, add a DataLock for compliance and regulatory requirements, to ensure that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users, including those who have the Data Security role in Cohesity, from modifying or deleting any Snapshots that were generated by theProtection Groups that use this Policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the Protection Group fails. The default time between retries is five minutes. You can customize the number of retries and how long to wait between each attempt.
Periodic full backup	Use this option to choose when to run a periodic full backup.

Field	Description
Quiet Times	Quiet Times defines time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturday or Sunday. Quiet Times only prevents new protection runs from starting during the specified time period. The protection runs that started before the quiet time period will be dependent on the quiet time settings configured in the additional settings of the protection Group. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	To replicate snapshots to another , click the replication icon from the floating menu. The copied snapshots are stored in the storage domain of the remote cluster.
Archival	To add an archive schedule, click the archival icon from the floating menu.
Storage Array Snapshot	Select **Snapshot Array Snapshot** from the floating menu to manage snapshots of VMs stored in NetApp NFS volumes. • **Every**: Specify how often the snapshots have to be taken. • **Retain for xx days**: Specify how long the snapshots are retained on the NetApp NFS volume. Consider the following when you are enabling this option: • **Storage Array Snapshot** can be enabled only on new or existing Protection Policies with Scheduled Backup configured. • If **Replication** or **Archive** is enabled for a Protection Policy along with **Storage Array Snapshot**, replication or archival will only work on the regular backups and not on the Storage Array Snapshots. • **Storage Array Snapshot** cannot be enabled with backup options such as **Continuous Data Protection**.

Use an Existing Protection Policy

You can use an existing Protection Policy protecting NetApp ONTAP. If you are selecting an existing protection policy, then the following options are ignored if they are already configured in the protection policy:

- **Log Backup (Databases)**
- **Continuous Data Protection**
- **Add BMR Backup (Physical Server)**
- **CloudSpin**

NAS File DataLock

Many organizations now rely on WORM (Write Once Read Many) data storage such as Nutanix Files, NetApp Snaplock, and Isilon Smartlock volumes to retain their files and folders in unmodified form for compliance and regulatory purposes and to protect the files from accidental deletion or modification.

When you backup these WORM-capable volumes to the Cohesity cluster, it is essential to have the ability to easily lock the files and folders of the volumes while protecting themselves from accidental or malicious attempts to delete or modify these files. The backup should also provide a strong guarantee for data non-mutability and non-erasability for a certain period of time.

This capability is achieved using Cohesity's File DataLock feature.

FileDataLock is a Cohesity feature that preserves the WORM attributes of any files and folders created in the Cohesity View. You can apply the File DataLock feature for NAS backup (NFS exports and SMB shares) while creating a NAS Protection Group.

With File DataLock enabled in a NAS Protection Group, Cohesity preserves the future AccessTime (ATime) of the files and folders. The AccessTime provides the lock period for files and folders. You can also override this lock period in the File DataLock settings of the NAS Protection Group. This lock period is applied to these files and folders when they are recovered to the Cohesity View. This prevents the users or administrator from modifying or deleting these locked files until the locking period expires.

You can also set lock periods for the new files and folders created on recovered Cohesity View in the File DataLock settings of the NAS Protection Group.

When you recover File DataLock enabled storage volumes to the Cohesity View, the View details page will display a lock icon for the files and folders having File DataLock.

File DataLock is supported for all types of NAS sources including NetApp 7-mode and is agnostic of the NAS protocol.

For information on how to enable File DataLock, see Protect NAS.

File DataLock Modes

You can use the File DataLock feature in one of the two modes:

- **Compliance mode**: The WORM files are protected at the disk level. This mode implements strict regulatory requirements for data retention and complies with SEC Rule 17a-4. In this mode, files and folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files and folders. However, a root user can delete the files and folders. Only a user with a Data Security role can edit the File DataLock settings in the Compliance mode.
- **Enterprise mode**: The WORM files are protected at the file level. This mode does not implement strict regulatory requirements for data retention. In this mode, files and folders cannot be modified by any user before or after the lock retention period

expiration. However, a root user can delete the files and folders during or after the lock retention period. An admin user or a user with the Data Security role can edit the File DataLock settings in the Enterprise mode

> **Note:** An admin user or a user with the Data Security role can edit the File DataLock settings and switch from the Enterprise mode to Compliance mode. But switching from Compliance mode to Enterprise mode is not supported.

Things to remember for File DataLock

- You cannot enable or disable File DataLock once the View is created.
- If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed.
- You can delete Protection Groups and snapshots having File DataLock enabled.

  File DataLock also preserves the ATimes of the hardlinks and symlinks.
- **Recovery**
  - Recovering a File DataLock enabled backed-up volume as a Cohesity View will preserve ATimes and the locking period for the file.
  - Recovering a File DataLock enabled volume to the original location with 'Overwrite existing files/folder' as the recovery option will encounter an error if the files and folders are read-only. However, the ATime and other attributes of other files and folders of the source volumes are preserved.
  - Recovering a File DataLock enabled volume to an alternate NAS recovery will preserve ATimes and other attributes of files and folders of the source volumes.
  - Cloning a recovered View having File DataLock enabled will not inherit the File DataLock properties.
- **Replication**

  When you replicate to a remote Cohesity cluster, the File DataLock settings are also replicated.
- **Archival**

  File DataLock settings are preserved when you recover volume from an archived snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.
- **Cloud Archive Direct**

  File DataLock settings are preserved when you recover from a Cloud ArchiveDirect snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.

Exclusions and Inclusions for NAS Protection Group

In a NAS Protection Group, all the files and folders are included by default for data protection. While creating a Protection Group, you can toggle on **Exclusions and Inclusions** from the **Advanced Settings** options if you want to exclude or include locations. By creating the exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories. As a result, you can minimize the disk space used to store the data.

> **Note:**
> - To add an exclusion or inclusion rule, you must prefix a forward slash (‘/’) or suffix an asterisk (‘*’) to the path or to a particular file within the protected object. For example, ‘**/test**’ or ‘***.txt**’.
> - In exclusion or inclusion rules, file names and folder names are case-sensitive, which means ‘**test.txt**’ and ‘**Test.txt**’ are considered separate files. Use the correct capitalization of file or folder names to ensure that the rules recognize your file or folder.

Add Inclusion

From the Additional Settings options on the Protection Group page, toggle on **Exclusions and Inclusions**. Click **Add Inclusion** to include a particular path or a particular file within the protected object.

For example, consider four directories - **test**, **test1**, **test2**, and **test3** under the protected object with path `/ifs/TestShare1/Folder1`. The table below lists the input types for an inclusion list of folders and their respective outcome:

Input	Outcome
/	Everything inside the protection path is backed up if exclusions are not defined.
/test	When the inclusion path is not suffixed by '/', Protection Group will include all folders starting from the protection path. In this example, /test will behave as /test* and Protection Group will include '**test**', '**test1**', '**test2**', and '**test3**' from the protection path for backup.

Input	Outcome
/test/	When the inclusion path is suffixed by '/', it will only include the specific folder. In this example, /test/ behaves as /test/* and Protection Group will only include '**test**' from the protection path for backup.

- Do not provide the folder name along with the backup path while including files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not include any folder.
- If exclusions are defined, they will take precedence over inclusions.

Add Exclusion

From the Additional Settings options on the Protection Group page, toggle on **Exclusions and Inclusions**. Click **Add Exclusion** to exclude a particular path or a particular file within the protected object. You can also specify regular expressions for excluding files in the format `regex:<regex pattern>`.

For example, consider there are three directories - **Test**, **Test1**, and **Folder1** and a file, **TestFile.tmp** under the protected object with path **/ifs/TestShare1**.

- The **Test1** folder has **TestFile1.tmp**, **TestFile2.tmp** files.
- The **Folder1** folder has **TestFile3.tmp**, **TestFile4.tmp files**, and **SubFolder1** which further has a **Test** folder with **File.txt**, **File1.txt** files in it.

The following figure illustrates the directory structure under the protected object with path **/ifs/TestShare1**.

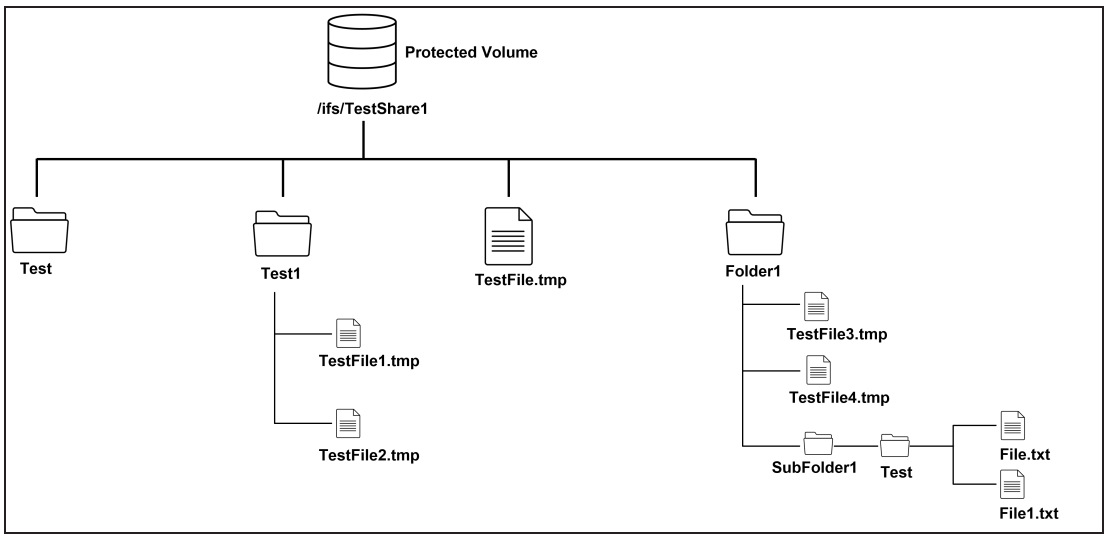

The table below lists the input types for an exclusion list of folders and their respective outcome:

Input	Description	Outcome
/Test	When the exclusion path is not suffixed by '**/**', the Protection Group will exclude all files and folders with **Test** as the prefix. For exclusions, **/Test** behaves as **/Test***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes **Test** and **Test1** folders, and **TestFile1.tmp** file from the protection path. **TestFile3.tmp** and **TestFile4.tmp** files in the **Folder1** folder will not be excluded.
/Test/	When the exclusion path is suffixed by '**/**', the Protection Group excludes only the files and folders with the name **Test**. For exclusions, **/Test/** behaves as **/Test/***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes the **Test** folder of the parent directory. The **Test** sub-folder in **/Folder1/SubFolder1** will not be excluded.
*/Test	Excludes all the files and directories with the name '**Test**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes the **Test** folder of the parent directory and the **Test** sub-folder in **/Folder1/SubFolder1**.
*.tmp	Excludes all the files that with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.
regex:.*\.tmp	Excludes the files and folders with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.

Input	Description	Outcome
regex:/Folder1/.*/Test/.*\.txt$	Excludes all files with **.txt** extension from all the **Test** folders directory in **/Folder1/**.	Excludes **File.txt** and **File1.txt** files in the **/Folder1/SubFolder/Test** folder.

> **Note:** Do not provide the folder name along with the backup path while excluding files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not exclude any folder.

Cohesity automatically excludes the following NetApp system files:

- .vtoc_internal and .bplusvtoc_internal files
- .copy-offload directory and .tokens file

Considerations

Cohesity does not support inclusion or exclusion of paths in:

- an IsilonProtection Group with Changelist enabled.
- a NetApp SnapDiff-basedProtection Group.

Adaptive Backup for NetApp

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

You can now apply a threshold on the CPU capacity utilization for the NetApp source. The Cohesity cluster will adaptively adjust the number of backup IO requests sent to the NetApp source such that the CPU capacity utilized does not exceed the set threshold.

For example, if the threshold value specified is 40% and the primary workload utilizes 30% of the CPU capacity, then the Cohesity cluster will send backup IO requests so as not to exceed 40% CPU on NetApp by consuming only 10% CPU for backup. If the primary load on the NetApp now increases to 35%, then Cohesity cluster will adapt to drop its IO request to consume only 5% of CPU capacity instead of 10% of the CPU capacity. To summarize, Cohesity backup is adaptive, source aware, and sensitive to the load condition on the NetApp primary device.

> **Note:** Since the Cohesity cluster adaptively increases or decreases the percentage of the CPU capacity utilized for the backup operations, setting up a low CPU utilization threshold value might affect the backup duration.

Manage the NetApp ONTAP Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirement.

> **Note:** You cannot update the registered NetApp ONTAP Server and the assigned Storage Domain in the Protection Group settings. You cannot enable or disable the NetApp SnapDiff option on an existing Protection Group.

Add NetApp ONTAP volumes to the Protection Group

You can add new NetApp ONTAP volumes to the existing Protection Group for backup. The newly added volumes are backed up from the next protection run and it will be a full backup.

To add NetApp ONTAP volumes to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
   For NetApp Snapdiff backup, select the NetApp SnapDiff-enabled Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Select the NetApp ONTAP volumes you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added NetApp ONTAP volumes will be backed in the next scheduled protection run.

Remove NetApp ONTAP volumes from the Protection Group

From the **Edit Protection** page, you can remove NetApp ONTAP volumes from the Protection Group that you do not want to back up in the next protection run.

To remove NetApp ONTAP volumes from the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.

For NetApp Snapdiff backup, select the NetApp SnapDiff-enabled Protection Group that you had created for NetApp ONTAP.

3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Unselect the NetApp ONTAP volumes that you want to remove from the Protection Group.
6. Click **Save Selection** > **Save**.

The removed NetApp ONTAP volumes will not be backed in the next scheduled protection run.

Enable Auto Protect

The Auto Protect option enables you to automatically protect the new NetApp ONTAP volumes that are added to the selected parent NetApp ONTAP server. However, you cannot enable auto-protect at the server-level containing both Read-Write volumes and Data-Protect volumes.

To enable auto protect:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click **Edit Objects**.
5. Click the shield icon to enable Auto Protect.
6. Click **Save Selection** > **Save**.

When this option is enabled on the parent NetApp ONTAP server, all the NetApp ONTAP volumes that are added to the parent NetApp ONTAP server in the future are automatically added to the Protection Group and are protected from the next protection run.

For example, you create a Protection Group that backs up a NetApp ONTAP SVM. The Protection Group runs for two weeks and then a new volume is added to the NetApp ONTAP SVM. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new volume is also backed up even though the new volume has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list, and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirement. In the Protection Policy, you can edit the settings listed in the Create or Use an Existing Protection Policy.

> **Note:** You cannot enable or disable the NetApp SnapDiff option on an existing Protection Group.

To the policy settings:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Policy** field, click **Edit**.

   The Edit Policy page appears.
5. Edit the settings listed in the Create or Use an Existing Protection Policy and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, from the drop-down list in the **policy** field, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to back up NetApp ONTAP volumes. It provides you options to back up all the objects or only selected objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- Is configured with replication targets, you can select those targets.
- Is configured with archival targets, you can select those targets.

To start a protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    a. In the **Backup** drop-down list, select any one of the following:
        - **Backup all Objects in the Protection Group**
        - **Backup only selected Objects in the Protection Group**
    b. In the **Backup Type** drop-down list, select the backup type.
    c. In the Replicate section, select the replication target, the retention duration, and enable the **Replicate only fully successful runs** checkbox if you want to replicate the data only after a successful protection run.
    d. In the Archive section, select the archive target, the retention duration, and enable the **Archive only fully successful runs** checkbox if you want to archive the data only after a successful protection run
5. Click **Run Now**.

Pause Protection Runs Instantly

You can instantly pause a protection run of the Protection Group that you have created to back up the NetApp ONTAP sources.

Perform the following steps to pause a currently executing protection run of a Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, click the NetApp ONTAP Protection Group that has been created.
3. Click the action icon next to the protection run and click **Pause Run**.

Pause Future Protection Runs

You can stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the action icon next to the Protection Group and select **Resume**.

Delete the Protection Group

Cohesity provides the following ways to delete the Protection Group associated with NetApp ONTAP:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Protection Group are not deleted
- **Delete Protection Group and Snapshots**: Delete the Protection Group and all Snapshots, any local Snapshots captured by this Protection Group and stored on this Cluster are deleted. In addition, any archived Snapshots stored on external targets that were captured by this Protection Group are also deleted.

For NetApp SnapDiff-based backup:

- You cannot delete the latest snapshot available on both local and remote clusters.
- You cannot delete the snapshot when indexing is in progress.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Group Run.

If the Protection Group is replicating to a remote cluster, the inactive job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the

remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the inactive job.

NetApp ONTAP SnapDiff-based Backup

In the SnapDiff-based backup approach, Cohesity leverages NetApp's SnapMirror cloud functionality to move the data to the Cohesity cluster and NetApp's SnapDiff functionality to get the changed file list. This section provides information on how Cohesity performs the backup and recovery of NetApp ONTAP data protection (DP) and Read-Write (RW) volumes using the SnapDiff-based backup approach.

> **Note:** In this document, NetApp SnapDiff means Cohesity's implementation of SnapDiff using NetApp SnapDiff v3 APIs.

Supplemental terms regarding the Cohesity implementation of the SnapDiff API. Please review before use.

Plan and Prepare for NetApp ONTAP SnapDiff-based Backup

This topic provides information on the supported NetApp ONTAP versions and volumes, minimum permissions, and the prerequisites to register and protect NetApp ONTAP with Cohesity cluster using SnapDiff-based backup approach. For more information, see Register NetApp ONTAP.

Before you register your NetApp ONTAP with the Cohesity cluster and protect the DP volumes and read-write volumes using the NetApp SnapDiff-based backup approach, ensure the following prerequisites:

Prerequisites for NetApp

- Obtain the SnapMirror Cloud license or ONTAP One license for enabling NetApp SnapDiff-based backup.
- Ensure the Cohesity cluster has the FQDN setup configured and all VIPs resolved to the FQDN of the NetApp cluster.
- Enable NetApp Snapdiff V3 on your NetApp Vservers/SVMs by running the following commands:

```
::> set adv
```

```
::*> vserver snapdiff-rpc-server on -vserver <svm_name>
```

- NFS must be enabled on the data LIFs to run SnapDiff APIs for File Level Recovery and to allow RPC initialization.
- To ensure protection using the SnapDiff approach, all NetApp ONTAP volumes must have the 'snapdif-access' capability enabled.

Support Matrix

Before you register your NetApp ONTAP with the Cohesity cluster, ensure that the Cohesity cluster supports the NetApp ONTAP versions and volumes you want to protect. For more information, see Supported Software.

Supported NetApp ONTAP Volumes

The following table lists the supported NetApp ONTAP versions and volume types for NetApp SnapDiff-based backup:

Volume Type	Volume Sub-Type
Data Protection Volume (DP volume)	SnapMirror Destination Volume
	SnapVault Destination Volume
Read-Write Volume (RW volume)	

Volumes that are configured to use NetApp Volume Encryption are supported.

Cohesity does not support the backup of the following NetApp ONTAP volumes for NetApp Snapdiff backup:

- SnapLock Enterprise Volume
- SnapLock Compliance Volume
- Flex Group Volume

Supported NFS and SMB versions

For NetApp Snapdiff DP volume and read-write volume backup, Cohesity supports all the versions of NFS and SMB that NetApp supports.

Minimum Permissions

Ensure you have the required minimum permissions to register NetApp ONTAP as a source with the Cohesity cluster, and to perform backup and recovery.

Minimum Permissions for Registering NetApp ONTAP

The user account used for registering NetApp ONTAP SVM or NetApp ONTAP Cluster must have the required permissions to communicate with the Cohesity cluster

> **Note:** Cohesity does not support SnapDiff based backups when registering a NetApp SVM. SnapDiff based backups are only supported when registering the NetApp cluster.

You can use the REST APIs or the legacy Zephyr Application Programming Interface (ZAPI) roles to configure the permissions for the ONTAP Custer.

Configuring permissions using REST APIs

Considerations:

- To use only REST APIs for backup and recovery, you must create a REST role for configuring the permissions.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

The following tables describe the APIs.

For Cluster, Licensing, and Jobs

API	Access Level	Used By	Description
/api/cluster	Read-only	• GET /api/cluster • GET /api/cluster/nodes • GET /api/cluster/jobs/{uuid}	• Discovers cluster identity and version. • Polls job state for various asynchronous operations.
/api/cluster/licensing **Note:** The secondary path for this API is: /licensing	Read-only	GET /api/cluster/licensing/licenses	Inspects cluster licenses.

For SVM and Network

API	Access Level	Used By	Description
/api/svm/svms	Read-only	GET /api/svm/svms	Lists SVMs along with their supported data protocols within the process of entity-hierarchy discovery.

API	Access Level	Used By	Description
/api/network/ip/interfaces	Read-only	GET /api/network/ip/interfaces	Discovers LIFs, addresses, and services to determine usable NFS/CIFS/management interfaces.

For Protocol and Security Configuration

API	Access Level	Used By	Description
/api/protocols/nfs/services	Read-only	GET /api/protocols/nfs/services/{svm_uuid}	Determines NFS enablement (v3/v4.1) per SVM.
/api/protocols/nfs/kerberos/interfaces	Read-only	GET /api/protocols/nfs/kerberos/interfaces	Maps NFS Kerberos-enabled IPs and hostnames.
/api/protocols/cifs/services	Read-only	GET /api/protocols/cifs/services/{svm_uuid}	Retrieves CIFS server name and AD domain/FQDN.
/api/protocols/cifs/shares	Read-only	GET /api/protocols/cifs/shares	Lists CIFS shares, paths, and ACLs for hierarchy and access checks.
/api/protocols/cifs/users-and-groups/privileges	Read/Write	POST /api/protocols/cifs/users-and-groups/privileges	Adds or updates CIFS user/group privileges when NAS workflows require it.
/api/security/certificates	Read-only	GET /api/security/certificates	Inspects SVM SSL certificate configuration and server-authentication state.

For Volumes, Qtrees, and Export Policies

API	Access Level	Used By	Description
/api/storage/volumes	Read-only	• GET /api/storage/volumes • GET /api/storage/volumes/{uuid}	Discovers NetApp volumes (capacity, junction path, export policy, security style, and related attributes) for entity hierarchy and NAS backup targeting.
/api/storage/qtrees	Read/Write	• GET /api/storage/qtrees • POST /api/storage/qtrees	• Reads qtree security-style and exports policy information. • Creates qtrees as part of NAS export-related operations.
/api/protocols/nfs/export-policies	Read/Write	• GET/POST /api/protocols/nfs/export-policies • GET /api/protocols/nfs/export-policies/{id}/rules	Lists, creates, and modifies NFS export policies and rules that control NAS backup and access behavior.

For Snapshots (NAS backup and recovery)

API	Access Level	Used By	Description
/api/storage/volumes/*/snapshots **Note:** The secondary path for this API is: /*/snapshots	Read/Write	• POST /api/storage/volumes/{uuid}/snapshots • GET /api/storage/volumes/{uuid}/snapshots • GET /api/storage/volumes/{uuid}/snapshots/{snapshot_uuid} • PATCH /api/storage/volumes/{uuid}/snapshots/{snapshot_uuid} • DELETE /api/storage/volumes/{uuid}/snapshots/{snapshot_uuid}	Creates, lists, modifies (rename or change expiry), and deletes volume snapshots for NAS backups and recoveries. In some ONTAP versions, RBAC is purely prefix-based. In these versions, configuring api/storage/volumes with Read/Write access is sufficient to cover snapshots. If the UI allows the more specific /api/storage/volumes/*/snapshots entry, this document assumes it is granted Read/Write access.

For Private CLI Endpoints

API	Access Level	Used By	Description
/api/private/cli/snapshot **Note:** Secondary path: /private/cli/snapshot	Read/Write	GET /api/private/cli/snapshot?snapshot=...&volume=...&fields=instance-uuid	Retrieves CLI snapshot metadata like instance UUID as needed for NAS workflows.

API	Access Level	Used By	Description
/api/private/cli/vserver/export-policy/check-access **Note:** Secondary path: /private/cli/vserver/export-policy/check-access	Read-only	GET /api/private/cli/vserver/export-policy/check-access	Uses the private CLI layer to conduct export-policy access checks.

When using SnapDiff for source-initiated NAS backup with SnapMirror Cloud and object-store SnapMirror, these additional privileges are needed beyond the NAS-only role:

API	Access Level	Used By	Description
/api/cluster	Changes access from Read-only to Read/Write	Used internally by NetApp for the SnapMirror create command.	Changes access level to enable SnapMirror access. For more details, see ONTAP documentation.
/api/storage/volumes	Changes access from Read-only to Read/Write	POST /api/storage/volumes/{uuid}/snapdiff/sessions?return_records=true	Creates SnapDiff V3 sessions on a volume.

API	Access Level	Used By	Description
/api/snapmirror/policies	Read/Write	• POST /api/snapmirror/policies • GET /api/snapmirror/policies • GET /api/snapmirror/policies?name=... • GET /api/snapmirror/policies/{uuid} • PATCH /api/snapmirror/policies/ • DELETE /api/snapmirror/policies/{uuid}	Creates, inspects, updates, and deletes SnapMirror policies that manage SnapDiff and SnapMirror Cloud relationships.
/api/snapmirror/relationships	Read/Write	• POST /api/snapmirror/relationships • GET /api/snapmirror/relationships?source.path=...&destination.path=... • GET /api/snapmirror/relationships/{uuid} • PATCH /api/snapmirror/relationships/{uuid} • GET /api/snapmirror/relationships/{uuid}/transfers • GET /api/snapmirror/relationships/{uuid}/transfers/{transfer_uuid} • POST /api/snapmirror/relationships/{uuid}/transfers	Configures and manages SnapMirror relationships and transfers for SnapDiff-based source-initiated backup and restore.
/api/snapmirror/object-stores	Read/Write	• GET /api/snapmirror/object-stores/{obj_store_uuid}/endpoints/{endpoint_uuid}/snapshots • GET /api/snapmirror/object-stores/{obj_store_uuid}/endpoints/{endpoint_uuid}/snapshots?name=... • DELETE /api/snapmirror/object-stores/{obj_store_uuid}/endpoints/{endpoint_uuid}/snapshots/{uuid} • DELETE /api/snapmirror/object-stores/{obj_store_uuid}/endpoints/{endpoint_uuid}	Manages SnapMirror object-store endpoints and snapshots for SnapDiff source-initiated workflows.

API	Access Level	Used By	Description
/api/cloud/targets	Read/Write	• POST /api/cloud/targets • GET /api/cloud/targets • GET /api/cloud/targets?name=... • GET /api/cloud/targets/{uuid} • PATCH /api/cloud/targets/{uuid} • DELETE /api/cloud/targets?name=...	Creates, inspects, updates, and deletes cloud and object-store targets that are used as backing stores for SnapMirror Cloud relationships.

Configuring Permissions Using the Legacy ZAPI Roles

For NetApp SnapDiff-based backup, you must create two roles for the user account you use for registering the NetApp ONTAP Cluster with the Cohesity cluster:

- rest-role
- a role dealing with commands

For the role dealing with commands the user account must have the following permissions:

Access-level	Command*	Description	Protocol
All	volume snapshot	Allows fetching, creating, and deleting snapshots for volumes.	SMB / NFS
ReadOnly	cifs	Fetches information about CIFS/SMB shares for volumes.	SMB / NFS
	cluster identity	Fetches information about the cluster.	SMB / NFS
	network interface	Fetches information about network interfaces to which Cohesity cluster connects for mounting volumes.	SMB / NFS
	volume	Fetches information about volumes.	SMB / NFS
	vserver	Fetches information about SVM.	SMB / NFS
	version	Fetches information about ONTAP version	SMB/NFS
	system license	Fetches information about system licenses	SMB/NFS

*For the above command listed, the login method is ONTAPI.

For the rest-role, the user account must have the following API permissions:

Access-level	API	Description	Protocol
All	/api/cloud/targets	Allows fetching, creating and deleting object store configs.	SMB / NFS
	/api/cluster/licensing	Fetches access tokens for REST API calls.	SMB / NFS
	/api/snapmirror	Allows fetching, creating and deleting snapshots in a SnapMirror cloud.	SMB / NFS
	/api/snapmirror/policies	Manage SnapMirror policies of type 'async and 'sync'.	SMB / NFS
	/api/snapmirror/relationships	Creates and manages the SnapMirror relationships. SnapMirror transfer APIs are used to manage data transfers.	SMB / NFS
	/api/snapmirror/relationships/<relationship_ID>/transfer	Manage data transfers.	SMB / NFS
	/api/storage/volumes	Fetches information about volumes.	SMB / NFS
ReadOnly	/api/cluster/jobs	Fetches the status of long running jobs.	SMB / NFS

* For the above command listed, the login method is HTTP.

Credentials for NetApp ONTAP Backup with Multiple SVMs

To register NetApp ONTAP with multiple storage virtual machine (SVM), create a custom role with the required permissions and a local user at the SVM level. Assign the custom role to the local user. Use the respective local user account to register multiple SVMs.

Ports Used for Communication

For information on port numbers required, see Manage Firewall Ports.

Considerations

The following are the important considerations to be noted before you protect your NetApp ONTAP data using NetApp SnapDiff-based backup:

- The NAS storage FQDN must be resolvable. Even if registered by IP, SMB restore workflows depend on the FQDN for repeated SMB share connections. Ensure that the DNS or cluster host mapping resolves the storage FQDN.
- To use NetApp SnapDiff-based backup, ensure your Cohesity cluster has the FQDN setup configured and all VIPs resolve to the cluster name.
- Cohesity does not support NetApp SnapDiff-based backup if the NetApp source is registered as NetApp ONTAP SVM.
- You cannot protect the same NetApp volume in two Protection Groups (Protection Groups created within a Cohesity cluster or different Cohesity clusters) with NetApp SnapDiff-based backup. However, you can protect the same NetApp volume in two Protection Groups if one Protection Group is with SnapDiff enabled and the other Protection Group is without SnapDiff enabled.
- The backup of a NetApp volume will fail if there is already a SnapMirror cloud relationship created between the NetApp volume and another Cohesity S3 View.
- SnapDiff-based backups are immutable. That is, you cannot delete the SnapDiff-based backed-up data.
- When assigning or creating backup policies to a Protection Group for NetApp volumes, it is recommended to create the schedule of backup policy in-sync or less aggressive than the SnapMirror or SnapVault schedule. This ensures the presence of the snapshot in the volume, prior to the incremental backup triggered by the Cohesity backup policy schedule.
- CloudArchive Direct is not supported.
- CloudArchive with Periodic full is not supported. However, CloudArchive with incremental forever format to an S3-compatible target, NAS, and Tape is supported.
- You cannot delete the latest snapshot available on both local and remote clusters.
- You cannot delete the snapshot when indexing is in progress.
- NetApp requires an S3 Access key for accessing a Cohesity S3 view. For generating Cohesity S3 access keys, you must use one of the following accounts to create an S3 View:
    - A local Cohesity user.
    - An Active Directory user that was explicitly added to the Cohesity cluster and assigned a role. This user does not rely on an AD group for access to the Cohesity cluster.

  S3 access keys will not be generated for a Helios user or for an AD user that has Cohesity cluster access through an Active Directory group only.

- Pausing an ongoing SnapDiff-based protection run and resuming the protection run at a later time is not supported.
- Failover and failback are not supported.
- If the latest snapshot in the NetApp ONTAP volume is already backed up and there is no new snapshot in the NetApp volume, then the backup will fail.
- You should not manually perform 'snapmirror update'. Ie, you should not update the Cohesity S3 View. The following error message will be displayed when you perform SnapDiff-based backup after updating Cohesity S3 View:

  ```
  Unable to establish a SnapDiff session.
  ```

- Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.

The following are the considerations for Storage Volume Recovery:

- The recovery target must be a DP volume.
- Recover to the **Original Location** recovery workflow is not supported.
- You must select a new DP volume, which is not in a SnapMirror relationship. When a volume that was backed up via SnapDiff is recovered to an alternate NetApp DP volume, then the target DP volume will become an RW volume. Therefore, you cannot use the same volume as the target volume for the next volume recovery task.
- Recovering data to the following is not supported as NetApp SnapDiff stores the data in a proprietary format:
- Recovering to a non-NetApp NAS.
- Recovering to a new Cohesity View.

The following are the considerations for Directory-level Recovery:

- Directory-level recovery is supported through file and folder recovery.
- Due to SnapMirror limitations, you can recover up to eight individual files or one folder at a time.

The following are the considerations for File-level Recovery:

- The recovery target must be a Read-Write volume.
- To perform file recovery, indexing must be enabled on the Protection Group.
- You can select up to eight files for recovery at a time. If you select more than eight files for recovery, then the recovery task will fail.
- Downloading zip files is not supported.
- Recovery of hardlinks and symbolic links is not supported.
- Mtime is not preserved for SMB.

- File and folder recovery from an archive (incremental forever archival format) is not supported.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.

Best Practices for SnapDiff-based Backup

1. NFS protocol must be enabled on the SVM from which volumes are backed up using SnapDiff.
2. The SnapDiff API will run over data LIFs (logical interface) since they cannot be run over management LIF. Therefore, NFS must be enabled on the data LIFs used by the SVM to allow RPC initialization.
3. An inter-cluster LIF must be configured on the NetApp ONTAP cluster. NetApp ONTAP verifies connectivity prior to saving this configuration information.

Register and Manage NetApp ONTAP Source

To protect your NetApp ONTAP data, you need to register your NetApp ONTAP as a source on the Cohesity cluster. After registration, you can manage the source on the Cohesity cluster in the following ways:

- Update the NetApp ONTAP source configuration
- Refresh the source details
- Unregister the NetApp ONTAP source from the Cohesity cluster.

Register NetApp ONTAP as a Source

Before you register NetApp ONTAP as a source on the Cohesity cluster, ensure that you:

1. Review the requirements, considerations, and best practices.
2. Perform the prerequisite tasks.

For more information, see Plan and Prepare for File-Runner Based NetApp ONTAP Backup.

> **Note:** Before you register NetApp ONTAP as a source on the Cohesity cluster, ensure that you have reviewed and completed the instructions provided in the Plan and Prepare for File-Runner Based NetApp ONTAP Backup

To register NetApp ONTAP as a source:

> **Note:** To register the NetApp ONTAP cluster for NetApp SnapDiff-based backup, you must obtain the SnapMirror Cloud license from NetApp.

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click **Register** on the top-right of the page and then select **NAS**.
   The Register NAS page appears.
3. From the NAS Source drop-down list, select **NetApp**.
4. Enter the following information, and then click **Register**.

Field Name	Description
NetApp Source Type	The NetApp ONTAP cluster or Vserver/SVM to register. > **Note:** Cohesity does not support SnapDiff-based backup if the NetApp source is registered as NetApp ONTAP SVM.
Host or IP Address	The Host or IP address of the NetApp ONTAP cluster or SVM to register. For NetApp SnapDiff-based backup, use the Host or IP address of the NetApp ONTAP cluster to register.
Username	Username used to access the NetApp ONTAP cluster or SVM. Specify a user that has adequate privileges to perform actions on the Source. If you are providing a domain user account, then you must specify the user name in the following format: `fully_qualified_domain_name\username`. For example, `xydc.local\user1`. For more information, see Minimum Permissions. > **Note:** The username used to register the NetApp ONTAP cluster and SVM is case sensitive.
Password	Enter the password for the specified user.
Back Up SMB Volumes	Whether NetApp ONTAP will back up SMB volumes. If yes, provide the local or Active Directory user credentials that allow at least read access on the NetApp ONTAP cluster or SVM. > **Note:** The user must have full control on the restore target during restore.

Field Name	Description
Filter IPs	Enable this option to filter the IP addresses of the NetApp ONTAP cluster. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the NetApp ONTAP cluster. To filter the IP addresses of the NetApp ONTAP cluster, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the NetApp ONTAP source through which the communication to the Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the NetApp ONTAP source through which the communication to the Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.
Storage Array Snapshots	Enable this option to manage the snapshots of VMware VMs stored in NetApp NFS volumes from the Cohesity cluster. • **Cap number of snapshot per volume**: Enter the maximum number of snapshots that can be stored on each NetApp storage volume. • **Cap storage consumed by snapshot per volume**: Enter the percentage of the NetApp volume's storage you want the snapshots to consume. **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

When registration finishes, a new entry appears on the Sources page. You can now create a Protection Group to backup NetApp ONTAP volumes. For more information, see Backup NetApp ONTAP Using File-Runner .

Update the NetApp ONTAP Source Configuration

You can update the NetApp ONTAP details that you had provided during the registration process with the latest NetApp ONTAP configuration.

To edit the NetApp ONTAP source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the NetApp ONTAP source and select **Edit**.
3. In the Edit NAS page, update the respective configurations.
4. In the **Password** field, enter the password of the NetApp ONTAP user account, and then click **Update**.

Refresh the NetApp ONTAP Source

You can refresh the NetApp ONTAP configuration and fetch the latest changes on the NetApp ONTAP.

To refresh the NetApp ONTAP source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the NetApp ONTAP source and select **Refresh**.

Unregister NetApp ONTAP

If you plan to stop taking backups of the NetApp ONTAP volumes, you can unregister the NetApp ONTAP source from the Cohesity cluster. Before you unregister, you must delete all the Protection Groups associated with the NetApp ONTAP server. For more information, see Delete the Protection Group.

To unregister the NetApp ONTAP server on the Cohesity cluster:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the NetApp ONTAP source and select **Unregister**.
3. Click **Unregister**.

Backup NetApp ONTAP Using SnapDiff

After registering your NetApp ONTAP server as a source on the Cohesity cluster, you can create a Protection Group to back up your NetApp ONTAP DP and RW volumes. You can back up one or more NetApp ONTAP NFS or SMB shares in a Protection Group while preserving the ACLs and Extended attributes on SMB.

To perform SnapDiff-based backup, you must enable the **NetApp SnapDiff** option available while creating a Protection Group. For leveraging the existing snapshots present in a NetApp ONTAP volume using SnapDiff-based backup, you must create a Cohesity S3 View.

When a Protection Group with the **NetApp SnapDiff** option enabled is run, Cohesity will request NetApp to transfer the data. For the full backup, NetApp starts ingesting the list of files in the volume to Cohesity in parallel to the data transfer. For incremental backup, ingesting the list of files and data transfer occurs sequentially.

For ingesting the list of files, Cohesity creates a SnapDiff session and ingest files sequentially. For data transfer, Cohesity uses NetApp's SnapMirror cloud functionally which writes to the Cohesity S3 View.

Once the SnapDiff session is initialized, for NetApp ONTAP DP volume, a snapshot with the label, `snapmirror.<cluster_ID>.<Year-Month-Date-time>` is created on the NetApp ONTAP volume. Cohesity takes the full backup from this snapshot leveraging NetApp to write the data from this snapshot to the Cohesity S3 View. For taking a full backup of a DP

volume, you must initialize the DP volume with the RW volume before creating a Protection Group.

When you create a Protection Group for the NetApp ONTAP volumes with the NetApp SnapDiff option enabled, you must select an existing or create a new Cohesity S3 View. When creating an S3 view, ensure to select :

- **Object Services** as the **Category** value
- **Object ID** as the **Object Key Pattern** value.
- **S3** as the **Read/Write Protocol** value

For information see, Create Views Using a Custom Template.

For information on:

- The prerequisites for registering and protecting NetApp, see Plan and Prepare for NetApp ONTAP SnapDiff-based Backup.
- Registering NetApp ONTAP with Cohesity cluster, see Register and Manage NetApp ONTAP Source.

Protect NetApp ONTAP Using SnapDiff-based Backup

In Cohesity, NetApp ONTAP volumes are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

> **Note:** Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.

To protect NetApp ONTAP volumes with SnapDiff-based backup approach:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **NAS**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the NetApp source that contains the volumes you want to back up. You can also register a NAS new source with the Cohesity cluster so you can protect it. To register a new NAS source, click Register Source from the drop-down. For more information, see Register and Manage NetApp ONTAP Source.
4. From the list of volumes, select the NetApp ONTAP volumes for which you want to perform the SnapDiff-based backup.
5. Click **Continue**.

You can ignore the **Backup Preference for Mixed Mode Volumes** field, as it is not applicable for SnapDiff-based backup.

6. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
8. Select **More Options** to enable NetApp SnapDiff and configure settings, and additional settings.
9. Enable the **NetApp SnapDiff** option and provide the following details to leverage SnapMirror to take the full and incremental backup of the NetApp volume.
    1. From the **Cohesity S3 View** field, select an existing or create a new Cohesity S3 View. The snapshot taken from the volume will be written to the Cohesity S3 View. When creating an S3 view:
        1. Select **Object Services** for **Category**.
        2. Select **Object ID** for **Object Key Pattern**.
        3. Select **S3** for **Read/Write Protocol**.
        4. Select **Backup Target High** for **QoS**.
        5. Ensure that the created S3 view has the NetApp cluster IPs whitelisted.

        For information see, Create Views Using a Custom Template.
10. Optionally, from the **Protection Group** field, you can select an existing Protection Group. By selecting an existing Protection Group, the additional settings configured in the existing Protection Group will be applied to this Protection Group. To select an existing Protection Group, perform the following steps or else, go to step 12:
    1. From the **Protection Group** field, enable **Existing Group**.
    2. Select an existing Protection Group from the **Group** drop-down list.
    3. Click **Protect**.
11. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
12. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
13. If you need to change any of the advanced settings on the Protection Group page, click the down arrow icon next to **Advance Settings** and click **Edit** of the respective field. For more information, see Additional Settings.
14. Click **Protect**.

Your new Protection Group is now active and running

Additional Settings

Additional Settings	Description
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
Indexing	This option is toggled on by default to index the data of the Protection Group. Indexing improves files and folders search. Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.
Quiet Times	This option is available only if the selected policy has at least one quiet time period. Toggle the **Stop in-progress runs when quiet time starts** option to pause or abort the runs during quiet time. If this option is toggled off, after a protection run starts, execution continues even when the quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period. • **Pause Runs** - Select this option to pause the currently executing protection runs and resume the runs when quiet time ends. • **Abort Runs** - Select this option to abort the currently executing protection runs when the quiet time period specified for the Protection Group starts. > **Note:** If the quiet time window for multiple protection runs is scheduled at the same time, then when the protection run is resumed, prioritization of the protection run does not work.

Additional Settings	Description
Alerts	Select one or more of the following **Alerts On** settings if you want alerts to be created for the following triggers: • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • Success: Create an informational alert when the Protection Group completes successfully. Emails are not sent when informational alerts are created. To send emails when alerts are triggered for the Protection Group, in the **Email Recipients** field, add the email addresses of the recipients.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: a. High b. Medium c. Low
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

A Protection Policy is a reusable set of settings. It enables you to define how and when NetApp ONTAP data is protected, replicated, and archived. A standard policy provides various settings that you can configure for a Protection Group.

Create a Protection Policy with the settings listed in the Policy Settings for NetApp ONTAP. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for NetApp ONTAP

The following table lists the applicable policy settings for NetApp ONTAP:

Field	Description
Primary Copy	Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster.
DataLock	If applicable, add a DataLock for compliance and regulatory requirements, to ensure that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users, including those who have the Data Security role in Cohesity, from modifying or deleting any Snapshots that were generated by theProtection Groups that use this Policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Backup	Defines when backups are captured, how long they are retained and the type of backup to capture.
Periodic full backup	Use this option to choose when to run a periodic full backup.
Quiet Times	Quiet Times defines time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet time period will be dependent on the quiet time settings configured in the additional settings of the protection Group. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the Protection Group fails. The default time between retries is five minutes. You can customize the number of retries and how long to wait between each attempt.

Field	Description
Primary Copy	Select the primary copy to be on the local cluster. For NetApp SnapDiff-enabled Protection Group, only the metadata is copied onto the Cohesity cluster. The data will be backed up to the Cohesity S3 View. **Retain for**: Specify how many days backups are retained on the Cohesity cluster before they are deleted.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
Replication	To replicate snapshots to another , click the replication icon from the floating menu. The copied snapshots are stored in the storage domain of the remote cluster.
Archival	To add an archive schedule, click the archival icon from the floating menu.

Use an Existing Protection Policy

You can use an existing Protection Policy protecting NetApp ONTAP (also using NetApp Snapdiff backup). If you are selecting an existing protection policy, then the following options are ignored if they are already configured in the protection policy:

- **Log Backup (Databases)**
- **Continuous Data Protection**
- **Add BMR Backup (Physical Server)**
- **CloudSpin**

Manage the NetApp ONTAP Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirement.

> **Note:** You cannot update the registered NetApp ONTAP Server and the assigned Storage Domain in the Protection Group settings. You cannot enable or disable the NetApp SnapDiff option on an existing Protection Group.

Add NetApp ONTAP volumes to the Protection Group

You can add new NetApp ONTAP volumes to the existing Protection Group for backup. The newly added volumes are backed up from the next protection run and it will be a full backup.

To add NetApp ONTAP volumes to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
   For NetApp Snapdiff backup, select the NetApp SnapDiff-enabled Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Select the NetApp ONTAP volumes you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added NetApp ONTAP volumes will be backed in the next scheduled protection run.

Remove NetApp ONTAP volumes from the Protection Group

From the **Edit Protection** page, you can remove NetApp ONTAP volumes from the Protection Group that you do not want to back up in the next protection run.

To remove NetApp ONTAP volumes from the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
   For NetApp Snapdiff backup, select the NetApp SnapDiff-enabled Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Unselect the NetApp ONTAP volumes that you want to remove from the Protection Group.
6. Click **Save Selection** > **Save**.

The removed NetApp ONTAP volumes will not be backed in the next scheduled protection run.

Enable Auto Protect

The Auto Protect option enables you to automatically protect the new NetApp ONTAP volumes that are added to the selected parent NetApp ONTAP server. However, you cannot enable auto-protect at the server-level containing both Read-Write volumes and Data-Protect volumes.

To enable auto protect:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.

3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click **Edit Objects**.
5. Click the shield icon to enable Auto Protect.
6. Click **Save Selection** > **Save**.

When this option is enabled on the parent NetApp ONTAP server, all the NetApp ONTAP volumes that are added to the parent NetApp ONTAP server in the future are automatically added to the Protection Group and are protected from the next protection run.

For example, you create a Protection Group that backs up a NetApp ONTAP SVM. The Protection Group runs for two weeks and then a new volume is added to the NetApp ONTAP SVM. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new volume is also backed up even though the new volume has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list, and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirement. In the Protection Policy, you can edit the settings listed in the Create or Use an Existing Protection Policy.

> **Note:** You cannot enable or disable the NetApp SnapDiff option to an existing Protection Group.

To the policy settings:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Policy** field, click **Edit**.

   The Edit Policy page appears.
5. Edit the settings listed in the Create or Use an Existing Protection Policy and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, from the drop-down list in the **policy** field, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to back up NetApp ONTAP volumes. It provides you options to back up all the objects or only selected objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- Is configured with replication targets, you can select those targets.
- Is configured with archival targets, you can select those targets.

To start a protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    a. In the **Backup** drop-down list, select any one of the following:
        - **Backup all Objects in the Protection Group**
        - **Backup only selected Objects in the Protection Group**
    b. In the **Backup Type** drop-down list, select the backup type.
    c. In the Replicate section, select the replication target, the retention duration, and enable the **Replicate only fully successful runs** checkbox if you want to replicate the data only after a successful protection run.
    d. In the Archive section, select the archive target, the retention duration, and enable the **Archive only fully successful runs** checkbox if you want to archive the data only after a successful protection run
5. Click **Run Now**.

Pause Future Protection Runs

You can stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for NetApp ONTAP.
3. Click the action icon next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the action icon next to the Protection Group and select **Resume**.

Supplemental Terms for SnapDiff API Implementation

COHESITY, INC. (TOGETHER WITH ITS AFFILIATES, "**COHESITY**") AGREES TO SUPPLY AND/OR LICENSE CERTAIN OF SUBJECT IMPLEMENTATION TO YOUR BUSINESS OR ORGANIZATION ("**CUSTOMER**") ONLY IF (A) YOU REPRESENT AND WARRANT THAT YOU HAVE THE AUTHORITY TO LEGALLY BIND CUSTOMER AND (B) YOU ACCEPT AND AGREE ON BEHALF OF CUSTOMER TO BE BOUND BY ALL OF THE SUPPLEMENTAL TERMS (THIS "**SUPPLEMENTAL TERMS**", INCLUDING ALL DOCUMENTS INCORPORATED HEREIN BY REFERENCE), WHICH SHALL BE DEFINITIVELY EVIDENCED BY ANY ONE OF THE FOLLOWING MEANS: YOUR CLICKING THE "ACCEPT", "CONTINUE" OR A SIMILAR BUTTON, AS APPLICABLE; YOUR SIGNATURE ON A TANGIBLE COPY OF THIS LICENSE; YOUR INSTALLATION OR USE OF THE PRODUCTS AND/OR SOFTWARE; OR BY SUCH OTHER CONTRACT FORMATION MECHANISM AS MAY BE RECOGNIZED UNDER APPLICABLE LAW FROM TIME TO TIME. IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS IN THIS AGREEMENT, DO NOT COPY, INSTALL OR USE THE PRODUCT(S) OR ANY ASSOCIATED SOFTWARE. THESE SUPPLEMENTAL TERMS ARE IN ADDITION TO AND SUPPLEMENT THE TERMS AND CONDITIONS IN THE COHESITY EULA OR SIMILAR COMMERCIAL AGREEMENT AGREED TO BY COHESITY AND CUSTOMER WITH RESPECT TO COHESITY PRODUCTS AND SERVICES. THESE SUPPLEMENTAL TERMS APPLY ONLY WITH RESPECT TO THE SUBJECT IMPLEMENTATION.

Customer agrees not to, or authorize or assist a third party to:

1. decompile, disassemble, decrypt, extract, or otherwise reverse engineer or attempt to reconstruct, or discover any source code of the SnapDiff APIs or NetApp Format,

including any underlying ideas, algorithms, API Keys or file or data formats of, or used in, the SnapDiff APIs, or any portion of the SnapDiff APIs as implemented in the Subject Implementation, by any means whatever; provided that the foregoing restriction shall not limit any activity that is required to be permitted under the applicable law in a specific jurisdiction; further provided, however, that if Customer wishes to exercise its rights to engage in activities that are permitted under the applicable law in a specific jurisdiction that otherwise would be prohibited by the foregoing restrictions, Customer will provide prompt written notice to Cohesity and provide Cohesity and its licensors with a reasonable amount of time in which to perform the activity and provide related results to the Licensee;

2. remove or conceal any product identification, copyright or other notices contained in or on the Subject Implementation;
3. publish or provide to a third party any results of benchmark tests run on the Subject Implementation, any portion of the Subject Implementation, without Cohesity's prior written consent;
4. disable, disrupt, circumvent, interfere with, or otherwise violate or impair the security of the Subject Implementation;
5. access any application, system, service, computer, data, account, or network without authorization using the Subject Implementation;
6. combine, remix, link against, transform, build upon, create a derivative work with, or otherwise use the Subject Implementation or otherwise such that the Subject Implementation or NetApp Format would be licensed under a license that requires the SnapDiff APIs, Subject Implementation, or any portion thereof to be (i) disclosed or distributed in source code form, or (ii) licensed under terms that allow any licensee to copy, create derivative works and distribute the software without any fee or cost;
7. allow SnapDiff APIs to be visible, discoverable, usable, accessible, or otherwise expose, either directly or indirectly, by any third party; or
8. display, reproduce, publish, sell, offer for sale, modify, or create a derivative work of any part of the Subject Implementation

Troubleshoot NetApp Protection

- If backing up a NetApp volume fails with the following error: An entity being protected was not found in the global entity hierarchy, this could have been caused because the volume was moved to a different aggregate.

  Workaround - Edit the Protection Group that protects the volume that was moved and select the moved volume for protection.

- Registering and backing up a NetApp volume as an NFS NAS mount point results in a larger backup size because all files inside the .snapshot directory are backed up.

  Workaround - Do one of the following:

  - Register a path to a snapshot on that NetApp volume
  - Register it as a NetApp source instead of a NAS mount point
  - In the Protection Group's settings configure the filter to skip the .snapshot directory
  - Disable the configuration on that NetApp volume to show the .snapshot directory
- For an incremental backup, if there is no snapshot created in the DP volume after the last snapshot that is used for the full or incremental backup, then the protection run fails with the following error message:

```
No new snapshot found to take backup. Last snapshot used for backup:
<snaphot_name> with create time: <create_time> for volume: <volume_
name> on vserver: <server_name>
```

- If the snapshot label specified in the protection group does not match with the snapshot label in the DP volume, then the protection run fails with the following error message:

```
No snapshot with label: home found in the volume: <volume_name> on
vserver: <server_name>. Cohesity discovers snapshots for DataProtect
volumes and uses existing snapshots for taking backups. Please verify
the labels for the snapshot provided.
```

- If the backup fails for a protection group that has **Snapshot Array Management** enabled with the following error message, ensure the Storage Array backing the DataStore in which the VM resides is registered with the Cohesity Cluster:

```
VM cannot be backed up due to error: [kUnknownStorageDevice]: Storage
array of the datastore is not registered.
```

- If you do not enable the **Snapshot Prefix** option while creating the Protection Group, then the backup of NetApp DP volume may fail with the following error.

```
[kInvalidRequest]: Error: No new snapshot found in the volume: <volume_
name> on  vserver: <vserver_name>. Cohesity discovers snapshots and
uses existing snapshots for taking backups. Please verify the labels
for the snapshot provided
```

  Workaround - If the Snapshot Prefix option is not enabled, then ensure that an appropriate older snapshots are available in the NetApp source.

Nutanix Files

Nutanix Files is a hardware and software-based solution for file storage and retrieval. You can configure Cohesity to back up and restore data on Nutanix Files arrays.

You can back up and restore Nutanix Files mount targets using the Cohesity File-Runner backup method or the vendor-native Change File Tracking (CFT) method. Cohesity File-

Runner facilitates the backing up of Nutanix Files mount targets by leveraging snapshots. The tool specifically identifies differences between two snapshot copies, allowing for incremental backups and minimizing redundancy.

In the CFT method, Cohesity leverages the Nutanix Files CFT API to identify the delta between two snapshot copies.

See: Plan and Prepare for Nutanix Files Protection

Plan and Prepare for Nutanix Files Protection

This topic provides information on the supported Nutanix Files versions and mount targets, minimum permissions, and the prerequisites to register and protect Nutanix Files with Cohesity cluster. For more information, see Register and Manage Nutanix Files Source.

Before registering Nutanix Files arrays with the Cohesity cluster and protecting mount targets and data, verify the prerequisites in this section.

Considerations

- Prism central and Prism element are supported.
- Backup and recovery of WORM mount targets are supported.
- SMB DFS Referral supported: The Nutanix Files SMB client follows the Distributed File System (DFS) referral. The Distributed File System has a collection of shares on different machines that serve as a file system namespace. If one machine does not have a particular file requested by the client, it is referred to check the other machine.
- You can protect DFS shares using the NFS V4.1 and V3 protocols.

Required ports

Ensure the following ports are open for Nutanix Files.

Protocol	Primary Port	Key Services
NFS	2049 (TCP/UDP)	File access (v3/v4)
SMB	111 (TCP/UDP)	Portmapper/RPCBIND
SMB	445 (TCP)	Direct Host/Microsoft-DS
TCP	9440	REST API calls.

Supported NFS and SMB versions

The following table lists the supported NFS and SMB versions for backup and restore:

Protocol	Version
NFS	V3, V4.1
SMB	V3

Minimum Required Permissions

Before registering a Nutanix Files Cluster as the source type, ensure you have the following permissions:

- For Prism Central local user: Needs Files Admin permission.
- For Prism Element local user: Needs Cluster Admin permission.

Supported Nutanix Files Versions

Component	Version
Nutanix Prism Element	
FSM	5.1.1
File Server	5.1.1
Nutanix Prism Central	
Files Manager	5.1 or higher

Register and Manage Nutanix Files Source

To protect your Nutanix Files data, you need to register your Nutanix Files clusters as a source on the Cohesity cluster. After registration, you can manage the source on the Cohesity cluster in the following ways:

- Update the Nutanix Files source configuration
- Refresh the source details
- Unregister the Nutanix Files source from the Cohesity cluster.

Note: Service Account and Active Directory authentication methods are not supported for registering Nutanix Files source.

Register Nutanix Files as a Source

Before you register Nutanix Files as a source on the Cohesity cluster, ensure that you:

1. Review the requirements, considerations, and best practices.
2. Perform the prerequisite tasks.

For more information, see Plan and Prepare for Nutanix Files Protection.

To register Nutanix Files as a source:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click **Register** on the top-right of the page and then select **NAS**.

   The Register NAS page appears.
3. From the NAS Source drop-down list, select **Nutanix Files**.
4. Enter the following information, and then click **Register**.

Field Name	Description
NAS Source	Select **Nutanix Files**.
Type	Select either **Prism Central** or **Prism Element**. Prism Element manages one Nutanix cluster, while Prism Central lets you control several clusters from one interface. A single Prism Central can contain multiple Prism Elements.
Host or IP Address	The Host or IP address of the Nutanix Files cluster to register.
Username	Username used to access the Nutanix Files cluster . Specify a user that has adequate privileges to perform actions on the Source. If you are providing a domain user account, then you must specify the user name in the following format: `fully_qualified_domain_name\username`. For example, `xydc.local\user1`. For more information, see Plan and Prepare for Nutanix Files Protection **Note:** The username used to register the Nutanix Files cluster is case sensitive.
Password	Enter the password for the specified user.

Field Name	Description
Back Up SMB mount targets	Whether Nutanix Files will back up SMB mount targets. If yes, provide the local or Active Directory user credentials that allow at least read access on the Nutanix Files cluster . **Note:** The user must have full control on the restore target during restore.
Filter IPs	Enable this option to filter the IP addresses of the Nutanix Files cluster. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the Nutanix Files cluster. To filter the IP addresses of the Nutanix Files cluster, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the Nutanix Files source through which the communication to the Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the Nutanix Files source through which the communication to the Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.

When registration finishes, a new entry appears on the Sources page. You can now create a Protection Group to backup Nutanix Files mount targets. For more information, see Backup Nutanix Files .

Update the Nutanix Files Source Configuration

You can update the Nutanix Files details that you had provided during the registration process with the latest Nutanix Files configuration.

To edit the Nutanix Files source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Nutanix Files source and select **Edit**.
3. In the Edit NAS page, update the respective configurations.
4. In the **Password** field, enter the password of the Nutanix Files user account, and then click **Update**.

Refresh the Nutanix Files Source

You can refresh the Nutanix Files configuration and fetch the latest changes on the Nutanix Files.

To refresh the Nutanix Files source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Nutanix Files source and select **Refresh**.

Unregister Nutanix Files Source

If you plan to stop taking backups of the Nutanix Files mount targets, you can unregister the Nutanix Files source from the Cohesity cluster. Before you unregister, you must delete all the Protection Groups associated with the Nutanix Files cluster. For more information, see Delete the Protection Group.

To unregister the Nutanix Files cluster on the Cohesity cluster:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Nutanix Files source and select **Unregister**.
3. Click **Unregister**.

Backup Nutanix Files

After you register the Nutanix Files cluster as a source on the Cohesity cluster, create a Protection Group to back up Nutanix Files. You can include one or more Nutanix Files NFS or SMB shares in a Protection Group. You can preserve ACLs on SMB shares, if required.

Select Nutanix Files mount targets for Backup

Cohesity supports backing up Nutanix Files read-write (RW) mount targets. The mount target can be an NFS export or an SMB share. Based on the selected mount target and protocol, you must provide additional details when creating a Protection Group.

Considerations:

- NFS only or SMB only mount targets are supported. Mixed mount targets are not supported.
- Supported versions of NFS are 3 and 4.1.
- SMB supported version is 3.
- SMB DFS Referral is supported.
- DFS shares are supported using the NFS V4.1 and V3 protocols.

User Access Credentials to Back Up SMB Shares

To back up SMB shares of any mount target type, you must provide the local or Active Directory user credentials that allow at least read access to the Nutanix Files SMB share. You can provide these credentials while registering the Nutanix Files cluster with the Cohesity cluster or later when creating the Protection Group. For more information, see Register and Manage Nutanix Files Source.

Protect Nutanix Files

In Cohesity, Nutanix Files mount targets are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

This topic provides information on how to protect Nutanix Files mount targets using a Protection Group and a Protection Policy.

Protect Read-Write mount targets with SMB Share or NFS Export

You can protect Nutanix Files RW mount targets using either SMB or NFS protocols.

To protect RW mount targets with either SMB or NFS protocol:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **NAS**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the NAS source that contains the mount targets you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it. For more information, see Register and Manage Nutanix Files Source.
4. From the Object list, select RW mount targets of NFS exports or SMB shares. You can filter the list using the **Protection Status** drop-down. Using the **Protocol** drop-down, filter the list for only SMB or NFS shares.
5. Click **Continue**.
6. To back up SMB shares, provide the user access credentials if you have not provided when registering the Nutanix Files source.
7. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
8. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Protect Nutanix Files.
9. Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
10. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, click **Protect**.
11. Optionally, from the **Protection Group** field, you can select an existing Protection Group. By selecting an existing Protection Group, the settings configured in the existing Protection Group will be applied to this Protection Group. To select an existing Protection Group, perform the following steps or else, go to step 13:

a. From the **Protection Group** field, enable **Existing Group**.

b. Select an existing Protection Group from the **Group** drop-down list.

c. Click **Protect**.

12. In the **Start Time** field, the default time zone is the browser's time zone. You can change the time zone of the job by selecting a different time zone.
13. If you need to change any of the additional settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Protect Nutanix Files.
14. Click **Protect**.

    Your new Protection Group is now active and running.

Additional Settings

Additional Settings

Additional Settings	Description
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
Use Nutanix Files CFT	Toggle on this option to use the vendor-native Change File Tracking (CFT) method for file discovery during an incremental backup. In this method, Cohesity leverages the Nutanix Files CFT API to identify the delta between two snapshot copies. It reduces the time taken to discover the files and folders to backup. If this option is toggled off, Cohesity uses the Cohesity's built-in CFT method. In this method, Cohesity uses its internal File-Runner to discover the delta between two snapshot copies. By default, this option is disabled. **Note:** Due to the limitations set by Nutanix Files, only one change list can be run at a time. If the Protection Group contains multiple shares to back up, Nutanix Files creates delta for each share serially.

Additional Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Pre & Post Scripts	Toggle **Enable** to run scripts on a Unix server before and/or after a Protection Group runs. If configured, scripts are run every time an object (or entity) is backed by a protection run. For example, if there are five mount/shares that are backed up as part of the NAS Protection Group and the Pre Script is configured, then five instances of the Pre Script is executed in parallel. You can distinguish between the scripts that are run on the different objects (entities) using the `COHESITY_BACKUP_ENTITY` environment variable. The Cohesity cluster populates this environment variable with the name of the backed object (or entity).
Skip Files on Errors	Toggled on by default. The Protection Group continues to run even if it encounters errors on files, such as permissions errors. If files are skipped, the protection run details page indicates a warning status and provides additional information. If toggled off, the Protection Group stops when it encounters an error.

Additional Settings	Description
File DataLock	Toggle on this option to preserve the AccessTime of the files and folders of the selected mount targets. The AccessTime provides the lock period of the files and folders that are applied when these files and folders are recovered to the Cohesity View. For more information on File DataLock, seeNAS File DataLock . Select one of the following **DataLock** modes: • **Compliance**: In this mode, files or folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files or folders. However, a root user can delete the files and folders. Only a user with Data Security role can edit the File DataLock settings in the Compliance mode. • **Enterprise**: In this mode, files or folders cannot be modified by any user before or after the lock retention period expiration. However, a root user can delete the files or folders before or after the lock retention period. An admin user or a user with a Data Security role can edit the File DataLock settings in the Enterprise mode. This mode is useful for giving storage administrators additional control over the file. You can further configure the File DataLock settings to: • Set the lock period for the new files and folders created on the recovered Cohesity View Provide the following information to set the lock period for the new files and folders that will be created on the recovered Cohesity View. These settings are not applicable for the existing NAS source files and folders recovered on the Cohesity View. • **Default Lock Period**: Provide the time period during which the files should remain locked. When the time period ends, files can be deleted but not changed. By default, the files are locked forever. • **Autolock Files**: By enabling this option, the files are automatically locked if they are not modified in the time period you specify. By default, this option is disabled. • **Manually Lock Using**: Select when the files are to be locked: • **Read-Only**: Select this option to lock files when the file permission is set to read only, for example; by setting the Windows file properties attribute or using

Additional Settings	Description
	the Linux `chmod` command. By default, the files are locked when the file permission is set to read only. • **FutureATime**: Select this option to lock files when the last access time (atime) is set, for example, by using Windows `PowerShell` commands or the Linux touch command. Optionally you can set a **Minimum Retention** and/or a **Maximum Retention** period to ensure that a file's last access time cannot be set outside this period. A request to do so results in a permission denied error if manual locking mode is FutureATime. Otherwise, the request will be silently rejected and the file will be locked for the minimum retention duration. • Override the locking period of the existing files and folders recovered on the Cohesity View: The following option can be applied to all the existing files and folders that will be recovered on the Cohesity View. • **Override (Lock all files until a specific date)**: Enable this option and select a date to override the locking period of the existing files and folders. When you recover the files to the Cohesity View, Cohesity applies the farthest lock period between the override lock period you specify and the source file's lock period. For example; • The lock period of the source file is 09/09/2023 and the override lock period you specified is 09/08/2022. When you recover the file to the Cohesity View, then the lock period of the source file is applied to the recovered file, which is 09/09/2023. • The lock period of the source file is 09/09/2022 and the override lock period you specified is 09/08/2024. When you recover the file to the Cohesity View, then the override lock period you specified is applied to the recovered file, which is 09/08/2024. In **Compliance** mode, a user with Data Security privileges can extend the override date into the future, but no user can shorten the date.

Additional Settings	Description
	In **Enterprise** mode, a user with admin or Data Security privileges can extend or shorten the override date. **Note:** • You cannot enable or disable File DataLock once the Protection Group is created. • If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed. • You can delete Protection Groups and snapshots having File DataLock enabled. File DataLock also preserves the ATimes of the hardlinks and symlinks.
Exclusions and Inclusions	Toggle on this option if you want to exclude or include locations. By creating exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories and therefore minimize the disk space used to store the data. All the files and folders are included for by default for protection. For more information, seeExclusions and Inclusions for NAS Protection Group.
Indexing	This option is toggled on by default to index the data of the Protection Group. Indexing improves files and folders search. Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.

Additional Settings	Description
Quiet Times	This option is available only if the selected policy has at least one quiet time period. Toggle the **Stop in-progress runs when quiet time starts** option to pause or abort the runs during quiet time. If this option is toggled off, after a protection run starts, execution continues even when the quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period. • **Pause Runs** - Select this option to pause the currently executing protection runs and resume the runs when quiet time ends. • **Abort Runs** - Select this option to abort the currently executing protection runs when the quiet time period specified for the Protection Group starts. **Note:** If the quiet time window for multiple protection runs is scheduled at the same time, then when the protection run is resumed, prioritization of the protection run does not work. 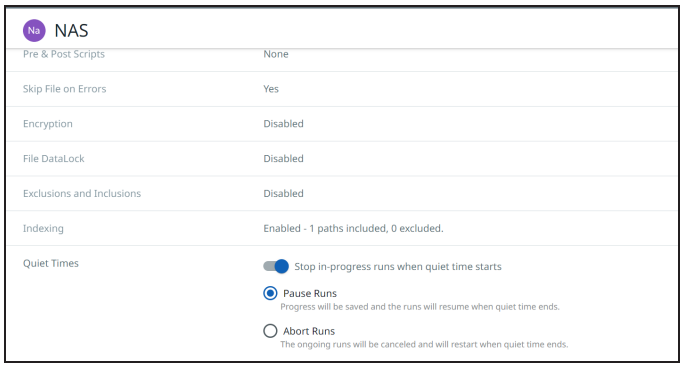

Additional Settings	Description
Alerts	Select one or more of the following **Alerts On** settings if you want alerts to be created for the following triggers: • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • Success: Create an informational alert when the Protection Group completes successfully. Emails are not sent when informational alerts are created. To send emails when alerts are triggered for the Protection Group, in the **Email Recipients** field, add the email addresses of the recipients.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: a. High b. Medium c. Low
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

A Protection Policy is a reusable set of settings. It enables you to define how and when Nutanix Files data is protected, replicated, and archived. A standard policy provides various settings that you can configure for a Protection Group.

Create a Protection Policy with the settings listed in the Protect Nutanix Files. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Nutanix Files

The following table lists the applicable policy settings for Nutanix Files:

Field	Description
Primary Copy	Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster. Select any external target for performing CloudArchive Direct wherein the primary copy of the data is stored on the external targets. For more information, see CloudArchive Direct.
DataLock	If applicable, add a DataLock for compliance and regulatory requirements, to ensure that your protected data, including local backups, replication, and archives, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users, including those who have the Data Security role in Cohesity, from modifying or deleting any Snapshots that were generated by theProtection Groups that use this Policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the Protection Group fails. The default time between retries is five minutes. You can customize the number of retries and how long to wait between each attempt.
Periodic full backup	Use this option to choose when to run a periodic full backup.

Field	Description
Quiet Times	Quiet Times defines time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturday or Sunday. Quiet Times only prevents new protection runs from starting during the specified time period. The protection runs that started before the quiet time period will be dependent on the quiet time settings configured in the additional settings of the protection Group. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	To replicate snapshots to another, click the replication icon from the floating menu. The copied snapshots are stored in the storage domain of the remote cluster.
Archival	To add an archive schedule, click the archival icon from the floating menu.

Use an Existing Protection Policy

You can use an existing Protection Policy protecting Nutanix Files. If you are selecting an existing protection policy, then the following options are ignored if they are already configured in the protection policy:

- **Log Backup (Databases)**
- **Continuous Data Protection**
- **Add BMR Backup (Physical Server)**
- **CloudSpin**

NAS File DataLock

Many organizations now rely on WORM (Write Once Read Many) data storage such as Nutanix Files, NetApp Snaplock, and Isilon Smartlock volumes to retain their files and folders in unmodified form for compliance and regulatory purposes and to protect the files from accidental deletion or modification.

When you backup these WORM-capable volumes to the Cohesity cluster, it is essential to have the ability to easily lock the files and folders of the volumes while protecting themselves from accidental or malicious attempts to delete or modify these files. The backup should also provide a strong guarantee for data non-mutability and non-erasability for a certain period of time.

This capability is achieved using Cohesity's File DataLock feature.

FileDataLock is a Cohesity feature that preserves the WORM attributes of any files and folders created in the Cohesity View. You can apply the File DataLock feature for NAS backup (NFS exports and SMB shares) while creating a NAS Protection Group.

With File DataLock enabled in a NAS Protection Group, Cohesity preserves the future AccessTime (ATime) of the files and folders. The AccessTime provides the lock period for files and folders. You can also override this lock period in the File DataLock settings of the NAS Protection Group. This lock period is applied to these files and folders when they are recovered to the Cohesity View. This prevents the users or administrator from modifying or deleting these locked files until the locking period expires.

You can also set lock periods for the new files and folders created on recovered Cohesity View in the File DataLock settings of the NAS Protection Group.

When you recover File DataLock enabled storage volumes to the Cohesity View, the View details page will display a lock icon for the files and folders having File DataLock.

File DataLock is supported for all types of NAS sources including NetApp 7-mode and is agnostic of the NAS protocol.

For information on how to enable File DataLock, see Protect NAS.

File DataLock Modes

You can use the File DataLock feature in one of the two modes:

- **Compliance mode**: The WORM files are protected at the disk level. This mode implements strict regulatory requirements for data retention and complies with SEC Rule 17a-4. In this mode, files and folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files and folders. However, a root user can delete the files and folders. Only a user with a Data Security role can edit the File DataLock settings in the Compliance mode.
- **Enterprise mode**: The WORM files are protected at the file level. This mode does not implement strict regulatory requirements for data retention. In this mode, files and folders cannot be modified by any user before or after the lock retention period expiration. However, a root user can delete the files and folders during or after the lock retention period. An admin user or a user with the Data Security role can edit the File DataLock settings in the Enterprise mode

> **Note:** An admin user or a user with the Data Security role can edit the File DataLock settings and switch from the Enterprise mode to Compliance mode. But switching from Compliance mode to Enterprise mode is not supported.

Things to remember for File DataLock

- You cannot enable or disable File DataLock once the View is created.
- If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed.
- You can delete Protection Groups and snapshots having File DataLock enabled.

  File DataLock also preserves the ATimes of the hardlinks and symlinks.

- **Recovery**
  - Recovering a File DataLock enabled backed-up volume as a Cohesity View will preserve ATimes and the locking period for the file.
  - Recovering a File DataLock enabled volume to the original location with 'Overwrite existing files/folder' as the recovery option will encounter an error if the files and folders are read-only. However, the ATime and other attributes of other files and folders of the source volumes are preserved.
  - Recovering a File DataLock enabled volume to an alternate NAS recovery will preserve ATimes and other attributes of files and folders of the source volumes.
  - Cloning a recovered View having File DataLock enabled will not inherit the File DataLock properties.
- **Replication**

  When you replicate to a remote Cohesity cluster, the File DataLock settings are also replicated.
- **Archival**

  File DataLock settings are preserved when you recover volume from an archived snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.
- **Cloud Archive Direct**

  File DataLock settings are preserved when you recover from a Cloud ArchiveDirect snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.

Exclusions and Inclusions for NAS Protection Group

In a NAS Protection Group, all the files and folders are included by default for data protection. While creating a Protection Group, you can toggle on **Exclusions and Inclusions** from the **Advanced Settings** options if you want to exclude or include locations. By creating the exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories. As a result, you can minimize the disk space used to store the data.

> **Note:**
> - To add an exclusion or inclusion rule, you must prefix a forward slash ('/') or suffix an asterisk ('*') to the path or to a particular file within the protected object. For example, '**/test**' or '***.txt**'.
> - In exclusion or inclusion rules, file names and folder names are case-sensitive, which means '**test.txt**' and '**Test.txt**' are considered separate files. Use the correct capitalization of file or folder names to ensure that the rules recognize your file or folder.

Add Inclusion

From the Additional Settings options on the Protection Group page, toggle on **Exclusions and Inclusions**. Click **Add Inclusion** to include a particular path or a particular file within the protected object.

For example, consider four directories - **test**, **test1**, **test2**, and **test3** under the protected object with path `/ifs/TestShare1/Folder1`. The table below lists the input types for an inclusion list of folders and their respective outcome:

Input	Outcome
/	Everything inside the protection path is backed up if exclusions are not defined.
/test	When the inclusion path is not suffixed by '/', Protection Group will include all folders starting from the protection path. In this example, /test will behave as /test* and Protection Group will include '**test**', '**test1**', '**test2**', and '**test3**' from the protection path for backup.
/test/	When the inclusion path is suffixed by '/', it will only include the specific folder. In this example, /test/ behaves as /test/* and Protection Group will only include '**test**' from the protection path for backup.

- Do not provide the folder name along with the backup path while including files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not include any folder.
- If exclusions are defined, they will take precedence over inclusions.

Add Exclusion

From the Additional Settings options on the Protection Group page, toggle on **Exclusions and Inclusions**. Click **Add Exclusion** to exclude a particular path or a particular file within the protected object. You can also specify regular expressions for excluding files in the format `regex:<regex pattern>`.

For example, consider there are three directories - **Test**, **Test1**, and **Folder1** and a file, **TestFile.tmp** under the protected object with path **/ifs/TestShare1**.

- The **Test1** folder has **TestFile1.tmp**, **TestFile2.tmp** files.
- The **Folder1** folder has **TestFile3.tmp**, **TestFile4.tmp files**, and **SubFolder1** which further has a **Test** folder with **File.txt**, **File1.txt** files in it.

The following figure illustrates the directory structure under the protected object with path **/ifs/TestShare1**.

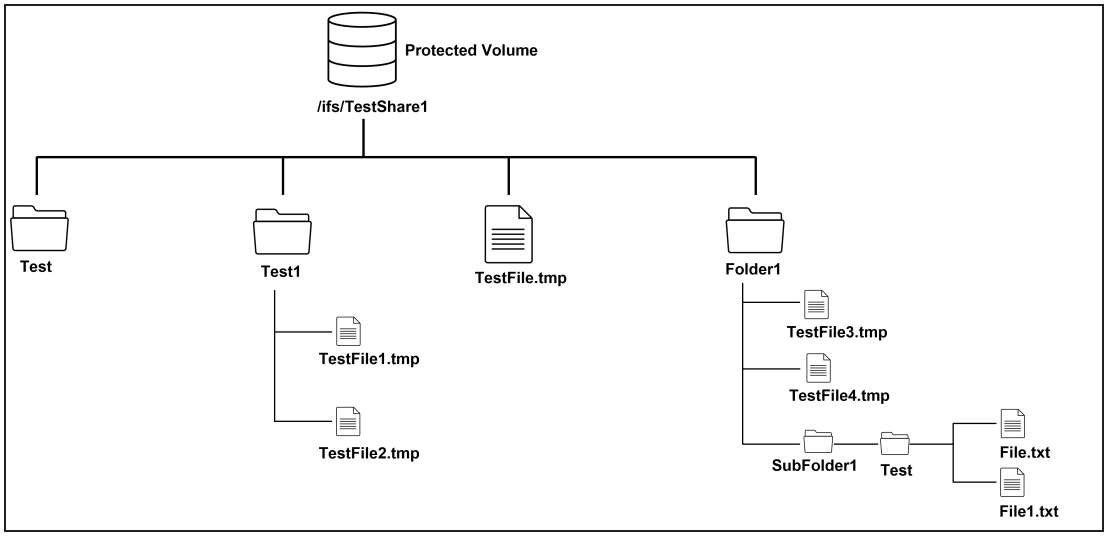

The table below lists the input types for an exclusion list of folders and their respective outcome:

Input	Description	Outcome
/Test	When the exclusion path is not suffixed by '/', the Protection Group will exclude all files and folders with **Test** as the prefix. For exclusions, **/Test** behaves as **/Test***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes **Test** and **Test1** folders, and **TestFile1.tmp** file from the protection path. **TestFile3.tmp** and **TestFile4.tmp** files in the **Folder1** folder will not be excluded.
/Test/	When the exclusion path is suffixed by '/', the Protection Group excludes only the files and folders with the name **Test**. For exclusions, **/Test/** behaves as **/Test/***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes the **Test** folder of the parent directory. The **Test** sub-folder in **/Folder1/SubFolder1** will not be excluded.

Input	Description	Outcome
*/Test	Excludes all the files and directories with the name '**Test**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes the **Test** folder of the parent directory and the **Test** sub-folder in **/Folder1/SubFolder1**.
*.tmp	Excludes all the files that with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.
regex:.*\.tmp	Excludes the files and folders with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.
regex:/Folder1/.*/Test/.*\.txt$	Excludes all files with **.txt** extension from all the **Test** folders directory in **/Folder1/**.	Excludes **File.txt** and **File1.txt** files in the **/Folder1/SubFolder/Test** folder.

Note: Do not provide the folder name along with the backup path while excluding files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not exclude any folder.

Cohesity automatically excludes the following NetApp system files:

- .vtoc_internal and .bplusvtoc_internal files
- .copy-offload directory and .tokens file

Considerations

Cohesity does not support inclusion or exclusion of paths in:

- an IsilonProtection Group with Changelist enabled.
- a NetApp SnapDiff-basedProtection Group.

Manage the Nutanix Files Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirement.

> **Note:** You cannot update the registered Nutanix Files cluster and the assigned Storage Domain in the Protection Group settings.

Add Nutanix Files mount targets to the Protection Group

You can add new Nutanix Files mount targets to the existing Protection Group for backup. The newly added mount targets are backed up from the next protection run and it will be a full backup.

To add Nutanix Files mount targets to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Select the Nutanix Files mount targets you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added Nutanix Files mount targets will be backed in the next scheduled protection run.

Remove Nutanix Files mount targets from the Protection Group

From the **Edit Protection** page, you can remove Nutanix Files mount targets from the Protection Group that you do not want to back up in the next protection run.

To remove Nutanix Files mount targets from the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Unselect the Nutanix Files mount targets that you want to remove from the Protection Group.
6. Click **Save Selection** > **Save**.

The removed Nutanix Files mount targets will not be backed in the next scheduled protection run.

Enable Auto Protect

The Auto Protect option enables you to automatically protect the new Nutanix Files mount targets that are added to the selected parent Nutanix Files cluster. However, you cannot enable auto-protect at the cluster-level containing Read-Write mount targets.

To enable auto protect:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click **Edit Objects**.
5. Click the shield icon to enable Auto Protect.
6. Click **Save Selection** > **Save**.

When this option is enabled on the parent Nutanix Files cluster, all the Nutanix Files mount targets that are added to the parent Nutanix Files cluster in the future are automatically added to the Protection Group and are protected from the next protection run.

For example, you create a Protection Group that backs up a Nutanix Files cluster. The Protection Group runs for two weeks and then a new mount target is added to the Nutanix Files cluster. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new mount target is also backed up even though the new mount target has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list, and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirement. In the Protection Policy, you can edit the settings listed in the Create or Use an Existing Protection Policy.

To the policy settings:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Edit**.

4. On the **Edit Protection** page, in the **Policy** field, click **Edit**.

   The **Edit Policy** page appears.

5. Edit the settings listed in the Create or Use an Existing Protection Policy and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, from the drop-down list in the **policy** field, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to back up Nutanix Files mount targets. It provides you options to back up all the objects or only selected objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- Is configured with replication targets, you can select those targets.
- Is configured with archival targets, you can select those targets.

If the last protection run contains failed objects, you can backup only the failed objects.

To start a protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
   a. In the **Backup** drop-down list, select any one of the following:
      - **All objects in the protection group**
      - **All objects that failed last run**: Available when the last backup run contains failed objects.

- **All objects that succeeded the last run**: Available when the last backup run does not contain failed objects.
- **Selected objects in the protection run**

b. In the **Backup Type** drop-down list, select the backup type.

c. In the **Replication** section, select the replication target, and the retention duration in the **Replicate** drop-down list. Select **All objects that succeeded last run**, if you want to replicate the data only after a successful protection run. This option is not available when the last protection run contains failed objects.

d. In the **Archive** section, select the archive target, and the retention duration. You can register an external target, if required.

5. Click **Run**.

Pause Protection Runs Instantly

You can instantly pause a protection run of the Protection Group that you have created to back up the Nutanix Files sources.

Perform the following steps to pause a currently executing protection run of a Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, click the Nutanix Files Protection Group that has been created.
3. Click the action icon next to the protection run and click **Pause Run**.

Pause Future Protection Runs

You can stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the action icon next to the Protection Group and select **Resume**.

Delete the Protection Group

Cohesity provides the following ways to delete the Protection Group associated with Nutanix Files:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Protection Group are not deleted
- **Delete Protection Group and Snapshots**: Delete the Protection Group and all Snapshots, any local Snapshots captured by this Protection Group and stored on this Cluster are deleted. In addition, any archived Snapshots stored on external targets that were captured by this Protection Group are also deleted.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Nutanix Files.
3. Click the action icon next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Group Run.

Dell EMC Isilon

Isilon is a network-attached storage platform offered by Dell EMC and provides a cluster-based storage array based on industry-standard hardware. Using the policy-based data protection solution from Cohesity, you can protect the data stored on Isilon. This section provides information on how to back up and recover Isilon data using Cohesity.

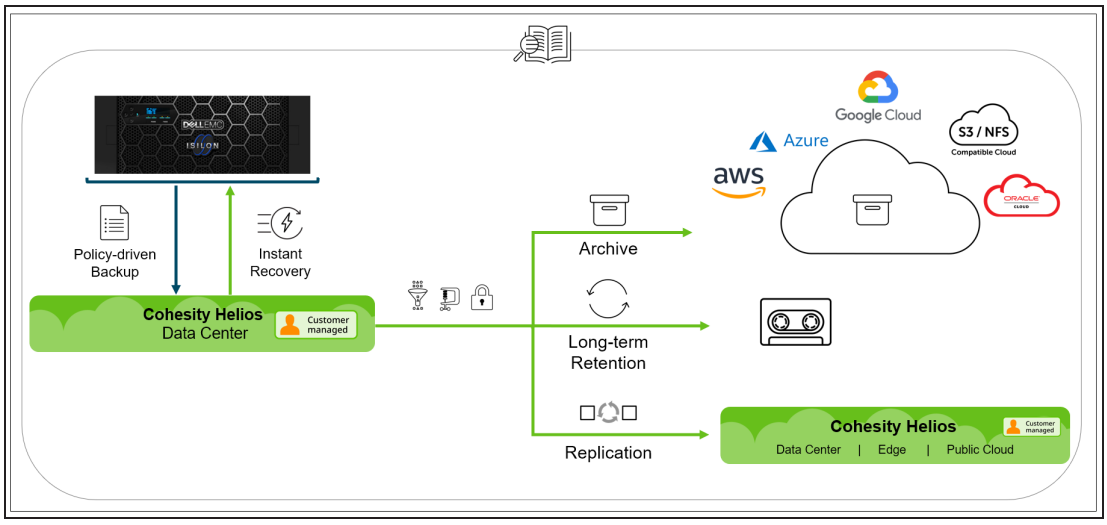

Plan and Prepare for Isilon Protection

Before you register your Isilon cluster with the Cohesity cluster and protect the volumes and data, ensure that:

- Cohesity supports the Isilon cluster versions and volumes you plan to protect.
- You have the minimum permissions required to register and manage your Isilon cluster.
- The required firewall ports are open to enable communication between the Cohesity cluster and the Isilon cluster.
- You meet the prerequisites to register and manage your Isilon cluster.
- You have reviewed and understood the considerations of Isilon cluster backup and recovery.

Support Matrix

Before you register your Isilon cluster with the Cohesity cluster, ensure that Cohesity supports the Isilon versions and volumes you want to protect. For more information, see Supported Software.

Supported NFS and SMB versions

The following table lists the supported NFS and SMB versions for backup:

Protocol	Version	Notes
NFS	NFSv3 NFSv4.1	NFS v4.1 Kerberos authentication is supported.
SMB	SMB v2.x, v3	

Minimum Permissions

The user account that will be used for registering the Isilon cluster with the Cohesity cluster does not require administrator privileges on the Isilon cluster. The user account only requires the following permissions necessary to back up and restore your Isilon data:

Access-level	Command	Description	Protocol
ReadOnly	Platform API	For access to Isilon's APIs.	SMB/NFS
	Auth	To verify users and passwords.	
	Cluster	To obtain cluster identity and settings.	
	Network	To obtain the network interfaces.	
	SMB	To read the settings in the SMB server.	
Read/Write	Job Engine	To read and write Changelist jobs.	
	Snapshot	To fetch, create, and delete snapshots for shares and exports.	
	NFS	To read and write settings to and from the NFS server.	

Ports Used for Communication

Ensure that TCP/UDP ports 300, 302, 304, 443, 445, 8080, 111, and 2049 are open in the firewall between the Cohesity cluster and your Isilon device. For more information, see Manage Firewall Ports.

Prerequisites

> **Note:** The sub-settings for Snapshot Self-Service are not mandatory. However, when the Global Visibility and Access is enabled, the sub-settings for NFS Settings, SMB Settings, and Local Settings are also enabled by default. These sub-settings for Snapshot Self-Service are not mandatory for Cohesity and can be disabled if you choose to do so.

Ensure to meet the following prerequisites to register and manage your Isilon cluster:

- Cohesity leverages SnapshotIQ to perform the backup of the Isilon source. Therefore, ensure to enable the SnapshotIQ license with the following settings on the Isilon cluster:

    - Enable Snapshot Service
    - Enable global visibility and access
- The user provided during the Isilon source registration must be added to the SMB shares as access with Read-Only permission.

Considerations

- The NAS storage FQDN must be resolvable. Even if registered by IP, SMB restore workflows depend on the FQDN for repeated SMB share connections. Ensure that the DNS or cluster host mapping resolves the storage FQDN.
- The Cohesity Dashboard estimates the total size of protected and unprotected Isilon shares by multiplying the number of shares by about 3.5 GiB. Cohesity cannot determine the actual size because each share is a subdirectory of the root /ifs directory.
- For Isilon OneFS 9.2.x or earlier versions, the maximum filename length or full path length limits to 255 bytes.
- Root permissions are required to set the ACLs. To set the permissions, you can perform the following:
    - for the share that belongs to System Access Zone, add "Backup files from /ifs" and "Restore files to /ifs" privileges for the role to which the user or group belongs.
    - for non-System Access Zone, set the share permissions for the user as "Run as root".
- Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.
- If the SMB credentials on the source change during an active backup, Isilon SMB backup jobs automatically pause and display a pause note: "Auto-paused due to credential/logon failure."

  The jobs resume after you update the SMB credentials on the Cohesity cluster. This behavior prevents loss of overall backup progress. Any missed files are backed up in the next incremental run.

Register and Manage Isilon Source

To protect your Isilon data, you must register your Isilon cluster as a source on the Cohesity cluster. After registration, you can manage the source on Cohesity in the following ways:

- Update the Isilon source configuration
- Refresh the source details
- Unregister the Isilon cluster source from the Cohesity cluster.

Register Isilon Cluster as a Source

Before you register an Isilon cluster as a source on the Cohesity cluster, ensure that you:

1. Review the requirements, considerations, and best practices.
2. Perform the prerequisite tasks.

For more information, see Plan and Prepare for Isilon Protection.

To register Isilon cluster as a source:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click **Register** on the top-right of the page and then select **NAS**.

   The Register NAS page appears.
3. From the NAS Source drop-down list, select **Isilon (Cluster)**.
4. Enter the following information and then click **Register**.

Field Name	Description
Host or IP Address	Provide the management IP address or hostname of the Isilon cluster.
Username	Provide the user account that you configured with the required permissions on the Isilon cluster. If you are providing a domain user account, then you must specify the user name in the following format: `fully_qualified_domain_name\username`. For example, `xydc.local\user1`.
Password	Provide the password of the user account.

Field Name	Description
Back Up SMB Volumes	If you want to protect Isilon SMB shares, provide the local or Active Directory user credentials that allow read access to the Isilon SMB share you want to back up. Ensure to provide the Active Directory user credentials in the `domain\user` format. **Note:** • You can assign the local or AD user to the built-in "BackupAdmin" role to allow the user to read the SMB data for backup without modifying the access control lists (ACLs). • To provide access at the share level, grant the "Run as root" and "Full Control" permissions at the share level.
Filter IPs	Enable this option to filter the IP addresses of the Isilon cluster. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the Isilon cluster. To filter the IP addresses of the Isilon cluster, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the Isilon source through which the communication to the Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the Isilon source through which the communication to the Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.

For more information, see Protect Dell EMC Isilon with Cohesity DataProtect.

When registration finishes, a new entry appears on the Sources page. You can now create a Protection Group to backup Isilon volumes. For more information, see Protect Isilon.

After registering the Isilon source, you can choose the **Configure** option in the **Source** to configure the **Static Network Pool** and **Dynamic Network Pool**, as well as the SmartConnect for each Isilon Access Zone.

Update the Isilon Cluster Source Configuration

You can update the Isilon cluster details that you provided during the registration process with the latest Isilon configuration.

To edit the Isilon cluster source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Isilon cluster source and select **Edit**.
3. In the Edit NAS page, update the respective configurations.

4. In the **Password** field, enter the password of the Isilon user account, and then click **Update**.

Refresh the Isilon cluster Source

You can refresh the Isilon cluster configuration and fetch the latest changes on the Isilon cluster.

To refresh the Isilon cluster source configuration:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Isilon cluster source and select **Refresh**.

Unregister Isilon cluster

If you plan to stop taking backups of the Isilon cluster volumes, you can unregister the Isilon cluster source from the Cohesity cluster. Before you unregister, you must delete all the Protection Groups associated with the Isilon cluster. For more information, see Delete the Protection Group

To unregister the Isilon cluster server on Cohesity cluster:

1. In the Cohesity Dashboard, select **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Isilon cluster source and select **Unregister**.
3. Click **Unregister**.

Isilon SmartConnect

Isilon SmartConnect is a feature that offers load balancing and failover capabilities for connections to an Isilon cluster. It effectively manages requests from multiple connections, manages DNS requests, and oversees multiple IP addresses.

SmartConnect is a licensed software module of the EMC Isilon OneFS, designed to enhance performance and availability by enabling intelligent client connection load balancing and failover support. If the Isilon cluster has SmartConnect enabled, it can be used for backup operations.

To utilize SmartConnect, you need to configure it after registering the Isilon source. You must register the Isilon Adapter Source using the SmartConnect IP (SIP), which provides multiple addresses for the Cohesity cluster to access

SmartConnect can be configured under **Cluster Management** > **Network Configuration**, where you can view or edit the subnet settings. Cohesity supports using the SmartConnect Fully Qualified Domain Name (FQDN) to discover and connect to the Isilon cluster.

Configure Isilon SmartConnect

To configure Isilon SmartConnect:

1. Navigate to **Data Protection > Sources**.
2. Select the registered Isilon source, click the 3 dots on the right, and click **Configure**.
3. On the **Configure Isilon Source** page, select the **Access Zone** you want to use.
4. From the **Static Network Pool**, select the desired Cohesity network pool to use for communicating with Isilon.
5. Enter or confirm the **SmartConnect FQDN** of the Isilon cluster.
6. Enable the **Use Smart Connect** option.
7. Click **Save**.

   The **SmartConnect FQDN** provided is used to do the backup.

Backup Isilon

After registering your Isilon cluster as a source on the Cohesity cluster, you can create a Protection Group to protect the Isilon volumes.

The volume can be of NFS export only, SMB share only, or of both NFS and SMB (mixed) type. To back up SMB shares of any volume type, you must provide the local or Active Directory user credentials that allow at least read access to the Isilon SMB share. You can provide these credentials while registering the Isilon cluster with the Cohesity cluster or later when creating the Protection Group. For more information, see, Register and Manage Isilon Source.

You can backup one or more Isilon NFS mount points, SMB shares, or mixed volumes in a Protection Group while preserving ACLs and Extended attributes on SMB.

Protect Isilon

In Cohesity, Isilon volumes are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

You can protect Isilon volumes with only SMB share, NFS export, or the volumes that are exposed to both NFS export and SMB share (mixed volumes). For Isilon mixed volumes, Cohesity will back up both NFS POSIX ACLs and SMB access control lists (ACLs).

> **Note:** Cohesity recommends a first full and incremental forever backup approach to back up your NAS sources.

To protect Isilon volumes:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **NAS**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the registered Isilon source that contains the volumes you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it. For more information, seeRegister and Manage Isilon Source.
4. From the **Objects** list, select Isilon volumes NFS exports, SMB shares, or the volumes that are exposed to both NFS and SMB protocols.
5. Click **Continue**.
6. To back up SMB shares, provide the user access credentials if you have not provided when registering the Isilon source.
7. If you have selected the volumes that are exposed to both NFS and SMB protocols, then select the backup preference mode of the volume you want to backup.
8. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
9. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Use an Existing Protection Policy.
10. Select the appropriate Storage Domain. For more information, see Manage Storage Domains.
11. If you want to change or configure any of the additional settings, then select **More Options** and perform the below steps or else, click **Protect**.
12. If you want to add the selected objects to an existing Protection Group, then perform the following steps or else, go to step 13:
    a. From the **Protection Group** field, enableProtection Group **Existing Group**.
    b. Select an existing one from the **Group** drop-down list.
    c. Click **Protect**.

    The objects will be protected based on the policy settings applied to the selected Protection Group.
13. In the **Start Time** field, specify the time to run the Protection Group. The default time zone is the browser's time zone.

> **Note:** This option is available only if the selected policy is set to Backup Daily.

14. If you need to change any of the additional settings on the Protection Group page, click the down arrow icon next to **Additional Settings** and click **Edit** on the right. For more information, see Additional Settings.
15. Click **Protect**.

    Your new Protection Group is now active and running.

Additional Settings

Additional Settings	Description
End Date	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Pre & Post Scripts	Toggle **Enable** to run scripts on a Unix server before and/or after a Protection Group runs. If configured, scripts are run every time an object (or entity) is backed by a protection run. For example, if there are five mount/shares that are backed up as part of the NAS Protection Group and the Pre Script is configured, then five instances of the Pre Script are executed in parallel. You can distinguish between the scripts that are run on the different objects (entities) using the `COHESITY_BACKUP_ENTITY` environment variable. The Cohesity cluster populates this environment variable with the name of the backed object (or entity).

Additional Settings	Description
Skip Files on Errors	Toggled on by default. The Protection Group continues to run even if it encounters errors on files, such as permissions errors. If files are skipped, the protection run details page indicates a warning status and provides additional information. If toggled off, the Protection Group stops when it encounters an error.
Use Isilon Change Lists	Toggle on this option to use the vendor-native Change File Tracking (CFT) method for file discovery during an incremental backup. In this method, Cohesity leverages the Isilon CFT API to identify the delta between two snapshot copies. It reduces the time taken to discover the files and folders to backup. If this option is toggled off, Cohesity uses the Cohesity's built-in CFT method. In this method, Cohesity uses its internal File-Runner to discover the delta between two snapshot copies. By default, this option is disabled. > **Note:** Due to the limitations set by Isilon, only one change list can be run at a time. If the Protection Group contains multiple shares to back up, Isilon creates delta for each share serially.
Encryption	Toggle on this option to enable in-flight encryption of data that are backed up to the Cohesity cluster. This option is disabled by default. To use the encryption functionality for NFS volumes, you must join the Cohesity cluster to the active directory as the Kerberos server or add the Kerberos provider in the Cohesity cluster Access Management section. This option is disabled by default. > **Note:** Encrypting in-flight data will have a minor impact on the performance.

Additional Settings	Description
File DataLock	Toggle on this option to preserve the AccessTime of the files and folders of the selected volumes. The AccessTime provides the lock period of the files and folders that are applied when these files and folders are recovered to the Cohesity View. For more information on File DataLock, see NAS File DataLock. Select one of the following **DataLock** modes: • **Compliance**: In this mode, files or folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files or folders. However, a root user can delete the files and folders. Only a user with the Data Security role can edit the File DataLock settings in the Compliance mode. • **Enterprise**: In this mode, files or folders cannot be modified by any user before or after the lock retention period expiration. However, a root user can delete the files or folders before or after the lock retention period. An admin user or a user with the Data Security role can edit the File DataLock settings in the Enterprise mode. This mode is useful for giving storage administrators additional control over the file. You can further configure the File DataLock settings to: • Set the lock period for the new files and folders created on the recovered Cohesity View Provide the following information to set the lock period for the new files and folders that will be created on the recovered Cohesity View. These settings are not applicable for the existing NAS source files and folders recovered on the Cohesity View. • **Default Lock Period**: Provide the time period during which the files should remain locked. When the time period ends, files can be deleted but not changed. By default, the files are locked forever. • **Autolock Files**: By enabling this option, the files are automatically locked if they are not modified in the time period you specify. By default, this option is disabled. • **Manally Lock Using**: Select when the files are to be locked: • **Read-Only**: Select this option to lock files when the file permission is set to read only, for example; by setting the Windows file properties attribute or using the Linux `chmod` command. By default, the files are locked when the file permission is set to read only. • **FutureATime**: Select this option to lock files when the last access time (atime) is set, for example, by using Windows `PowerShell` commands or the Linux touch command. Optionally you can set a **Minimum Retention** and/or a **Maximum Retention** period to ensure that a file's last access time cannot be set outside this period.

Additional Settings	Description
	A request to do so results in a permission denied error if manual locking mode is FutureATime. Otherwise, the request will be silently rejected and the file will be locked for the minimum retention duration. • Override the locking period of the existing files and folders recovered on the Cohesity View: The following option can be applied to all the existing files and folders that will be recovered on the Cohesity View. • **Override (Lock all files until a specific date)**: Enable this option and select a date to override the locking period of the existing files and folders. When you recover the files to the Cohesity View, Cohesity applies the farthest lock period between the override lock period you specify and the source file's lock period. For example; • The lock period of the source file is 09/09/2023 and the override lock period you specified is 09/08/2022. When you recover the file to the Cohesity View, then the lock period of the source file is applied to the recovered file, which is 09/09/2023. • The lock period of the source file is 09/09/2022 and the override lock period you specified is 09/08/2024. When you recover the file to the Cohesity View, then the override lock period you specified is applied to the recovered file, which is 09/08/2024. In **Compliance** mode, a user with the Data Security privileges can extend the override date into the future, but no user can shorten the date. In **Enterprise** mode, a user with admin or the Data Security privileges can extend or shorten the override date. **Note:** • You cannot enable or disable File DataLock once the Protection Group is created. • If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed. • You can delete Protection Groups and snapshots having File DataLock enabled. File DataLock also preserves the ATimes of the hardlinks and symlinks.

Additional Settings	Description
Filter IPs	Toggle on option to filter the IP addresses of the Isilon cluster. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the Isilon cluster at the Protection Group level. To filter the IP addresses of the Isilon cluster, select one of the following options: • **Allow IPs**: Select this option and specify the IP addresses of the Isilon cluster source through which the communication to the Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. • **Deny IPs**: Select this option and specify the IP addresses of the Isilon cluster source through which the communication to the Cohesity cluster must not happen. You can provide the IP addresses in a comma-separated list or in a CIDR format. > **Note:** If you have specified the IP filtering at the source level also, then the IP filtering specified at the Protection Group level will take precedence over the IP filtering specified at the source level.
Exclusions and Inclusions	Toggle on this option if you want to exclude or include locations. By creating exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories and therefore minimize the disk space used to store the data. All the files and folders are included for by default for protection. For more information, see Exclusions and Inclusions for NAS Protection Group. > **Note:** Cohesity does not support inclusion or exclusion of paths in an Isilon Protection Group with the **Use Isilon Change Lists** option enabled.
Indexing	This option is toggled on by default to index the data of the Protection Group. Indexing improves the files and folders search. Cohesity performs incremental indexing where only the data that has changed since the previous snapshot is scanned and indexed.

Additional Settings	Description
Quiet Times	This option is available only if the selected policy has at least one quiet time period. Toggle the **Stop in-progress runs when quiet time starts** option to pause or abort the runs during quiet time. If this option is toggled off, after a protection run starts, execution continues even when the quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period. • **Pause Runs** - Select this option to pause the currently executing protection runs and resume the runs when quiet time ends. • **Abort Runs** - Select this option to abort the currently executing protection runs when the quiet time period specified for the Protection Group starts. **Note:** If the quiet time window for multiple protection runs is scheduled at the same time, then when the protection run is resumed, prioritization of the protection run does not work.
Alerts	Select one or more of the following **Alerts On** settings if you want alerts to be created for the following triggers: • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • Success: Create an informational alert when the Protection Group completes successfully. Emails are not sent when informational alerts are created. To send emails when alerts are triggered for the Protection Group, in the **Email Recipients** field, add the email addresses of the recipients.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection run exceeds the ability to process the protection run, this will be the priority of implementation of the protection run: a. High b. Medium c. Low

Additional Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
Description	Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

A Protection Policy is a reusable set of settings. It enables you to define how and when Isilon data is protected, replicated, and archived. A standard policy provides various settings that you can configure for a Protection Group.

Create a Protection Policy with the settings listed in the Policy Settings for Isilon. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Isilon

The following table list the applicable policy settings for Isilon:

Field	Description
Primary Copy	Select **Local** from the drop-down to keep the primary copy data on the Cohesity cluster. Select any external target for performing CloudArchive Direct wherein the primary copy of the data is stored on the external targets. For more information, see CloudArchive Direct.

Field	Description
DataLock	If applicable, add a DataLock for compliance and regulatory requirements, to ensure that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users, including those who have the Data Security role in Cohesity, from modifying or deleting any Snapshots that were generated by the Protection Groups that use this Policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the Protection Group fails. The default time between retries is five minutes. You can customize the number of retries and how long to wait between each attempt.
Periodic full backup	Use this option to choose when to run a periodic full backup.
Quiet Times	Quiet Times defines time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturday or Sunday. Quiet Times only prevents new protection runs from starting during the specified time period. The protection runs that started before the quiet time period will be dependent on the quiet time settings configured in the additional settings of the protection Group. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	To replicate snapshots to another, click the Replication icon from the floating menu. The copied snapshots are stored in the storage domain of the remote cluster. .
Archival	To add an archive schedule, click archival icon from the floating menu.

Use an Existing Protection Policy

You can use an existing Protection Policy protecting Isilon. If you are selecting an existing protection policy, then the following options are ignored if they are already configured in the protection policy:

- **Log Backup (Databases)**
- **Add BMR Backup (Physical Server)**
- **CloudSpin**

NAS File DataLock

Many organizations now rely on WORM (Write Once Read Many) data storage such as Nutanix Files, NetApp Snaplock, and Isilon Smartlock volumes to retain their files and folders in unmodified form for compliance and regulatory purposes and to protect the files from accidental deletion or modification.

When you backup these WORM-capable volumes to the Cohesity cluster, it is essential to have the ability to easily lock the files and folders of the volumes while protecting themselves from accidental or malicious attempts to delete or modify these files. The backup should also provide a strong guarantee for data non-mutability and non-erasability for a certain period of time.

This capability is achieved using Cohesity's File DataLock feature.

FileDataLock is a Cohesity feature that preserves the WORM attributes of any files and folders created in the Cohesity View. You can apply the File DataLock feature for NAS backup (NFS exports and SMB shares) while creating a NAS Protection Group.

With File DataLock enabled in a NAS Protection Group, Cohesity preserves the future AccessTime (ATime) of the files and folders. The AccessTime provides the lock period for files and folders. You can also override this lock period in the File DataLock settings of the NAS Protection Group. This lock period is applied to these files and folders when they are recovered to the Cohesity View. This prevents the users or administrator from modifying or deleting these locked files until the locking period expires.

You can also set lock periods for the new files and folders created on recovered Cohesity View in the File DataLock settings of the NAS Protection Group.

When you recover File DataLock enabled storage volumes to the Cohesity View, the View details page will display a lock icon for the files and folders having File DataLock.

File DataLock is supported for all types of NAS sources including NetApp 7-mode and is agnostic of the NAS protocol.

For information on how to enable File DataLock, see Protect NAS.

File DataLock Modes

You can use the File DataLock feature in one of the two modes:

- **Compliance mode**: The WORM files are protected at the disk level. This mode implements strict regulatory requirements for data retention and complies with SEC Rule 17a-4. In this mode, files and folders cannot be deleted or modified during the retention period. After the lock expires, no user can modify the files and folders. However, a root user can delete the files and folders. Only a user with a Data Security

role can edit the File DataLock settings in the Compliance mode.

- **Enterprise mode**: The WORM files are protected at the file level. This mode does not implement strict regulatory requirements for data retention. In this mode, files and folders cannot be modified by any user before or after the lock retention period expiration. However, a root user can delete the files and folders during or after the lock retention period. An admin user or a user with the Data Security role can edit the File DataLock settings in the Enterprise mode

> **Note:** An admin user or a user with the Data Security role can edit the File DataLock settings and switch from the Enterprise mode to Compliance mode. But switching from Compliance mode to Enterprise mode is not supported.

Things to remember for File DataLock

- You cannot enable or disable File DataLock once the View is created.
- If you have disabled File DataLock while creating a Protection Group, then when you edit this Protection Group, the File DataLock option will not be displayed.
- You can delete Protection Groups and snapshots having File DataLock enabled.

  File DataLock also preserves the ATimes of the hardlinks and symlinks.
- **Recovery**
  - Recovering a File DataLock enabled backed-up volume as a Cohesity View will preserve ATimes and the locking period for the file.
  - Recovering a File DataLock enabled volume to the original location with 'Overwrite existing files/folder' as the recovery option will encounter an error if the files and folders are read-only. However, the ATime and other attributes of other files and folders of the source volumes are preserved.
  - Recovering a File DataLock enabled volume to an alternate NAS recovery will preserve ATimes and other attributes of files and folders of the source volumes.
  - Cloning a recovered View having File DataLock enabled will not inherit the File DataLock properties.
- **Replication**

  When you replicate to a remote Cohesity cluster, the File DataLock settings are also replicated.
- **Archival**

  File DataLock settings are preserved when you recover volume from an archived snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.
- **Cloud Archive Direct**

  File DataLock settings are preserved when you recover from a Cloud ArchiveDirect snapshot whose Protection Group has File Datalock enabled, to a Cohesity View.

Exclusions and Inclusions for NAS Protection Group

In a NAS Protection Group, all the files and folders are included by default for data protection. While creating a Protection Group, you can toggle on **Exclusions and Inclusions** from the **Advanced Settings** options if you want to exclude or include locations. By creating the exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories. As a result, you can minimize the disk space used to store the data.

> **Note:**
>
> - To add an exclusion or inclusion rule, you must prefix a forward slash (‘/’) or suffix an asterisk (‘*’) to the path or to a particular file within the protected object. For example, ‘**/test**’ or ‘***.txt**’.
> - In exclusion or inclusion rules, file names and folder names are case-sensitive, which means ‘**test.txt**’ and ‘**Test.txt**’ are considered separate files. Use the correct capitalization of file or folder names to ensure that the rules recognize your file or folder.

Add Inclusion

From the Additional Settings options on the Protection Group page, toggle on **Exclusions and Inclusions**. Click **Add Inclusion** to include a particular path or a particular file within the protected object.

For example, consider four directories - **test**, **test1**, **test2**, and **test3** under the protected object with path `/ifs/TestShare1/Folder1`. The table below lists the input types for an inclusion list of folders and their respective outcome:

Input	Outcome
/	Everything inside the protection path is backed up if exclusions are not defined.
/test	When the inclusion path is not suffixed by '/', Protection Group will include all folders starting from the protection path. In this example, /test will behave as /test* and Protection Group will include '**test**', '**test1**', '**test2**', and '**test3**' from the protection path for backup.

Input	Outcome
/test/	When the inclusion path is suffixed by '/', it will only include the specific folder. In this example, /test/ behaves as /test/* and Protection Group will only include '**test**' from the protection path for backup.

- Do not provide the folder name along with the backup path while including files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not include any folder.
- If exclusions are defined, they will take precedence over inclusions.

Add Exclusion

From the Additional Settings options on the Protection Group page, toggle on **Exclusions and Inclusions**. Click **Add Exclusion** to exclude a particular path or a particular file within the protected object. You can also specify regular expressions for excluding files in the format `regex:<regex pattern>`.

For example, consider there are three directories - **Test**, **Test1**, and **Folder1** and a file, **TestFile.tmp** under the protected object with path **/ifs/TestShare1**.

- The **Test1** folder has **TestFile1.tmp**, **TestFile2.tmp** files.
- The **Folder1** folder has **TestFile3.tmp**, **TestFile4.tmp files**, and **SubFolder1** which further has a **Test** folder with **File.txt**, **File1.txt** files in it.

The following figure illustrates the directory structure under the protected object with path **/ifs/TestShare1**.

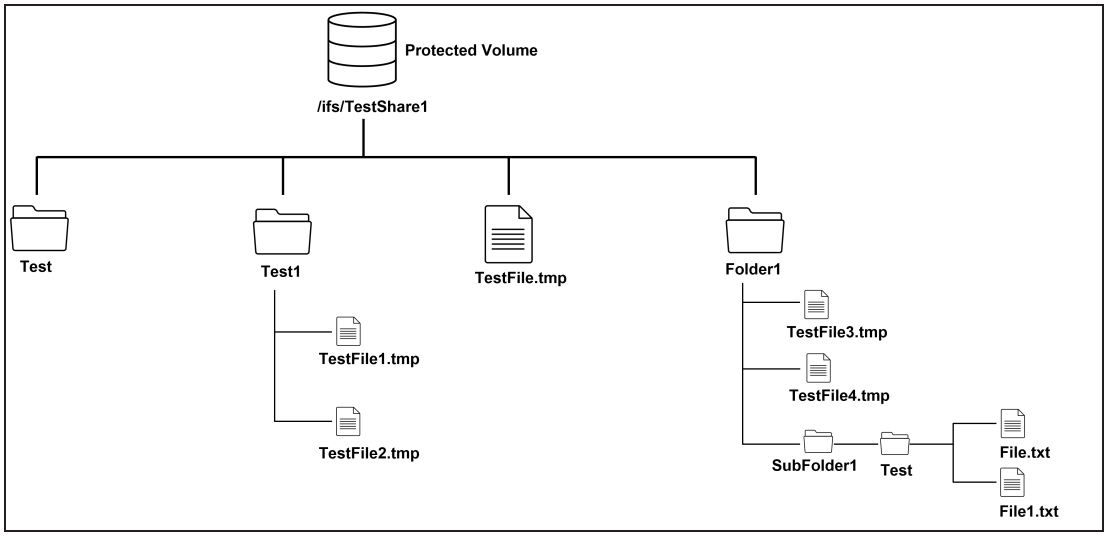

The table below lists the input types for an exclusion list of folders and their respective outcome:

Input	Description	Outcome
/Test	When the exclusion path is not suffixed by '**/**', the Protection Group will exclude all files and folders with **Test** as the prefix. For exclusions, **/Test** behaves as **/Test***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes **Test** and **Test1** folders, and **TestFile1.tmp** file from the protection path. **TestFile3.tmp** and **TestFile4.tmp** files in the **Folder1** folder will not be excluded.
/Test/	When the exclusion path is suffixed by '**/**', the Protection Group excludes only the files and folders with the name **Test**. For exclusions, **/Test/** behaves as **/Test/***. **Note:** Exclusion will not be effective on the sub-directory level.	Excludes the **Test** folder of the parent directory. The **Test** sub-folder in **/Folder1/SubFolder1** will not be excluded.
*/Test	Excludes all the files and directories with the name '**Test**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes the **Test** folder of the parent directory and the **Test** sub-folder in **/Folder1/SubFolder1**.
*.tmp	Excludes all the files that with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.
regex:.*\.tmp	Excludes the files and folders with '**.tmp**'. **Note:** Exclusion will be effective on all levels of sub-directories.	Excludes **TestFile.tmp**, **TestFile1.tmp**, **TestFile2.tmp**, **TestFile3.tmp**, **TestFile4.tmp** files.

Input	Description	Outcome
regex:/Folder1/.*/Test/.*\.txt$	Excludes all files with **.txt** extension from all the **Test** folders directory in **/Folder1/**.	Excludes **File.txt** and **File1.txt** files in the **/Folder1/SubFolder/Test** folder.

> **Note:** Do not provide the folder name along with the backup path while excluding files and folders for a Protection Group. The backup path is considered by default. For example, `/ifs/TestShare1/Folder1/test` does not exclude any folder.

Cohesity automatically excludes the following NetApp system files:

- .vtoc_internal and .bplusvtoc_internal files
- .copy-offload directory and .tokens file

Considerations

Cohesity does not support inclusion or exclusion of paths in:

- an IsilonProtection Group with Changelist enabled.
- a NetApp SnapDiff-basedProtection Group.

Manage the Isilon Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirement.

> **Note:** You cannot update the registered Isilon Server and the assigned Storage Domain in the Protection Group settings.

Add Isilon volumes to the Protection Group

You can add new Isilon volumes to the existing Protection Group for backup. The newly added volumes are backed up from the next protection run and it will be a full backup.

To add Isilon volumes to the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.

5. Select the Isilon volumes you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added Isilon volumes will be backed in the next scheduled protection run.

Remove Isilon volumes from the Protection Group

From the **Edit Protection** page, you can remove Isilon volumes from the Protection Group that you do not want to back up in the next protection run.

To remove Isilon volumes from the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Unselect the Isilon volumes that you want to remove from the Protection Group.
6. Click **Save Selection** > **Save**.

The removed Isilon volumes will not be backed in the next scheduled protection run.

Enable Auto Protect

The Auto Protect option enables you to automatically protect the new Isilon volumes that are added to the selected parent Isilon server. However, you cannot enable auto-protect at the cluster level containing both Read-Write volumes and Data-Protect volumes.

To enable auto protect:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the **Edit** icon.
5. Click the shield icon to enable Auto Protect.
6. Click **Save Selection** > **Save**.

When this option is enabled on the parent Isilon cluster, all the Isilon volumes that are added to the parent Isilon cluster in the future are automatically added to the Protection Group and are protected from the next protection run.

For example, you create a Protection Group that backs up an Isilon cluster. The Protection Group runs for two weeks and then a new volume is added to the Isilon cluster. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new volume is also backed up even though the new volume has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select

Data Protection > Sources, find the source in the list, and click . To exclude an object from Auto Protect, click .

- Indicates an Auto Protected object.
- Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirements. In the Protection Policy, you can edit the settings listed in the Create or Use an Existing Protection Policy.

To the policy settings:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Policy** field, click **Edit**.

   The Edit Policy page appears.
5. Edit the settings listed in the Create or Use an Existing Protection Policy and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, from the drop-down list in the **policy** field, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to back up Isilon volumes. It provides you options to back up all the objects or only selected objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- Is configured with replication targets, you can select those targets.

- Is configured with archival targets, you can select those targets.

To start a protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    a. In the **Backup** drop-down list, select any one of the following:
        - **Backup all Objects in the Protection Group**
        - **Backup only selected Objects in the Protection Group**
    b. In the **Backup Type** drop-down list, select the backup type.
    c. In the Replicate section, select the replication target, the retention duration, and enable the **Replicate only fully successful runs** checkbox if you want to replicate the data only after a successful protection run.
    d. In the Archive section, select the archive target, the retention duration, and enable the **Archive only fully successful runs** checkbox if you want to archive the data only after a successful protection run.
5. Click **Run Now**.

Pause Protection Runs Instantly

You can instantly pause a protection run of the Protection Group that you have created to back up the Isilon sources.

Perform the following steps to pause a currently executing protection run of a Protection Group:

1. In the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, click the Isilon Protection Group that has been created.
3. Click the action icon next to the protection run and click **Pause Run**.

Pause Future Protection Runs

You can stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run.

To resume a paused protection run:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the action icon next to the Protection Group and select **Resume**.

Delete the Protection Group

Cohesity provides the following ways to delete the Protection Group associated with Isilon:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Protection Group are not deleted
- **Delete Protection Group and Snapshots**: Delete the Protection Group and all Snapshots, any local Snapshots captured by this Protection Group and stored on this Cluster are deleted. In addition, any archived Snapshots stored on external targets that were captured by this Protection Group are also deleted.

To delete the Protection Group:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Isilon.
3. Click the action icon next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Group Run.

If the Protection Group is replicating to a remote cluster, the inactive job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the inactive job.

Recover NAS Source

Cohesity recovers NAS data from snapshots of storage volumes created earlier by a Protection Group. Cohesity offers two levels of recovery granularity:

- Recover Storage Volume
- Recover Files or Folders

NAS volumes and files and folders can be recovered to their original location or to a newly specified location, which can be in the original source or a different NAS source. The following figure illustrates the phases and choices you encounter in a typical NAS data restore workflow.

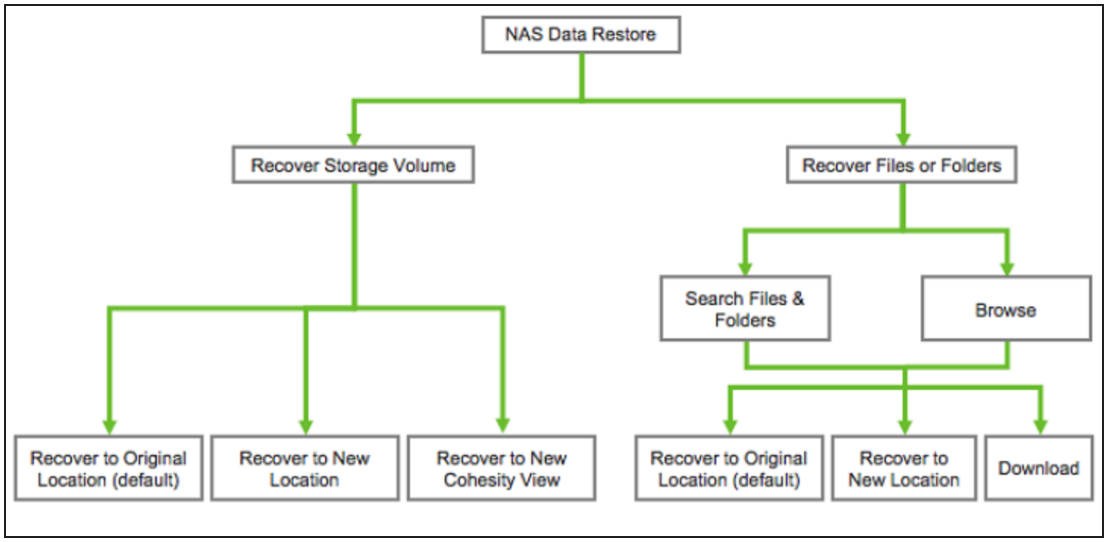

Recover Storage Volume

Cohesity's recover storage volume capability allows you to select and restore specific NAS shares from any previous backup. Before you can recover storage volumes, ensure that snapshots of those volumes exist in the Protection Groups on your Cohesity cluster.

For more information, see Recover NAS Storage Volumes.

Recover Files and Folders

Cohesity’s recover files or folders capability allows you to search and recover specific files and whole folders from any previous NAS backup. Before you try to recover files or folders, ensure that snapshots of these files and folders appear on the Cohesity cluster in a Protection Group. For backed-up SMB volumes, the user account must have full control on the target as Cohesity uses that user to perform recovery.

Cohesity now allows the recovery of up to eight files or one folder from SnapDiff-based backups.

For more information, see Recover NAS Files or Folders.

Hardlink

A hardlink set consists of file references (hardlinks) with the same inode, pointing to the data's location on the filesystem.

For example, hardlinks can be created from file A such as B, C, D, and so on. All these files will be a part of the hardlink set.

You can restore the hardlinks to their original location or restore multiple directories containing common hardlinks. For more information, see Recover hardlinks.

Recover NAS Files or Folders

You can restore the files or folders to their original location or to a newly specified location, which can be within the original source or a different one, without losing the original permissions and attributes. For more information, see:

- Recover NAS Files or Folders to the Original Location
- Recover NAS Files or Folders to a New Location

You can also download specific files and folders directly from any snapshot that was created by a Cohesity Protection Group.

Before you try to recover files or folders, ensure that snapshots of these files and folders appear on the Cohesity cluster in a Protection Group. For backed-up SMB volumes, the user account must have full control on the target as Cohesity uses that user to perform recovery.

Recover NAS Files or Folders to the Original Location

To recover files and folders to the original NAS:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **NAS > Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the source or Protection Group from any backup available on the Cohesity cluster.

If you select..	Then..
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recovery point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial source or Protection Group name. You can optionally specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. By default, only the files and folders that are indexed are displayed. To display all the available files and folders, toggle off the **Browse on Indexed Data** option. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery.

If you select..	Then..
	4. By default, the snapshot is recovered from the local cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next Recover Option**: If you select this option, then continue with step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
     Example:
     - The original name is "a:b" then the archived name is "a_COLON_b"
     - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

     If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
     Example:

- The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
- The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **Original Location** to recover the files or folders to the original source.
6. To recover to a different location in the original NAS, enable the **Recover to Alternate Path** option, and then provide the location to which the files or folders are to be recovered in the **Alternate Path** field. By default, the files and folders will be recovered to the original location.
7. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable the **Overwrite** option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

     - C:\restore\1.txt
     - C:\restore\cohesity<taskid>_0\1.txt
     - C:\restore\bar
     - C:\restore\cohesity<taskid>_1\bar

     If **Overwrite** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
   - **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
   - **Encryption**: Enable this option if you want the data to be encrypted in-flight during recovery. By default, this option is disabled. For more information on encryption, see NAS Solution Guide.
   - **Filter IPs**: Enable this option to filter the IP addresses of the NAS source. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the NAS source.
   - **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option,

then select a configured interface group from the **Interface Group** drop-down.

- **Task Name**: Change the default name of the recovery task.

8. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover NAS Files or Folders to a New Location

To recover files and folders to a new NAS:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **NAS > Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the source or Protection Group from any backup available on the Cohesity cluster.

If you select..	Then..
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: **Note:** The option to change the recovery point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial source or Protection Group name. You can optionally specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. By default, only the files and folders that are indexed are displayed. To display all the available files and folders, toggle off the **Browse on Indexed Data** option. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option.

If you select..	Then..
	When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next Recover Option**: If you select this option, then continue with step 5 to configure the file recovery options.
   - **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recovery task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

     Character replacement explanation for : and \ in file and directory names in zip files

     If the original file or directory name contains a ':', it is replaced with '_COLON_' in the archive.
     Example:
     - The original name is "a:b" then the archived name is "a_COLON_b"
     - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

     If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
     Example:
     - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
     - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **New Location** to recover the files or folders to a new NAS source.
6. From the **Registered Source** drop-down menu, select a registered NAS source to which you want to recover the files and folders. You can also register a new source by clicking **Register Source** from the **Registered Source** drop-down menu.
7. From the **Volume** drop-down menu, select a volume to which you want to recover the files and folders.
8. In the **Path** field, enter the path where the files and folders will be recovered. By default, the files and folders will be recovered in the `/temp/recover-file-<month_date_time>` folder of the selected volume.
9. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable the **Overwrite** option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

     - C:\restore\1.txt
     - C:\restore\cohesity<taskid>_0\1.txt
     - C:\restore\bar
     - C:\restore\cohesity<taskid>_1\bar

     If **Overwrite** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

     > **Note:** For Pure Storage FlashBlade, the recovery of SMB data may fail if the files and folders being recovered have restrictive permissions. Cohesity recommends you either create folders or directories manually in the target location before initiating the recovery or use an SMB user with read/write permissions when registering the Pure Storage FlashBlade source.

   - **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.

- **Encryption**: Enable this option if you want the data to be encrypted in-flight during recovery. By default, this option is disabled. For more information on encryption, see NAS Solution Guide.
- **Filter IPs**: Enable this option to filter the IP addresses of the NAS source. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the NAS source.
- **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
- **Task Name**: Change the default name of the recovery task.

10. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover NAS Storage Volumes

Cohesity provides the ability to recover storage volumes from a snapshot created earlier by a Protection Group. To learn more about recovering volumes, see About Recovering Storage Volumes.

Before you can recover storage volumes, snapshots of those volumes must exist on the Cohesity cluster. To create snapshots, create and run a Protection Group.

The options to recover storage volumes are as follows:

- Recover NAS Storage Volumes to the Original Location
- Recover NAS Storage Volumes to a New Location
- Recover NAS Storage Volumes as a Cohesity View

Recover NAS Storage Volumes to the Original Location

Before you can recover volumes, snapshots of those volumes must exist on the Cohesity cluster. To create snapshots, create and run a Protection Group.

To recover a storage volume to the original NAS location:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **NAS > Storage Volume**.

   The Storage Volumes page appears.
3. Search for the NAS volume that you want to recover by entering the characters of the volume name, the Protection Group name that protects the volume, or by specifying the wildcard character *.

   A list of NAS volume displays.

You can optionally filter the search results by source, Protection Group, Storage Domain, or date.

4. From the list of search results, select the volume you want to recover.

   By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <Volume_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

     b. Select the snapshot to be recovered by clicking the corresponding blue dot.

     c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     d. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the **List** mode:
     a. Select the **List** tab from the top-right corner of the Edit recovery point for <Volume_name> page.

        By default, the page displays the snapshots available in the last month.

     b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

     c. Select the snapshots to be recovered.

d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the volume to the original location.
8. Optional. Edit the default settings of the following **Recovery options**:
   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.
   - **Preserve file/folder attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders of the NAS volume. If you disable this option, then the ACLs and the permissions are not preserved. If recovering folders and files, folders receive new timestamps, but files retain their original timestamps. If recovering files only, files receive new timestamps.
   - **Continue on Error**: Enable this option if you want to continue recovery even if an error occurs during the recovery task. By default, this option is disabled and the recovery operation will fail if it encounters an error.
   - **Encryption**: Enable this option if you want the data to be encrypted in-flight during recovery. By default, this option is disabled. For more information on encryption, see NAS Solution Guide.
   - **Filter IPs**: Enable this option to filter the IP addresses of the NAS source. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the NAS source.

     To filter the IP addresses of the NAS source, select one of the following options:

     - **Allow IPs**: Select this option and specify the IP addresses of the NAS source through which the communication to the Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.
     - **Deny IPs**: Select this option and specify the IP addresses of the NAS source through which the communication to the Cohesity cluster.

> **Note:**
> If you have specified the IP filtering at the source-level also, then the IP filtering specified at the recovery task-level will take precedence over the IP filtering specified at the source-level.

- **Cluster interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
- **Task name**: Change the default name of the recovery task.

9. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover NAS Storage Volumes to a New Location

Before you can recover volumes to an alternate location as a new storage volume, snapshots of those volumes must exist on the Cohesity cluster. To create Snapshots, create and run a Protection Group.

To recover a storage volume to a new location:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **NAS > Storage Volume**.

   The Storage Volumes page appears.
3. Search for the NAS volume that you want to recover by entering the characters of the volume name, the Protection Group name that protects the volume, or by specifying the wildcard character *.

   A list of NAS volume displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, or date.
4. From the list of search results, select the volume you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <Volume_name> page appears.

> **Note:** In this workflow, deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process.

By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date.

     The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.
  d. Click **Select Recovery Point** and then go to step 6.
- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the Edit recovery point for <Volume_name> page.

     By default, the page displays the snapshots available in the last month.
  b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
  c. Select the snapshots to be recovered.
  d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.
  e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** to recover the volume to an alternate location as a new storage volume.

8. From the **Registered Source** drop-down menu, select a registered NAS source to which you want to recover the NAS volume. You can also register a new NAS source by clicking **Register Source** from the Source drop-down menu.
9. From the **Target** drop-down menu, select a volume to which you want to recover the volume.
10. Optional. Edit the default settings of the following **Recovery options**:
    - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.
    - **Preserve file/folder attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders of the NAS volume. If you disable this option, then the ACLs and the permissions are not preserved. If recovering folders and files, folders receive new timestamps, but files retain their original timestamps. If recovering files only, files receive new timestamps.
    - **Continue on Error**: Enable this option if you want to continue recovery even if an error occurs during the recovery task. By default, this option is disabled and the recovery operation will fail if it encounters an error.
    - **Encryption**: Enable this option if you want the data to be encrypted in-flight during recovery. By default, this option is disabled. For more information on encryption, see NAS Solution Guide.
    - **Filter IPs**: Enable this option to filter the IP addresses of the NAS source. By filtering IP addresses, you can allow or deny the communication of the Cohesity cluster to specific IP addresses or subnets of the NAS source.

      To filter the IP addresses of the NAS source, select one of the following options:
      - **Allow IPs**: Select this option and specify the IP addresses of the NAS source through which the communication to the Cohesity cluster must happen. You can provide the IP addresses in a comma-separated list or in a CIDR format.
      - **Deny IPs**: Select this option and specify the IP addresses of the NAS source through which the communication to the Cohesity cluster.

    > **Note:**
    > If you have specified the IP filtering at the source-level also, then the IP filtering specified at the recovery task-level will take precedence over the IP filtering specified at the source-level.

- **Cluster interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
- **Task name**: Change the default name of the recovery task.

11. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover NAS Storage Volumes as a Cohesity View

When recovering a NAS storage volume to a Cohesity View, Cohesity clones the volume's backed up data to a new Cohesity View within the Cohesity cluster and provides instant access to the data.

> **Note:** This workflow is not supported for NAS backups performed using the NFS4.1 protocol.

Before you can recover volumes, snapshots of those volumes must exist on the Cohesity cluster. To create Snapshots, create and run a Protection Group.

To recover a storage volume as a Cohesity View:

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **NAS > Storage Volume**.

   The Storage Volumes page appears.
3. Search for the NAS volume that you want to recover by entering the characters of the volume name, the Protection Group name that protects the volume, or by specifying the wildcard character *.

   A list of NAS volume displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, or date.
4. From the list of search results, select the volume you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <Volume_name> page appears.

> **Note:** In this workflow, deleted snapshots may be displayed as valid recovery points for a brief period until they are removed from the search index by a background process.

By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date.

     The snapshots available on the selected date are displayed in blue dots on the timeline bar.
  b. Select the snapshot to be recovered by clicking the corresponding blue dot.
  c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.
  d. Click **Select Recovery Point** and then go to step 6.
- To select a snapshot from the **List** mode:
  a. Select the **List** tab from the top-right corner of the Edit recovery point for <Volume_name> page.

     By default, the page displays the snapshots available in the last month.
  b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.
  c. Select the snapshots to be recovered.
  d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.
  e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Cohesity View** to recover the volume to a new view within the Cohesity cluster.

8. Optional. In the **Name** field, change the default name of the new view where you want to recover the volume to. This cannot be an existing view name. Using UNIX special characters for the view name may cause the recovery to fail.
9. From the **QoS Policy** drop-down, select a QoS policy. Each Cohesity view is assigned a Quality of Service (QoS) Policy. A QoS Policy determines to which storage media it is written and the priority of I/O (when contention occurs). For details and recommendations, see QoS Policies for Cohesity Views.
10. Optional. Edit the default settings of the following **Recovery options**:
    a. **Cluster interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
    b. **Task name**: Change the default name of the recovery task.
11. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Once the recovery is complete, you must access or mount the View using SMB or NFS mount path to the system of your choice and then perform the desired operation with the recovered data.

To locate the mount point information, navigate to **Data Protection > Recoveries** and click the completed recovery task to view its results page. The **NFS Mount Path** for NFS systems and **SMB Mount Path** for Windows systems are displayed.

Recover hardlinks

A hardlink set consists of file references (hardlinks) with the same inode, pointing to the data's location on the filesystem.

You can restore the hardlinks to their original location or restore multiple directories containing common hardlinks.

Recover hardlinks to the original location

When recovering, you can enable or disable the "Overwrite" option. If the option is disabled, only the missing hardlinks from any hardlink set will be recreated after the restore. In this case, a new inode pointing to the backed-up version of the data will be created for the missing hardlink file, and all the other hardlinks to be restored will point to this inode. However, if the option is enabled, all hardlinks from any hardlink set are recreated after restoration, pointing to a new inode that points to the backed-up data version.

For example, let's say there is a hardlink set containing files from A to F, all of which are backed up. If we delete the hardlinks A, B, and C, and then restore to the original path with the “Overwrite” option disabled, a new inode will be created. This inode will point to a new backed-up version of the data, and hardlinks A, B, and C will be recreated and point to this new inode. The hardlinks D, E, and F will remain unchanged and have the original inode.

With the Overwrite option enabled, all files (A to F) will be recreated as new hardlinks, pointing to a new inode that contains the backed-up data version.

Recover multiple directories containing common hardlinks

If multiple folders are being restored, and they contain common hard links, each folder is restored independently. This means that a new inode is created for each hardlink set within a folder, pointing to a new backed-up version of the data. Any missing hardlinks are recreated, pointing to the new inodes of their respective folders.

Consider a scenario where directories A, B, and C have multiple hardlinks: A1 and A2 in A, B1 and B2 in B, and C1 and C2 in C, all pointing to the same file (with inode 100). With these hardlinks missing, if we restore only the directories A, B, and C instead of the entire volume, all these hard links will be recreated. The newly created A1 and A2 will point to a new inode (101), while B1 and B2 will point to a new inode (102), and C1 and C2 will point to a new inode (103). Each of these inodes will then reference separate backed-up versions of the data.

Replicate NAS Data

You can replicate Isilon data backup to the Cohesity cluster to another Cohesity cluster for disaster recovery use cases. For more details about replication, see Replication and Remote Access Setup.

Cohesity provides a policy-based replication solution for Isilon data. In the Protection Policy, you can specify the remote Cohesity cluster, define conditions for replication, specify the retention period of the data on the remote Cohesity cluster, and other parameters.

To replicate Isilon data to another Cohesity cluster:

1. Configure replication to another Cohesity cluster. For more details, see Replication and Remote Access Setup.
2. Edit the existing Protection Policy that you had created for backup or create a new Protection Policy for backup. Configure the replication settings as per your requirement. For more information, see Create or Edit a Standard Policy.
3. Run the Protection Group that you have created for Isilon based on your requirement or the replication settings will take effect in the next scheduled Protection Group run.

Archive NAS Data

Cohesity provides the option to store a copy of your data backed up onto the Cohesity cluster to a designated external target, such as tape or cloud storage for long-term retention purposes. The process of storing a copy of the data backed up on the Cohesity cluster to an external target is called Archival. For more information, see About Cloud Archival.

Cohesity provides a policy-based archival solution for NAS data. In the Protection Policy, you can define the archive schedule, select the external target, and specify the retention period of the data on the external target.

To archive NAS data:

1. Register an external target. For more details, see Register an External Target.
2. Edit the existing Protection Policy that you had created for backup or if you are creating a new Protection Policy for backup. Configure the Archival settings as per your requirement. For more information, see **Step 14** in Create or Edit a Standard Policy.
3. Run the Protection Group that you have created for the respective NAS source based on your requirement or the Archival settings will take effect in the next scheduled Protection Group run.

Chapter 15: Microsoft 365

Microsoft 365 is a subscription service that bundles the traditional office productivity applications and delivers them as SaaS applications. Microsoft 365 includes Exchange Online, OneDrive for Business, SharePoint Online, Teams, and other applications. Cohesity provides enterprise-grade backup and recovery solutions for Microsoft 365.

Cohesity's software-defined platform provides a simple, fast, and cost-effective data protection solution for the following Microsoft 365 applications:

- Exchange Online

> **Note:** Starting October 2026, Cohesity will discontinue support for Public Folder protection, following Microsoft's retirement of APIs for managing Public Folders.

- OneDrive for Business
- SharePoint Online
- Groups (Early Access)
- Microsoft Teams

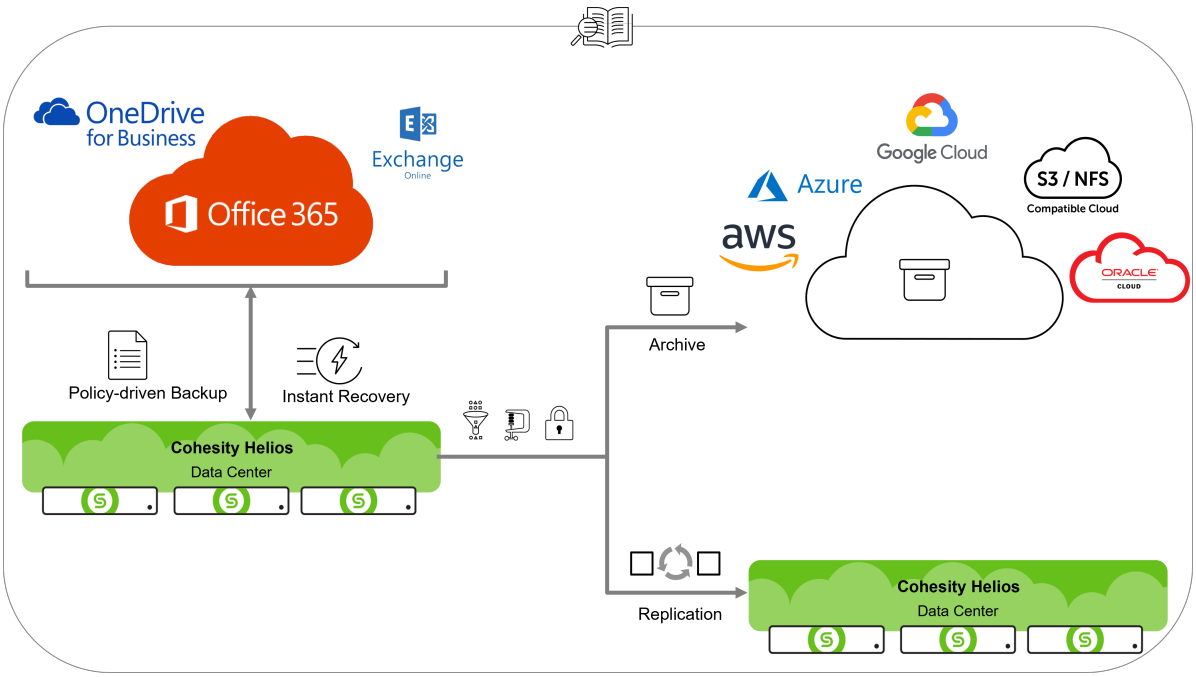

> **Important:** Azure ACS reached full retirement across all Microsoft 365 tenants on April 2, 2026. Authentication requests relying on ACS are no longer supported. If your Azure application is currently using ACS-based authentication, Cohesity recommends migrating to Certificate-Based Authentication (CBA) to ensure continued support for SharePoint backup and recovery. The Azure application itself remains in use only; the authentication method needs to be updated. For more information, see Microsoft 365 Certificate-Based Authentication.

Microsoft 365 Data Backup and Restore: Process Overview

This section provides a high-level overview of the steps required to protect data in your Microsoft 365 domain. These steps include completing the prerequisites, registering the Microsoft 365 domain as a source with Cohesity, performing backups, and restoring data when needed.

- **Complete the following prerequisites**:
  - Review and understand the supportability for Microsoft 365 backup and recovery workflows.
  - Review and understand the considerations and best practices for each application captured in its **Plan and Prepare** section.

    - Review the following general considerations:
        - When multiple clusters are deployed, ensure that the Microsoft 365 source is registered on all the clusters to avoid recovery failures.
        - Cohesity does not support Basic Auth for Microsoft 365 source registration.
    - Add Microsoft 365 Service User Account with Required Roles
    - Add a Custom Application for Cohesity in Microsoft Entra ID (Azure AD)
    - Configure Microsoft Organization Setting for Mailbox Size Reporting
    - Configure Firewall and Port Requirements
    - (Optional) Configure Windows Connector
    - Configure Enhanced Backup of Large Microsoft 365 Data

    For more information, see Prerequisites.
- **Register your Microsoft 365 domain as a source in Cohesity** after completing the above prerequisites. You can register using either of the following methods:
    - Manual registration method
    - Express registration method

    Once the registration is complete, your Microsoft 365 domain is available as a source for the backup solution. You can then back up your Microsoft 365 domain data and restore it in case of data loss.
- **Perform backup**
- **Perform restore**

GCC High: Microsoft 365 Data Backup and Recovery

Cohesity supports the protection and recovery of Microsoft 365 in a U.S. Government Community Cloud High (GCC High) tenant. Consider the following points:

- Backup and recovery of Public Folders is not supported.
- Recovery of Private Channels is not supported.

Using Microsoft Graph APIs

By September 2026, Cohesity will migrate from Exchange Web Services (EWS) APIs to Microsoft Graph APIs to protect Microsoft 365 applications. This is as per the Retirement of Exchange Web Services in Exchange Online announcement from Microsoft. For more information, see Deprecation of Exchange Web Services in Exchange Online.

As part of the migration to Microsoft Graph APIs, Cohesity now supports Exchange Online mailbox backup and recovery using Microsoft Graph APIs as an Early Access feature,

replacing Exchange Web Services (EWS) APIs. This transition ensures seamless and uninterrupted mailbox protection while providing greater flexibility and compatibility. Existing backups remain fully usable, with no need for data re-ingestion. The EWS to Graph transition supports only user and shared primary mailboxes in this release. Teams mailboxes, Groups mailboxes and user/shared mailboxes that protect the in-place archive continue to use EWS.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Prerequisites

You must complete the following prerequisites to back up and restore data in your Microsoft 365 domain:

- Review and understand the supportability for Microsoft 365 backup and recovery workflows.
- Review and understand the considerations and best practices for each application captured in its **Plan and Prepare** section.
- Review the following general considerations:
  - When multiple clusters are deployed, ensure that the Microsoft 365 source is registered on all the clusters to avoid recovery failures.
  - Cohesity does not support Basic Auth for Microsoft 365 source registration.
- Add Roles to Microsoft 365 Service User Account
- Register Custom App in Microsoft Entra ID
- Update Microsoft Organization Setting for Mailbox Size Reporting
- Firewall and Port Requirements for Microsoft 365 Protection
- Set Up a Windows Connector for Microsoft 365
- Enhance Backup of Large Microsoft 365 Data

Add Roles to Microsoft 365 Service User Account

To register your Microsoft 365 domain with Cohesity, you must have a service user account with the required roles. You can either create a new user account with the required roles or assign the required roles to an existing account. To add the Microsoft 365 service user account, ensure the following requirements are met:

- The account must have a valid SharePoint license and a Mailbox license. A Microsoft 365 license is not required unless you plan to protect Public Folders. Assign the roles

to the service user account as described in the following procedure.

- If you want to back up and restore Public Folders, the service user account must have the required root permissions

> **Important:** Ensure that multi-factor authentication is not enabled for the user account.

To add roles to the Microsoft 365 service user account:

1. Log in to Microsoft 365.
2. In the **Office 365** page, click **Admin**.
3. In the **Microsoft 365 admin center** page, select **Admin centers**, and then click **Exchange**.
4. Select **Roles** > **Admin roles**.
5. In the **Admin roles** page, click **Add role group**.
6. On the **Add role group** page , do the following:
    1. Under **Basics**, enter a **Name** and **Description** for the admin role. From the **Write scope** drop-down, select **Default**, and then click **Next**.
    2. Under **Permissions**, select the following and click **Next**:
        - `Mail Recipients`
        - `Mailbox Search`
        - `View-Only Configuration`
        - `View-Only Recipients`
    3. Under **Admins**, search and select the user account you plan to use to register the Microsoft 365 domain with Cohesity Platform, then click **Next**.
    4. Under **Review and finish**, review the configuration and click **Add role group**.
    5. After the role group is added, click **Done**.

       You are ready to update your Microsoft 365 Org setting for Mailbox size reporting. For more information, see Update Microsoft Organization Setting for Mailbox Size Reporting.

Here is the GIFF you can refer to add roles to the Microsoft 365 service user account:

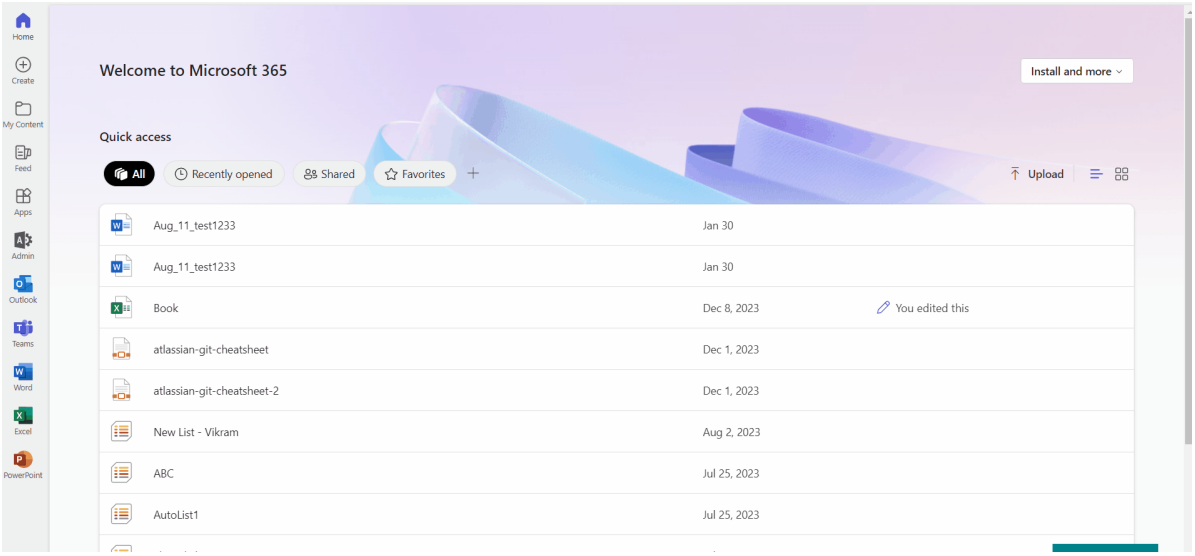

Add Required Public Folder Root Permissions for Microsoft 365 Service User Account

Your Microsoft 365 service user account must be configured with **Publishing Editor** root permission on `PublicFoldersRoot (IPM_SUBTREE)` for Cohesity to backup and recover public folders in Exchange Online.

To add the required root permissions to Microsoft 365 user account:

1. Log in to Microsoft 365.
2. In the **Exchange admin center**, select **public folders**.
3. In the **public folders** tab, click (**...**) menu and select **Root permissions**.
4. In the **public folder permissions** page, perform the following:
    1. In the **User** field, click **Browse** and select your Microsoft 365 user account.
    2. From the **Permission level** drop-down list, select **Publishing Editor**.
    3. Ensure that the following permissions are selected:
        - `Create items`
        - `Read items`
        - `Create subfolders`
        - `Edit own`
        - `Edit all`
        - `Folder visible`
        - `Delete own`
        - `Delete all`
    4. Click **Save**.

5. Select **Apply changes to this public folder and all its subfolders** and click **Save**.

Register Custom App in Microsoft Entra ID

To get started, you'll register a custom app below to add the necessary permissions.

Go to the Azure portal, register a new app, add the permissions, and capture the App ID and Access Key. For more on registering and configuring apps in Microsoft Entra ID, see Register an application with the Microsoft identity platform and Configure a client application to access a web API in the Microsoft documentation.

> **Note:** Make sure that you make note of the App ID and Access Key while registering the app. You'll need them to register your Microsoft 365 domain as a source in Cohesity Platform.

To register your custom app for Cohesity Platform:

1. Open Microsoft Entra ID.
    1. To manage Microsoft Entra ID using the Azure Portal:
        1. Log in to the Azure portal with your Microsoft 365 administrator user credentials.
        2. Click the main menu (≡) in the top left corner and select **Microsoft Entra ID**.
    2. To manage Microsoft Entra ID using Microsoft 365:
        1. Log in to Microsoft 365.
        2. On the **Microsoft 365** page, click **Admin**.
        3. On the **Microsoft 365 admin center** page, select **Admin centers** and then click **Microsoft Entra**.
2. Create a new custom app.
    1. Under the **Manage** section, select **App Registrations**, then click **New Registration**. In the **Register an application** page:
        1. Enter a **Name** for your app.
        2. Select the **Supported account types** that can access the app.
        3. In the **Redirect URI** drop-down, select **Web** and enter `https://localhost`.
        4. Click **Register**.

3. After the custom app has been created, click **Overview** and copy the **Application (client) ID**. You need to use the **Application (client) ID** to register Microsoft 365 as a source in Cohesity Platform.

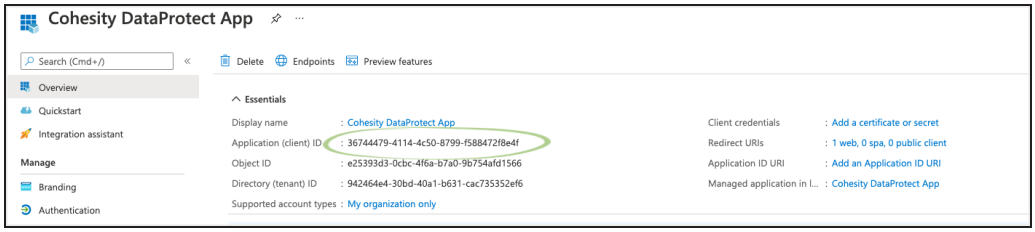

4. Add API permissions to the custom app:
    1. **Add Oauth API permission** if the Microsoft 365 source tenant has OAuth enabled for secure communication:
        1. Under the **Manage** section, select **App Registrations** and click **Add a permission**.
        2. In the **Request API permissions** page, click the **APIs my organization uses** tab.
            a. In the search bar, enter **Office 365 Exchange Online** then **click the API**. (Use the complete app name.)

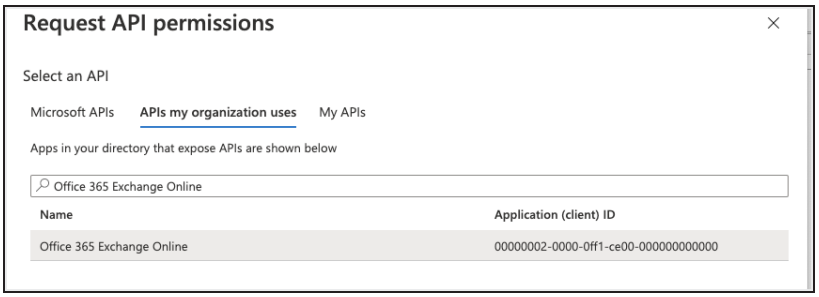

            b. In the Office 365 Exchange Online API, click **Application Permissions**.

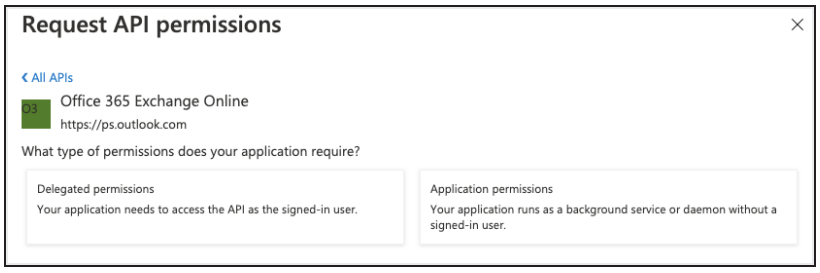

            c. Under **Other Permissions**, select `full_access_as_app` to enable

OAuth and click **Add Permissions**.

App Permissions	Permission Type
full_access_as_app	Application

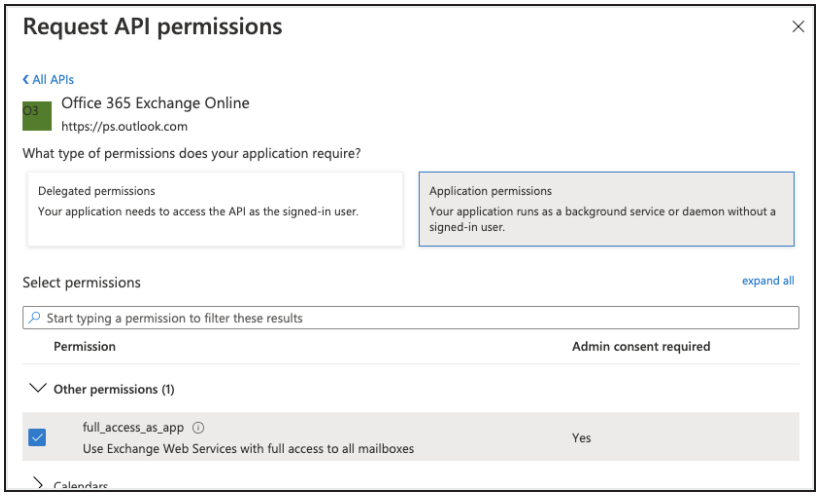

2. Add Graph API permissions:
    1. Under the **Manage** section, select **App Registrations**, and then click **Add a permission**.
    2. In the **Request API permissions** page, select **Microsoft Graph API**.

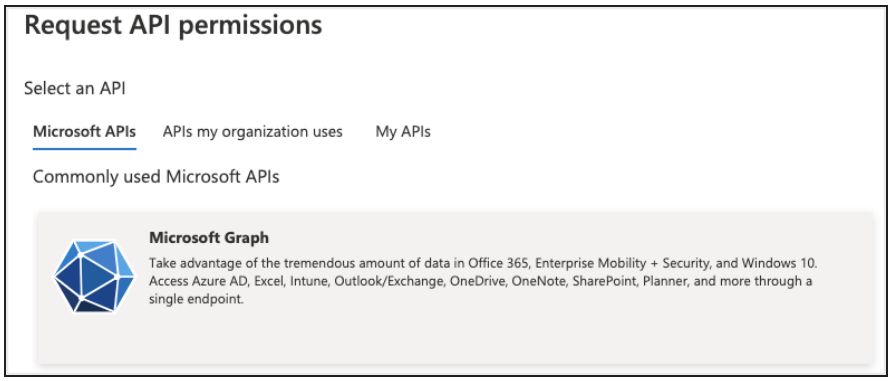

    3. Click **Application Permissions** and add the permissions. See, Permission Descriptions.

App Permissions	Permission Type	Mailboxes	OneDrive	SharePoint Online Sites	MS Group	MS Teams
Application.Read.All	Application	Yes	Yes	Yes	Yes	Yes
Application.ReadWrite.OwnedBy	Application	Yes	Yes	Yes	Yes	Yes
ChannelMember.ReadWrite.All	Application	N/A	N/A	N/A	N/A	Yes
ChannelMessage.Read.All	Application	N/A	N/A	N/A	N/A	Yes
Chat.Read.All	Application	Yes	N/A	N/A	N/A	N/A
Files.ReadWrite.All	Application	N/A	Yes	Yes	Yes	Yes
Group.ReadWrite.All	Application	N/A	N/A	N/A	Yes	Yes
MailboxSettings.Read	Application	Yes	N/A	N/A	N/A	N/A
Organization.Read.All	Application	Yes	Yes	Yes	Yes	Yes
Reports.Read.All	Application	Yes	Yes	Yes	Yes	Yes
Sites.FullControl.All	Application	N/A	Yes	Yes	Yes	Yes
Sites.ReadWrite.All	Application	N/A	Yes	Yes	Yes	Yes

App Permissions	Permission Type	Mailboxes	OneDrive	SharePoint Online Sites	MS Group	MS Teams
TeamsAppInstallation.ReadForChat.All	Application	Yes	N/A	N/A	N/A	N/A
User.Read.All	Application	Yes	Yes	Yes	Yes	Yes

4. Click **Add permissions**.

3. Add SharePoint permissions to the custom app:
    1. Under the **Manage** section, select **App Registrations** and click **Add a permission**.
    2. In the **Request API permissions** page, select **SharePoint**. (If you don't see it, scroll further down.)
        a. Click **Application Permissions** and add the permissions listed below, then click **Add permissions**.

Permission Type	Permissions Name
Application	Sites.FullControl.All

5. Grant admin consent for the API permissions.
    1. Under **Configured permissions**, click **Grant admin consent**.
    2. On the Grant admin consent confirmation, click **Yes**.

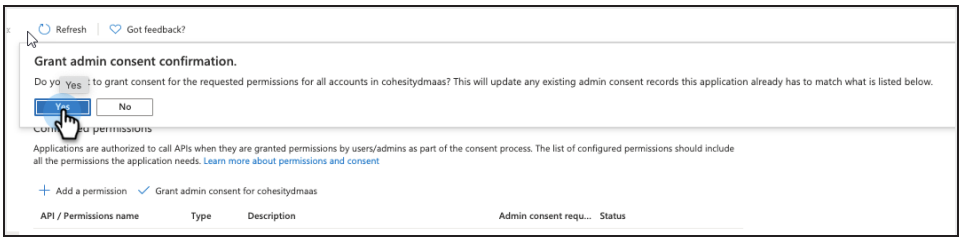

6. Create a new client secret that will be used to register Microsoft 365 as a source in Cohesity Platform.

1. Under the **Manage** section, select **Certificates & secrets**.
    1. In the Client secrets section, click **New client secret**. Enter a **Description**.
    2. In the **Expires** drop-down, select how long the secret key will be valid.
    3. Click **Add**.

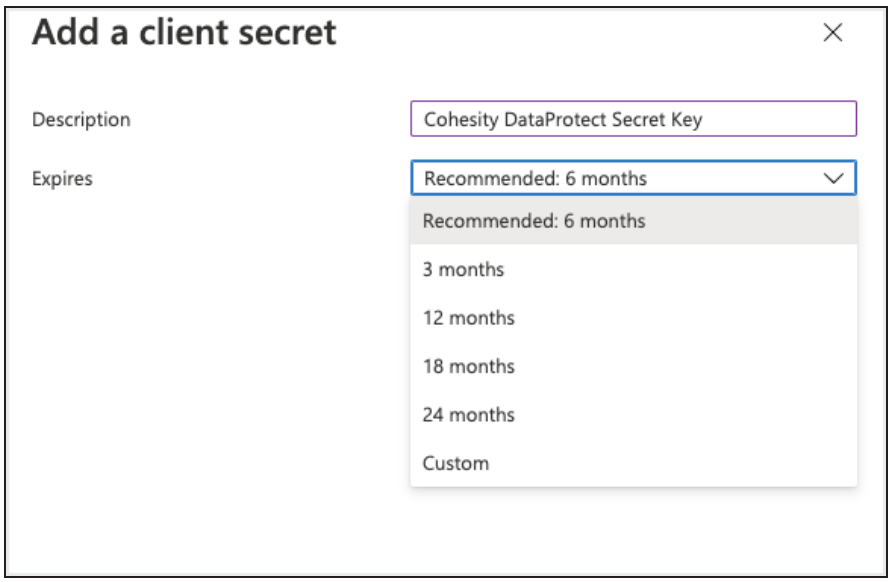

2. Under **Client secrets**, click the **Copy** button next to the string under **VALUE**. You need the Value key of the client secret to register Microsoft 365 as a source in Cohesity Platform.
3. Store the Value key in a secure location. After you exit this page, you will not be able to see the Value key again. If you lose your value key, you will need to create a new client secret.

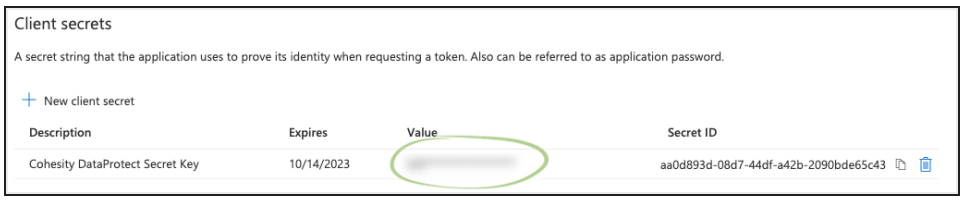

When you finish, your custom app should include the permissions as shown below:

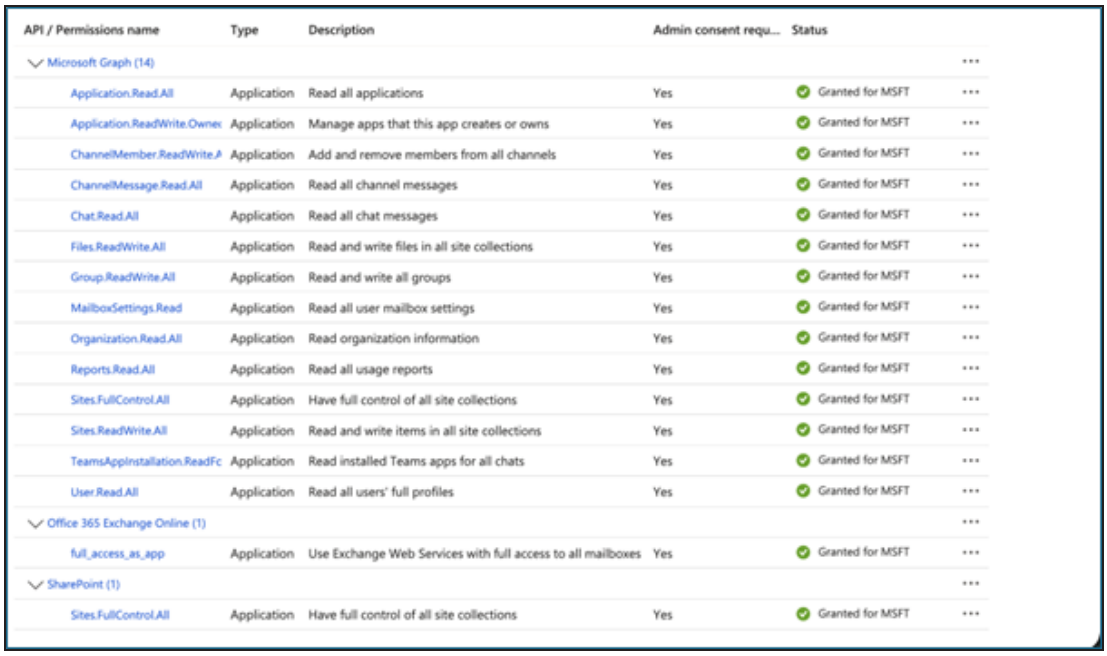

API / Permissions name	Type	Description	Admin consent requ...	Status
Microsoft Graph (14)				
Application.Read.All	Application	Read all applications	Yes	Granted for MSFT
Application.ReadWrite.Owner	Application	Manage apps that this app creates or owns	Yes	Granted for MSFT
ChannelMember.ReadWrite.A	Application	Add and remove members from all channels	Yes	Granted for MSFT
ChannelMessage.Read.All	Application	Read all channel messages	Yes	Granted for MSFT
Chat.Read.All	Application	Read all chat messages	Yes	Granted for MSFT
Files.ReadWrite.All	Application	Read and write files in all site collections	Yes	Granted for MSFT
Group.ReadWrite.All	Application	Read and write all groups	Yes	Granted for MSFT
MailboxSettings.Read	Application	Read all user mailbox settings	Yes	Granted for MSFT
Organization.Read.All	Application	Read organization information	Yes	Granted for MSFT
Reports.Read.All	Application	Read all usage reports	Yes	Granted for MSFT
Sites.FullControl.All	Application	Have full control of all site collections	Yes	Granted for MSFT
Sites.ReadWrite.All	Application	Read and write items in all site collections	Yes	Granted for MSFT
TeamsAppInstallation.ReadFc	Application	Read installed Teams apps for all chats	Yes	Granted for MSFT
User.Read.All	Application	Read all users' full profiles	Yes	Granted for MSFT
Office 365 Exchange Online (1)				
full_access_as_app	Application	Use Exchange Web Services with full access to all mailboxes	Yes	Granted for MSFT
SharePoint (1)				
Sites.FullControl.All	Application	Have full control of all site collections	Yes	Granted for MSFT

Permission Descriptions

> **Note:** The following permissions are required based on the features selected for protection.

The purpose to grant the application permissions is described as follows:

App Permissions	API Type	Permission Type	Purpose
Application.Read.All	Graph	Application	Read all applications and service principals.
Application.ReadWrite.OwnedBy	Graph	Application	Used for rotation of certificates.
ChannelMember.ReadWrite.All	Graph	Application	Add and remove members from channels.
ChannelMessage.Read.All	Graph	Application	Read all channels chat messages.
Chat.Read.All	Graph	Application	Read all users chat messages.
Files.ReadWrite.All	Graph	Application	Back up and recover the OneDrive for Business files.

App Permissions	API Type	Permission Type	Purpose
Group.ReadWrite.All	Graph	Application	Create groups, read all group properties and memberships, and update group properties and memberships.
MailboxSettings.Read	Graph	Application	Read mailbox type.
Organization.Read.All	Graph	Application	Read the organization and related resources. Related resources include things like subscribed skus and tenant branding information.
Reports.Read.All	Graph	Application	Read all usage reports.
Sites.FullControl.All	Graph	Application	Used to avoid resync error in getting incremental changes.
Sites.ReadWrite.All	Graph	Application	Create, read, update, and delete documents and list items in all site collections.
TeamsAppInstallation.ReadForChat.All	Graph	Application	Read installed teams apps for all user chats, for example, this allows identifying apps/bots in messaging.
User.Read.All	Graph	Application	Read full profiles of all users.
full_access_as_app	Exchange	Application	Use Exchange Web Services with full access to all mailboxes.
Sites.FullControl.All	SharePoint	Application	Have full control of all site collections.

Enhance Backup of Large Microsoft 365 Data

When it comes to ingesting Microsoft 365 data, the size can vary from hundreds of terabytes (TBs) to a few petabytes (PBs) in some cases. One of the issues that may arise while performing the initial full backup of such large data is throttling from the Microsoft 365 APIs.

Cohesity recommends using only one application ID for Microsoft 365 backups. However, when you back up a large amount of data, using only one application ID may result in a prolonged ETA for the first full backup. Contact your Cohesity account team to determine

the appropriate number of application IDs to optimize the performance while ensuring that the Microsoft Tenant Level Throttling limits are not exceeded, which may cause service failures for Microsoft 365.

After completing the first full backup, using multiple application IDs is not usually required. However, in certain situations, a single application ID may be inadequate. For guidance on the appropriate number of application IDs for incremental backups following the initial full backup, Contact your Cohesity account team.

Cohesity also supports Microsoft paid APIs to perform the first full backups faster. These APIs come with higher limits, ensuring the backups are performed faster at a higher API rate. The paid APIs are configured to be used in addition to the existing backup capacity allocated by Microsoft, rather than as a substitute, to help minimize additional costs.

Update Microsoft Organization Setting for Mailbox Size Reporting

By default, Microsoft reports, using Graph API, display information as de-identified names for users, groups, and sites. However, for Mailbox size reporting to work in Cohesity, you need to have identifiable information in the Email activity reports. To do that, you need to disable de-identified names for users, groups, and sites in Microsoft 365 reports.

Update the following organization setting in your Microsoft 365 admin center:

1. Log in to your Microsoft 365 admin center as a Microsoft 365 tenant administrator.
2. Go to **Settings** > **Org settings** > **Services** > **Reports**.
3. In **Reports**, ensure the information is identifiable by deselecting **Display concealed user, group, and site names in all reports**.
4. Click **Save**.

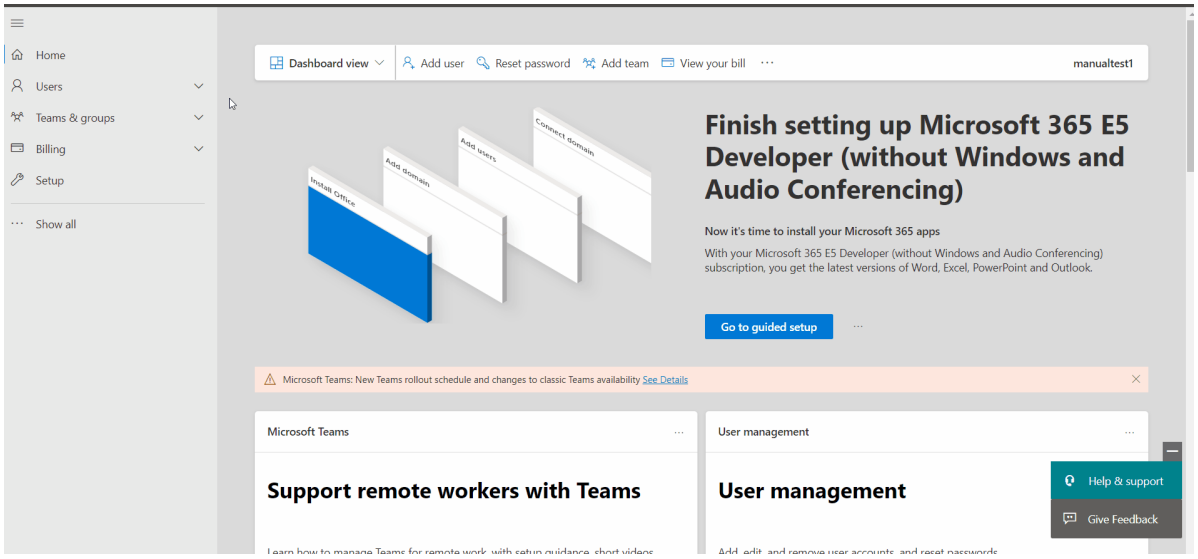

Firewall and Port Requirements for Microsoft 365 Protection

You must open specific ports in the firewall to allow the Cohesity cluster to protect Microsoft 365 applications and the following table lists the required ports. For details on the Microsoft 365 endpoint URLs and IP address ranges, see Office 365 URLs and IP address ranges in Microsoft 365 documentation.

Port	Source	Target	Direction	Network Protocol	Usage
80/443	Cohesity cluster	Microsoft 365 Services	Outgoing	TCP	For Cohesity cluster to communicate with Microsoft Graph and the following service endpoints: • https://accounts.accesscontrol.windows.net • https://<tenant>.onmicrosoft.com • https://graph.microsoft.com/v1.0 • https://graph.microsoft.com/beta • https://outlook.office365.com • https://<tenant>.sharepoint.com • https://<tenant>-admin.sharepoint.com • https://<tenant>-my.sharepoint.com • login.microsoftonline.com
50051	Windows Proxy	Cohesity cluster	Bidirectional	TCP	Used for Cohesity Agent, on the Windows Proxy, and Agentless data transfer to the Cluster. For more information, see Set Up a Windows Connector for Microsoft 365.
80/443	Windows Proxy	Microsoft 365 Services	Outgoing	TCP	For the Cohesity Agent to access the following SharePoint Online service endpoints: • https://accounts.accesscontrol.windows.net • https://<tenant>.onmicrosoft.com • https://<tenant>.sharepoint.com • https://<tenant>-admin.sharepoint.com • https://<tenant>-my.sharepoint.com

Set Up a Windows Connector for Microsoft 365

A Windows Connector is a Microsoft Windows VM that is required by the Cohesity cluster for downloading your backed-up mailboxes or email(s) in PST format.

You need to register a Windows Connector as a physical server source on Cohesity and connect it with your Microsoft 365 source. The following is the supported specification for the Windows Connector:

Specification	
Operating System	Microsoft Windows 10, 2016, 2019, or 2022.
vCPU	2
RAM	8 GB (minimum)
Disk	20 GB

Prerequisites

Ensure that:

- Cohesity Windows Agent is installed on Windows Connector.

  > **Note:** You do not have to install the Volume CBT or File CBT kernel drivers.

- Cohesity recommends that the Windows Connector is dedicated to Microsoft 365 data protection and is not part of any other workflows provided by Cohesity.
- The Windows Connector VM has internet connectivity.
- The Windows Connector VM is in English locale. Non-English locale is not supported.
- If your Cohesity cluster is joined to the Active Directory (AD) domain, ensure that the Windows Connector is also part of the AD domain.
- You exclude the following directory paths from Windows Security:
  - `C:\Program Files\Cohesity`
  - `C:\ProgramData\Cohesity`
  - `C:\Windows\Temp`
- You open specific ports in the firewall to allow data transfer between the Windows Connector and the Cohesity cluster.

Points to note:

- A Windows Connector can be used for a maximum of five Microsoft 365 tenants.
- To download your backed-up mailboxes or email(s) in PST format, ensure that you have installed Microsoft Outlook 64-bit version on the Windows Connector. From your Microsoft 365 homepage, you can install Office. For detailed instructions, see Microsoft Documentation on Office installation. After you have installed the Microsoft Outlook 64-bit version, ensure that you restart the Cohesity Windows Agent. To restart the Cohesity Windows Agent, use the command-line or PowerShell as administrator and run the following commands:

```
net stop cohesityagent

net start cohesityagent
```

Register a Windows Connector as a Physical Server Source

To register a Windows Connector as a physical server source on Cohesity, see Register or Edit a Physical Server.

> **Note:** After registering a Windows Connector as a physical server source on Cohesity, make sure to toggle on **Use Windows Connector** when registering or updating the Microsoft 365 source.

Microsoft 365 Certificate-Based Authentication

Cohesity supports Certificate-Based Authentication (CBA) to protect Microsoft 365 applications. This feature allows you to authenticate with an X.509 certificate and provides phishing resistant authentication. For more information on CBA, see Microsoft Entra Certificate-Based Authentication.

To register a new Microsoft 365 domain using CBA, see one of the following as required:

- Express Registration for Microsoft 365 Sources Using Certificate-Based Authentication
- Manual Registration for Microsoft 365 Sources Using Certificate-Based Authentication

If you have already registered a Microsoft 365 source using app-secret-based authentication, you can switch it to CBA. To do this, you need to migrate the Azure applications used for Microsoft 365 registration to CBA. You can use either of the following workflows:

- Manual migration workflow
- Express migration workflow

Prerequisites

- For Cohesity to automatically renew certificates, the Azure applications must be self-owned by Cohesity and has the **Application.ReadWrite.OwnedBy** permission.

  For more details, see Microsoft Graph permissions reference.

Considerations

Before you proceed with migration, review the following considerations:

- An error may occur during Microsoft 365 source registration or migration if the certificate selected on the source registration page is not uploaded to the respective Azure application in the Azure portal.
- Certificates generated by Cohesity are valid for one year. By default, Cohesity attempts to renew certificates 60 days before their expiration date. If the renewal attempt fails, alerts and warnings are displayed 45 days before the certificates expire.
- The certificate display name can be up to 32 characters long, and the maximum permitted validity is 397 days.
- When importing custom certificates, the certificates must be uploaded to the application before clicking **Register** in the UI, which will otherwise cause the registration to fail.

Certificate Generated Locally vs Certificate Generated in Cohesity Certificate Store

Certificate Generated Locally	Certificate Generated in Cohesity Certificate Store
If you choose to generate a certificate locally and use it for registration, follow this workflow: • Generate the certificate locally. The certificate display name can be up to 32 characters long, and the maximum permitted validity is 397 days. For more information on guidelines to generate certificates, see Certificates. • Map the certificate to the respective Azure application in the Azure portal. • Assign one additional permission to the Azure application, **Application.ReadWrite.OwnedBy**. • Import the certificate into the Cohesity Certificate Store. • Use this certificate during registration. The certificate type is displayed as **Imported**. **Note:** • Certificates generated locally are not renewed automatically; you must renew the certificate manually before it expires. • When importing the custom certificates, the certificates must be uploaded to the Azure application before clicking **Register** in the UI, which will otherwise cause the registration to fail.	If you choose to generate a certificate in the Cohesity Certificate Store and use it for registration, follow this workflow: • Generate the certificate directly in the Cohesity Certificate Store. • You download the certificate. • Map the certificate to the respective Azure application in the Azure portal. • Assign one additional permission to the Azure application, **Application.ReadWrite.OwnedBy**. • You use this certificate during registration. The certificate type is displayed as **Generated**. **Note:** Certificates generated by Cohesity are valid for one year. By default, Cohesity attempts to renew certificates 60 days before their expiration date. If the renewal attempt fails, alerts and warnings are displayed 45 days before the certificates expire.

Migrate Existing Azure Applications to Certificate-Based Authentication (Manual migration workflow)

> **Note:**
> - The express migration workflow generates a single certificate that is mapped to all Azure applications used for Microsoft 365 registration. To use different certificates for individual applications, use the manual migration method.
> - Certificates generated by Cohesity are valid for one year. By default, Cohesity attempts to renew certificates 60 days before their expiration date. If the renewal attempt fails, alerts and warnings are displayed 45 days before the certificates expire.
> - Certificates generated locally are not renewed automatically; you must renew the certificate manually before it expires.

To migrate existing Azure applications from app-secret-based authentication to CBA (Manual migration workflow):

1. In **Cloud Protection Service**, navigate to **Sources**.
2. On the **Sources** page, click the actions menu next to the required registered Microsoft 365 domain and then click **Edit**.
3. On the **Register Microsoft 365 Source** page, under the **Authentication Options** section, click **Certificate**.
4. Click **Migrate existing apps**.
5. In the **Migrate Azure Application** form, copy the device code and click the **Microsoft Azure App** link to open the Microsoft Azure App authorization service form in a new tab.
6. In the Microsoft Azure Application authorization service form, paste the copied code and click **Next**.
7. Log in to Microsoft Azure, enter the **Username** and **Password** of your Microsoft 365 account and click **Sign in**.

   > **Note:** Ensure that your Microsoft 365 account has Global administrator access.

8. Click **Continue**.
9. Wait for Microsoft Azure Authorization to complete, and then click **Update**.

The express migration workflow automatically:

- Generates a certificate.
- Uploads the generated certificate to all the Azure applications used for Microsoft 365 registration.
- Assigns the **Application.ReadWrite.OwnedBy** permission to the Azure applications.
- Populates the generated certificate and the Azure application on the **Register Microsoft 365 Source** page.

Migrate Existing Azure Applications to Certificate-Based Authentication (Express migration workflow)

To migrate existing Azure applications from app-secret-based authentication to CBA using express migration workflow:

1. In **Cloud Protection Service**, navigate to **Sources**.
2. On the **Sources** page, click the actions menu next to the required registered Microsoft 365 domain and then click **Edit**.
3. On the **Register Microsoft 365 Source** page, under the **Authentication Options** section, click **Certificate**.
4. In the **Azure application** section, the App ID is automatically populated. If required, you can enter the App ID of a different Azure Application to use for this registration. From the **Associated Certificate** drop-down list, select the required certificate.

   You can generate certificates within the Cohesity Certificate Store and use them for registration. Alternatively, you can generate a certificate locally, import it to the Cohesity Certificate Store, and then map it to the required Azure application.

   For more information on certificate types, see Certificate Generated Locally vs Certificate Generated in Cohesity Certificate Store.
5. Click **Update**.

Register Microsoft 365 Sources

To protect data in Microsoft 365 applications, you need to register your Microsoft 365 domain as a source in Cohesity. Cohesity uses the Microsoft Graph API to discover, back up, and recover Microsoft 365 objects in your domain. To call the Graph API, Cohesity uses an Azure application that is registered in the Azure portal with the required permissions.

> **Note:**
> - Cohesity cluster supports a maximum of 10 Microsoft 365 Azure Applications (10 App ID registrations).
> - The terms App Secret, App Secret Key, and Client Secret are used interchangeably in the UI and documentation.

You can register your Microsoft 365 domain using one of the following workflows:

Registration Workflows and Authentication Methods

When registering your Microsoft 365 domain with Cloud Protection Service, you must decide which registration workflow and authentication method to use.

Workflow/Authentication Method	Description
Registration workflows	You can use one of the following workflows: • **Manual Registration:** In this workflow, you create the required number of Azure applications in your Azure portal, assign the recommended permissions to the applications, and then provide the application details such as App ID and Client Secret in the Cloud Protection Service UI while registering the source. • **Express Registration:** In this workflow, Cloud Protection Service automatically creates the Azure application in your Azure portal and assigns the required permissions. The same application is then used for registration.
Authentication methods	You can use one of the following authentication methods for the Azure application: • **Client Secret Authentication:** The Azure application authenticates using an App ID and a Client Secret. • **Certificate-Based Authentication (CBA):** The Azure application authenticates using an X.509 certificate instead of a Client Secret, providing phishing-resistant authentication.

The following combinations of workflow and authentication method are available for registration:

- Express Registration for Microsoft 365 Sources Using Client Secret
- Manual Registration for Microsoft 365 Sources Using Client Secret
- Express Registration for Microsoft 365 Sources Using Certificate-Based Authentication
- Manual Registration for Microsoft 365 Sources Using Certificate-Based Authentication

EWS to Graph API Migration for User and Shared Mailboxes

The user and shared primary mailboxes can be transitioned to use Microsoft Graph APIs instead of Exchange Web Services (EWS) APIs for backup and recovery.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Considerations

Before you enable EWS to Graph API Migration for user and shared mailboxes, you must review the following considerations.

- Microsoft is deprecating EWS in April 2027. To continue using EWS until then, you must configure your tenant by October 1, 2026. For more information, see Microsoft's EWS deprecation announcement.
- When the EWS-to-Graph API migration feature is enabled on your Cohesity cluster, Cohesity checks whether all Azure applications registered with the source have the required Graph API permissions.
- After transitioning from EWS APIs to Microsoft Graph APIs, the following mailbox folders will no longer be backed up because Microsoft Graph APIs do not support them.
    - Primary Mailbox folders
    - Archive Mailbox folders
- Cohesity now supports Exchange Online mailbox backup and recovery using Microsoft Graph APIs, replacing Exchange Web Services (EWS) APIs. This transition ensures seamless and uninterrupted mailbox protection while providing greater flexibility and compatibility. Existing backups remain fully usable, with no need for data re-ingestion. The EWS to Graph transition supports only user and shared primary mailboxes in this release. Teams mailboxes, Groups mailboxes and user/shared mailboxes that protect the in-place archive continue to use EWS.
- Data that was previously backed up from these folders using EWS remains available for recovery from existing snapshots taken before the Graph transition. New snapshots taken after the transition will not include these folders.

Prerequisites

Before starting, ensure all Azure applications registered with your Microsoft 365 source have the required Graph API permissions. For the full list of permissions, see Register Custom App in Microsoft Entra ID. The following new permissions must be added:

- MailboxItem.ImportExport.All
- MailboxFolder.ReadWrite.All

- MailboxItem.Read.All
- Mail.Read
- Tasks.Read.All
- Calendars.Read
- Contacts.Read

Add Graph API permissions to an existing Microsoft 365 source

1. Navigate to **Sources** and click **Register** > **Microsoft 365**.
2. Click the actions menu (⋮) next to the **Microsoft 365 source** and then click **Edit**.
3. If Graph API permissions are missing from any Azure application, the following warning banner appears in the Azure application ID section:

   "This source is not yet ready to migrate APIs for Mailbox from EWS to Graph due to missing permissions."
4. Do one of the following as required:
   1. To add permissions automatically (recommended): Click **Update App Permissions**. This initiates the Express workflow using a device code flow to add all required Graph permissions to each application ID.
   2. To add permissions manually: In the Azure portal, navigate to **App Registrations** and add the required Graph API permissions to each registered App ID. For the complete list of permissions, see the **Permission Descriptions** topic in Register Custom App in Microsoft Entra ID. You can also click the **Learn More** link in the banner to view the permissions page.
5. Click **Register** to save the updated source configuration.
6. Click the actions menu (⋮) next to the Microsoft 365 source and select **Refresh**.

Verify that the warning banner is no longer displayed on the source. The source is ready for the Graph API migration. The next incremental backup for each mailbox will run a transition backup (using both EWS and Graph APIs). Subsequent backups will use Graph APIs only.

Express Registration for Microsoft 365 Sources Using Client Secret

Following is the process overview of express registration using Client Secret:

- Cohesity auto-creates the Azure application in your Azure tenant.

  For a description of each application permissions that is assigned, see Permission Descriptions.
- Cohesity auto-assigns the required permissions to that application.
- The App ID and Client Secret of that application are used for authentication.

Use this method if you want Cohesity to handle Azure application creation and permission assignment for you, with a Client Secret used for authentication.

Before you register your Microsoft 365 domain, ensure that you have:

- Add Roles to Microsoft 365 Service User Account.
- Update Microsoft Organization Setting for Mailbox Size Reporting.
- Reviewed the considerations provided for each Microsoft 365 application.

Express Registration automatically creates Azure application with all required permissions, including the Graph API permissions needed for mailbox backup. If your tenant has the EWS-to-Graph migration feature enabled (Early Access), Express Registration is the recommended method to register a new source.

To register your Microsoft 365 domain using Client Secret (Express workflow):

1. Navigate to **Data Protection** > **Sources**. Click **Register** in the upper-right corner of the page > **Microsoft 365**.
2. On the **Register Microsoft 365 Source** page, select checkboxes of the required Microsoft 365 Applications to register.

   > **Note:**
   > Discovery selection change is not allowed for applications with protected objects.
   >
   > If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.
   >
   > This is an Early Access feature. Contact your Cohesity account team to enable the feature.

3. *[Optional]* You can enable the below options based on your requirement:

Options	Description
Fetch Mailbox Info	Toggle ON this option to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage. **Note:** • Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated. • You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.
Fetch OneDrive Info	Toggle ON this option to fetch and process the OneDrive information including the provisioning status and storage quota.
Include Users without MySite	Toggle ON this option to include users who have unprovisioned OneDrive or do not have MySite.
Enable Site Tagging	Toggle ON this option to tag SharePoint Sites whether they are a Group Site or a Teams Site. **Note:** • Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection. • Site tagging can be enabled by default. Sites tagged to Teams or Groups can be directly protected during SharePoint protection if Teams or Groups workflow is not enabled. This ensures that no Teams Sites or Groups Sites are missed during recovery.

4. In the **Account Credentials** section, enter the Microsoft 365 username with a valid SharePoint license.
5. In the **Authentication Options** section, select **App Secret**.
6. In the **Number of Azure Applications Required** section, enter the number of Azure applications that you want to create based on your requirements, and then click **Create Apps**.

> **Note:** For the better management of Microsoft 365 throttling, Cohesity recommends at least one Azure app.

7. In the **Add Azure Application** form, click the **You can also add Azure App manually** link.
8. Enter the **App ID** and **App Secret** of your Azure application.
9. Log in to Microsoft Azure, enter the **Username** and **Password** of your Microsoft 365 account and click **Sign in**.

   > **Note:** Ensure that your Microsoft 365 account has a global administrator access.

10. Follow the instructions to complete the authorization on the Microsoft Azure portal. Then, wait for Microsoft Azure Authorization to complete.
11. Optional. To enable exporting and downloading mailboxes or email(s) in PST format, toggle on **Use Windows Connector** and specify the hostname (FQDN) or server IP of the Windows Connector added as a Physical server.

    For more information, see Set Up a Windows Connector for Microsoft 365.
12. Click **Register**.

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Manual Registration for Microsoft 365 Sources Using Client Secret

Following is the process overview of manual registration using Client Secret:

- You need to create the Azure application in your Azure portal.
- You assign the recommended permissions to that application.

  For a description of each application permissions that is assigned, see Permission Descriptions.
- While registering the Microsoft 365 domain as a source in Cohesity , you need to provide the App ID and Client Secret of that application on the Cohesity UI.

Use this method if you want full control over the Azure application that Cohesity uses, and you are comfortable managing the application and its Client Secret in your own Azure portal.

Before you register your Microsoft 365 domain using this method, ensure that you have:

- Register Custom App in Microsoft Entra ID with the required Graph API, and also noted the Application ID and Client Secret.

  For a description of each application permissions that is assigned, see Permission Descriptions.
- Add Roles to Microsoft 365 Service User Account.
- Update Microsoft Organization Setting for Mailbox Size Reporting
- Review the considerations provided for each Microsoft 365 application.

To register your Microsoft 365 domain manually using Client Secret authentication:

1. Navigate to **Data Protection** > **Sources**. Click **Register** in the upper-right corner of the page > **Microsoft 365**.
2. On the **Register Microsoft 365 Source** page, select checkboxes of the required Microsoft 365 Applications to register.

   > **Note:**
   > Discovery selection change is not allowed for applications with protected objects.
   >
   > If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.
   >
   > This is an Early Access feature. Contact your Cohesity account team to enable the feature.

3. *[Optional]* You can enable the below options based on your requirement:

Options	Description
Fetch Mailbox Info	Toggle ON this option to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage. **Note:** • Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated. • You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.
Fetch OneDrive Info	Toggle ON this option to fetch and process the OneDrive information including the provisioning status and storage quota.
Include Users without MySite	Toggle ON this option to include users who have unprovisioned OneDrive or do not have MySite.
Enable Site Tagging	Toggle ON this option to tag SharePoint Sites whether they are a Group Site or a Teams Site. **Note:** • Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection. • Site tagging can be enabled by default. Sites tagged to Teams or Groups can be directly protected during SharePoint protection if Teams or Groups workflow is not enabled. This ensures that no Teams Sites or Groups Sites are missed during recovery.

4. In the **Account Credentials** section, enter the full username of the Microsoft 365 user account with a valid SharePoint license.
5. In the **Authentication Options** section, select **App Secret**.
6. In the **Number of Azure Applications Required** section, enter the number of Azure applications that you want to create based on your requirements, and then click **Create Apps**.

> **Note:** For the better management of Microsoft 365 throttling, Cohesity recommends at least one Azure app.

7. In the **Add Azure Application** form, click the **You can also add Azure App manually** link.
8. Enter the **App ID** and **App Secret** of your Azure application.
9. Optional. To enable exporting and downloading mailboxes or email(s) in PST format, toggle on **Use Windows Connector** and specify the hostname (FQDN) or server IP of the Windows Connector added as a Physical server.

   For more information, see Set Up a Windows Connector for Microsoft 365
10. Click **Register**.

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Express Registration for Microsoft 365 Sources Using Certificate-Based Authentication

Certificate-Based Authentication (CBA) lets you authenticate the Azure application using an X.509 certificate instead of a Client Secret, providing phishing resistant authentication.

Following is the process overview of express registration using certificate-based authentication:

- Cohesity auto-creates the Azure application in your Azure portal.
- Cohesity auto-generates a certificate in the Cohesity Certificate Store, and then map it to the Azure application in your Azure portal.
- Cohesity auto-assigns the application the additional permissions required for certificate-based authentication, **Application.ReadWrite.OwnedBy**.

Note:

- Certificates generated by Cohesity are valid for one year. By default, Cohesity attempts to renew certificates 60 days before their expiration date. If the renewal attempt fails, alerts and warnings are displayed 45 days before the certificates expire.
- This method generates a single certificate and maps it to all Azure applications used for Microsoft 365 registration. To use different certificates for individual Azure applications, you need to use the manual registration method.
  See, Manual Registration for Microsoft 365 Sources Using Certificate-Based Authentication.

Before you register your Microsoft 365 domain using this method, ensure that you have:

- Add Roles to Microsoft 365 Service User Account.
- Update Microsoft Organization Setting for Mailbox Size Reporting.
- Reviewed the considerations provided for each Microsoft 365 application.

To register your Microsoft 365 domain using certificate-based authentication with express registration method:

1. Navigate to **Data Protection** > **Sources**. Click **Register** in the upper-right corner of the page > **Microsoft 365**.
2. On the **Register Microsoft 365 Source** page, select checkboxes of the required Microsoft 365 Applications to register.

   Note:
   Discovery selection change is not allowed for applications with protected objects.

   If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.

   This is an Early Access feature. Contact your Cohesity account team to enable the feature.

3. *[Optional]* You can enable the below options based on your requirement:

Options	Description
Fetch Mailbox Info	Toggle ON this option to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage. **Note:** • Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated. • You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.
Fetch OneDrive Info	Toggle ON this option to fetch and process the OneDrive information including the provisioning status and storage quota.
Include Users without MySite	Toggle ON this option to include users who have unprovisioned OneDrive or do not have MySite.
Enable Site Tagging	Toggle ON this option to tag SharePoint Sites whether they are a Group Site or a Teams Site. **Note:** • Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection. • Site tagging can be enabled by default. Sites tagged to Teams or Groups can be directly protected during SharePoint protection if Teams or Groups workflow is not enabled. This ensures that no Teams Sites or Groups Sites are missed during recovery.

4. In the **Account Credentials** section, enter the full **Username** of the Microsoft 365 user account with a valid SharePoint license.
5. In the **Authentication Options** section, select **Certificate**.
6. In the **Number of Azure Applications Required** section, enter the number of Azure applications that you want to create based on your requirements, and then click **Create Apps**.

> **Note:** For the better management of Microsoft 365 throttling, Cohesity recommends at least one Azure app.

7. In the **Add Azure Application** form, copy the device code and click the **Microsoft Azure App** link to open the Microsoft Azure App authorization form. This form opens in a new tab.
8. In the **Microsoft Azure App authorization** form, paste the copied code and click **Next**.
9. Log in to Microsoft Azure, enter the **Username** and **Password** of your Microsoft 365 account and click **Sign in**.

   > **Note:** Ensure that your Microsoft 365 account has Global administrator access.

10. Wait for Microsoft Azure Authorization to complete.

    The App ID and the associated mapped certificate are populated on the registration page. The type of the certificate is displayed as **Generated**.
11. Optional. To enable exporting and downloading mailboxes or email(s) in PST format, toggle on **Use Windows Connector** and specify the hostname (FQDN) or server IP of the Windows Connector added as a Physical server.

    For more information, see Set Up a Windows Connector for Microsoft 365.
12. Click **Register**.

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Manual Registration for Microsoft 365 Sources Using Certificate-Based Authentication

Certificate-Based Authentication (CBA) lets you authenticate the Azure application using an X.509 certificate instead of a Client Secret, providing phishing resistant authentication.

In the manual registration workflow using CBA, you create the Azure application in your own Azure portal. Depending on where you want the certificate to originate, you can use one of the following two methods to generate a certificate:

Certificate Generated Locally	Certificate Generate in the Cohesity Certificate Store
To continue with this workflow, you need to: • Generate the certificate locally. The certificate display name can be up to 32 characters long, and the maximum permitted validity is 397 days. For more information on guidelines to generate Certificates. For more information, see Certificates. • Map the certificate to your Azure application in the Azure portal. • Assign one additional permission to the Azure application, **Application.ReadWrite.OwnedBy**. • Import the certificate into the Cohestity certificate store. • Use this certificate during registration. **Note:** • Certificates generated locally are not renewed automatically; you must renew the certificate manually before it expires. • The certificate must be uploaded on the app before clicking **Register** in the UI, which will otherwise cause the registration to fail.	To continue with this workflow, you need to: • Generate the certificate directly in the Cohesity certificate store. • You download the certificate. • Map the certificate to your Azure application in the Azure portal. • Assign one additional permission to the Azure application, **Application.ReadWrite.OwnedBy**. • You use this certificate during registration. **Note:** Certificates generated by Cohesity are valid for one year. By default, Cohesity attempts to renew certificates 60 days before their expiration date. If the renewal attempt fails, alerts and warnings are displayed 45 days before the certificates expire.

Registration Using CBA Manually (Use Locally Generated Certificate)

To continue with this workflow, you need to:

- Generate the certificate locally.

  The certificate display name can be up to 32 characters long, and the maximum permitted validity is 397 days. For more information on guidelines to generate certificates, see Certificates.
- Map the certificate to your Azure application in the Azure portal.
- Assign one additional permission to the Azure application, **Application.ReadWrite.OwnedBy**.
- Import the certificate into the Cohesity Certificate Store.
- Use this certificate during registration.

> **Note:**
> - Certificates generated locally are not renewed automatically; you must renew the certificate manually before it expires.
> - The certificates must be mapped to the Azure application before clicking **Register** in the UI, which will otherwise cause the registration to fail.

Before you register your Microsoft 365 domain using either method, ensure that you have:

- Register Custom App in Microsoft Entra ID with the required Graph API, and also noted the Application ID and Client Secret.
- Add Roles to Microsoft 365 Service User Account.
- Update Microsoft Organization Setting for Mailbox Size Reporting
- Review the considerations provided for each Microsoft 365 application.
- Generated a certificate.

  For more information, see Certificates.
- Mapped the certificate to your Azure application in the Azure portal.
- Assigned the **Application.ReadWrite.OwnedBy** permission to the Azure application.
- Imported the generated certificate to the Cohesity certificate store.

To register your Microsoft 365 domain manually by importing a certificate you created locally:

1. Navigate to **Data Protection** > **Sources**. Click **Register** in the upper-right corner of the page > **Microsoft 365**.
2. On the **Register Microsoft 365 Source** page, select checkboxes of the required Microsoft 365 Applications to register.

> **Note:**
> Discovery selection change is not allowed for applications with protected objects.
>
> If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.
>
> This is an Early Access feature. Contact your Cohesity account team to enable the feature.

3. *[Optional]* You can enable the below options based on your requirement:

Options	Description
Fetch Mailbox Info	Toggle ON this option to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage. **Note:** • Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated. • You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.
Fetch OneDrive Info	Toggle ON this option to fetch and process the OneDrive information including the provisioning status and storage quota.
Include Users without MySite	Toggle ON this option to include users who have unprovisioned OneDrive or do not have MySite.

Options	Description
Enable Site Tagging	Toggle ON this option to tag SharePoint Sites whether they are a Group Site or a Teams Site. **Note:** • Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection. • Site tagging can be enabled by default. Sites tagged to Teams or Groups can be directly protected during SharePoint protection if Teams or Groups workflow is not enabled. This ensures that no Teams Sites or Groups Sites are missed during recovery.

4. In the **Account Credentials** section, enter the full **Username** of the Microsoft 365 user account with a valid SharePoint license.
5. In the **Authentication Options** section, select **Certificate**.
6. In the **Number of Azure Applications Required** section, enter the number of Azure applications that you want to create based on your requirements, and then click **Create Apps**.

   > **Note:** For the better management of Microsoft 365 throttling, Cohesity recommends at least one Azure app.

7. In the **Add Azure Application** form, click the **You can also add Azure App manually** link.
8. Enter the **App ID** and then select the required certificate from the **Associated Certificate** list.

   The certificate type is displayed as **Imported**.
9. Click **Register**.

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Registration Using CBA Manually (Use Certificate Generated in Certificate Store)

To continue with this workflow, you need to:

- Generate the certificate directly in the Cohesity Certificate Store.
- You download the certificate.
- Map the certificate to your Azure application in the Azure portal.

- Assign one additional permission to the Azure application, **Application.ReadWrite.OwnedBy**.
- You use this certificate during registration.

> **Note:** Certificates generated by Cohesity are valid for one year. By default, Cohesity attempts to renew certificates 60 days before their expiration date. If the renewal attempt fails, alerts and warnings are displayed 45 days before the certificates expire.

Before you register your Microsoft 365 domain using either method, ensure that you have:

- Register Custom App in Microsoft Entra ID with the required Graph API, and also noted the Application ID and Client Secret.
- Add Roles to Microsoft 365 Service User Account.
- Update Microsoft Organization Setting for Mailbox Size Reporting
- Review the considerations provided for each Microsoft 365 application.
- Generated a certificate in the Cohesity Certificate Store.
- Downloaded the certificate and mapped it to the Azure application in your Azure Portal.
- Assigned the **Application.ReadWrite.OwnedBy** permission to the Azure application.

To register your Microsoft 365 domain manually using the certificate generated in Cohesity Certificate Store:

1. Navigate to **Data Protection** > **Sources**. Click **Register** in the upper-right corner of the page > **Microsoft 365**.
2. On the **Register Microsoft 365 Source** page, select checkboxes of the required Microsoft 365 Applications to register.

> **Note:**
> Discovery selection change is not allowed for applications with protected objects.
>
> If the **Private Chats** and **Teams Posts** option is enabled under the **Mailbox** and **Teams** apps respectively, the Private Chats and Teams Posts will be backed up along with the corresponding Users and Teams respectively.
>
> This is an Early Access feature. Contact your Cohesity account team to enable the feature.

3. *[Optional]* You can enable the below options based on your requirement:

> **Note:** Retaining the default values will speed up object discovery in the environment. The metadata required (which is fetched through these options) will be fetched during the backup of the objects and updated.

- **Fetch Mailbox Info** to fetch and process the Mailbox information including the provisioning status, mailbox type, and in-place archival usage.

  > **Note:** You can enable this option to discover the Mailboxes of the users that were converted into Shared Mailboxes by revoking the user Exchange Online licenses.

- **Fetch OneDrive Info** to fetch and process the OneDrive information including the provisioning status and storage quota.
- **Include Users without MySite** to include users who have unprovisioned OneDrive or do not have MySite.
- **Enable Site Tagging** to tag SharePoint Sites whether they are a Group Site or a Teams Site.

  > **Note:** Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection.

> **Note:** Site tagging can be enabled by default. Sites tagged to Teams or Groups can be directly protected during SharePoint protection if Teams or Groups workflow is not enabled. This ensures that no Teams Sites or Groups Sites are missed during recovery.

4. In the **Account Credentials** section, enter the full **Username** of the Microsoft 365 user account with a valid SharePoint license.
5. In the **Authentication Options** section, select **Certificate**.
6. In the **Number of Azure Applications Required** section, enter the number of Azure applications that you want to create based on your requirements, and then click **Create Apps**.

   > **Note:** For the better management of Microsoft 365 throttling, Cohesity recommends at least one Azure app.

7. In the **Add Azure Application** form, click the **You can also add Azure App manually** link.
8. Enter the **App ID** and then select the required certificate from the **Associated Certificate** list.

   The certificate type is displayed as **Generated**.
9. Optional. To enable exporting and downloading mailboxes or email(s) in PST format, toggle on **Use Windows Connector** and specify the hostname (FQDN) or server IP of the Windows Connector added as a Physical server.

   For more information, see Set Up a Windows Connector for Microsoft 365.
10. Click **Register**.

You can follow the Microsoft 365 source discovery and registration progress on the **Sources** page.

Manage Microsoft 365 Source Registration

After registering your Microsoft 365 domain as a source, you can:

- Update the Microsoft 365 source configuration.
- Refresh the source details.
- Unregister the Microsoft 365 domain from Cohesity.

Update the Microsoft 365 Source Configuration

After registering your Microsoft 365 source on Cohesity, you might have changed the Microsoft 365 domain configuration by:

- Changing the credentials of Microsoft 365 user account credentials.
- Updating the app secret by adding more permissions to the custom app.
- Creating a new app ID.
- Configure the Windows Proxy to enable export and download of your backed-up mailboxes or email(s) in PST format.

You can update the Microsoft 365 details provided during the registration process with the latest Microsoft 365 configuration details.

To edit the Microsoft 365 source configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Edit**.
3. On the **Update Microsoft 365 Source** page, update the required configurations.
4. Click **Register**.

Refresh the Microsoft 365 Source

You can refresh the Microsoft 365 domain configuration and fetch the latest changes on the Microsoft 365 domain.

To refresh the Microsoft 365 Source Configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Refresh**.

Unregister the Microsoft 365 Domain

If you plan to stop taking the backup of your Microsoft 365 domain, you can unregister the Microsoft 365 source from the Cohesity Cluster.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the Microsoft 365 source.

To unregister the Microsoft 365 domain:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Unregister**.

Exchange Online Mailboxes

Microsoft Exchange Online is a SaaS application that is bundled in your Microsoft 365 subscription service. It is a hosted messaging solution that delivers the capabilities of Microsoft Exchange Server as a cloud-based service. It gives users access to email, calendar, contacts, and tasks from PCs, the web, and mobile devices. Using the policy-

based data protection solution from Cohesity, you can protect Exchange Online data on Microsoft 365.

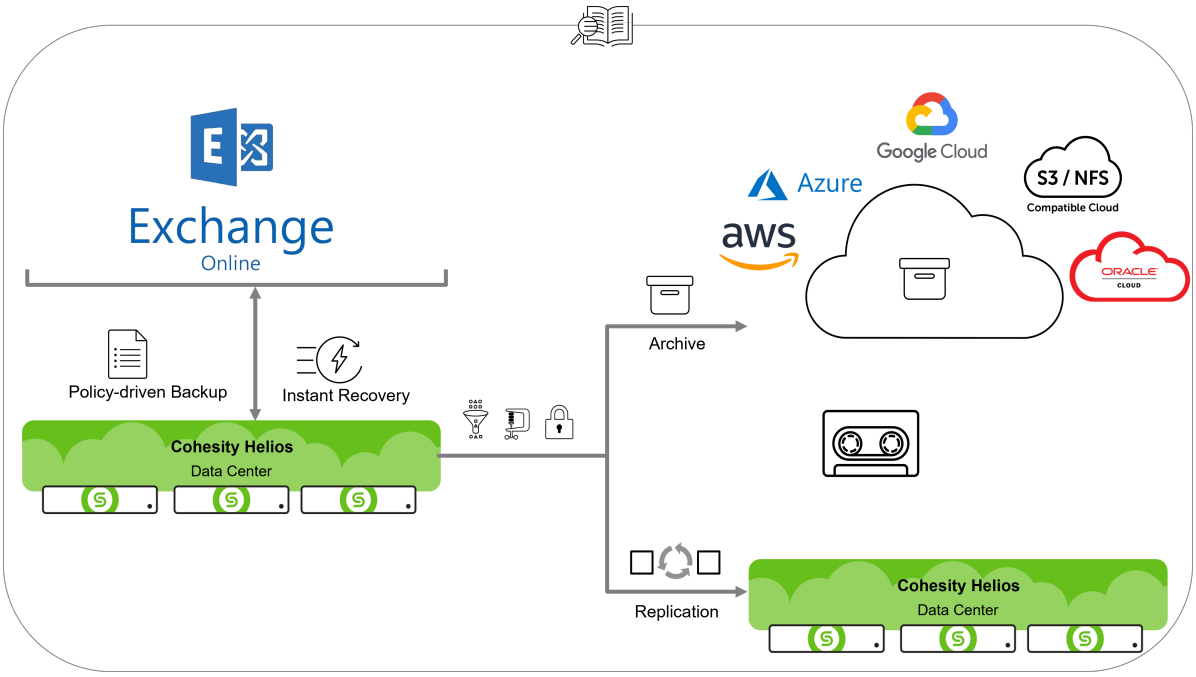

Cohesity now supports Exchange Online mailbox backup and recovery using Microsoft Graph APIs, replacing Exchange Web Services (EWS) APIs. This transition ensures seamless and uninterrupted mailbox protection while providing greater flexibility and compatibility. Existing backups remain fully usable, with no need for data re-ingestion. The EWS to Graph transition supports only user and shared primary mailboxes in this release. Teams mailboxes, Groups mailboxes and user/shared mailboxes that protect the in-place archive continue to use EWS.

If your Cohesity cluster has the EWS-to-Graph migration feature enabled, Express Registration is the recommended method to register a new source.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Plan and Prepare for Mailbox Protection

Before you register your Microsoft 365 domain with Cohesity to protect Microsoft 365 data, ensure:

- You have completed all the Prerequisites.
- You have reviewed and understood the Considerationsand Best Practices.

> **Important:** By September 2026, Cohesity will migrate from Exchange Web Services (EWS) APIs to Microsoft Graph APIs to protect Microsoft 365 Applications. This is as per the Retirement of Exchange Web Services in Exchange Online announcement from Microsoft. For more information, see Deprecation of Exchange Web Services in Exchange Online.

As part of the migration to Microsoft Graph APIs, Cohesity now supports Exchange Online mailbox backup and recovery using Microsoft Graph APIs as an Early Access feature, replacing Exchange Web Services (EWS) APIs. This transition ensures seamless and uninterrupted mailbox protection while providing greater flexibility and compatibility. Existing backups remain fully usable, with no need for data re-ingestion. The EWS to Graph transition supports only user and shared primary mailboxes in this release. Teams mailboxes, Groups mailboxes and user/shared mailboxes that protect the in-place archive continue to use EWS.

Considerations

Review and understand the following considerations before you protect your Microsoft 365 Exchange Online data:

- Recovery can be performed at the mailbox, folder, or item level. Currently, recovery at the Protection Group level is not supported.
- Recovery of items in the same mailbox from different recovery points is not supported in the same recovery task. You need to create different recovery tasks for different recovery points.
- If the email already exists at the destination during the recovery workflow, the email will be updated
- The backup or recovery of Mailboxes, OneDrives, SharePoint Online Sites, and MS Teams are executed in parallel as long as there are enough resources on the Cohesity cluster. However, the backup or recovery operations are serialized for the Mailboxes and OneDrives that belong to the same user, and the operations execute in first-in-first-out order.

  For example, if the backup of a user's mailbox is in progress, then the backup of the user's OneDrive starts only after the backup of the mailbox completes.
- Converted-shared mailboxes are not backed up by default. To enable backup of converted-shared mailboxes, contact Cohesity Support.
- The Retirement of RBAC Application Impersonation in Exchange Online has no known impact on the Microsoft 365 Exchange Online Mailboxes, Teams, and Groups protection workflow on the Cohesity cluster.

- By default, the Conversation History/Team Chat folder is excluded during the Exchange Online backup.
- Following the Microsoft changes for EWS APIs, Cohesity no longer backs up mailbox items of the type **IPM.Teams.SkypeMessage**. For more details, see the KB Article.
- Exchange Online Mailboxes protection is supported only for users present on Azure AD and not for users on the on-premises Active Directory.
- Exchange Online Mailboxes protection is:

  Supported for:

  - Users present in Azure AD
  - Hybrid setup with users present on Azure AD

  Not supported for:

  - Users on on-premises Active Directory
  - Users natively on on-premises Active Directory that are synchronized to Azure AD
- Microsoft 365 items with only special characters will not be listed in granular search. The item name must contain at least one alphanumeric character to be listed in the search result.
- Cohesity supports protecting 100k Microsoft 365 objects. This includes the following:
  - Objects on individual applications, for example, 100k Mailboxes.
  - Objects on combined applications, for example, 50k Mailboxes, 30k Teams, and 20k Groups.

  Validation messages are displayed when backing up more than 7500 objects and recovering more than 5000 objects.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- Backup and recovery of Recoverable Items:
  - Deletions, Purges, Discovery Holds, and SubstrateHolds folders are currently supported for backup and restore.
- Audits, Calendar Logging, and Versions folders are not supported.
- Backup and download of the following is not supported:
  - Self-message (messages sent to self)
  - Saved or pinned messages property in the conversation

  - Meeting recordings metadata from private chats
  - Emojis and Stickers
- Restoration at the protection group level is not supported for Microsoft 365.
- Recovery for items in the same mailbox from different recovery points is not supported in the same recovery task. You need to create different recovery tasks for different recovery points.
- If the size of the mailbox is large, then exporting the mailbox to PST format takes a longer duration.
- The exported PST of the mailbox or the email is valid for 72 hours. Ensure that you download the PST file within 72 hours of the recovery task completion.
- PST recovery of emails with more than 2000 recipients is not supported.
- The Cohesity cluster and connector must be a member of the same or any trusted AD for PST recovery.
- Backup and recovery of files with sensitivity labels is supported. For files with sensitivity labels,
  - Only a user with access via the label can open the file after downloading, provided the label is not permanently deleted in SharePoint.
  - When recovering to original location, if the file is already present in the destination, the recovery fails. If the file is not present in the destination, a new file will be created.
  - The files downloaded from SharePoint are encrypted. You can download the same file from the backup storage.
- Backup and download of Teams Chat are supported; however, restore is not supported currently.
- Backup and recovery of Exchange Online mailboxes using security groups is supported. You must review the following considerations before protecting mailboxes using a security group:

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  - Each protection group can contain only one security group. Mixing multiple security groups within the same protection group is not supported.
  - The member count for a security group includes only users present in the Entity Hierarchy, regardless of whether they have a valid mailbox or OneDrive license. If a user's license is removed, they remain in the count until the next source refresh filters them out. Members without a valid license will fail individually

during backup, but the protection job proceeds for the remaining members.

- Security group membership is fetched from Microsoft Entra ID. Newly added members are included in the next scheduled refresh; members removed from the group or deleted from Entra ID are automatically excluded after the next refresh. If security group optimization is enabled, members are not fetched during the entity hierarchy refresh and are instead fetched during the backup run.
- Security group protection is supported only for the members of the security group, not the owners of the group.
- Both mail-enabled and mail-disabled security groups are supported.
- Members of nested security groups are automatically included in protection. The member count displayed reflects the unique (deduplicated) count across all levels of nesting.
- If the selected security group has no members, the protection job fails. Jobs resume normally once members are added to the group and the next entity hierarchy refresh completes.
- Remove the protection group manually if the security group is deleted from Entra ID.
- If a user is a member of two security groups that are each assigned to a separate protection group, that user's mailbox or OneDrive is backed up by both protection jobs.
- The security group itself is manually selected for protection; individual members within the group are treated as auto-protected.
- Recovery of mailboxes or OneDrives protected via a security group follows the same workflow as recovery of individually protected objects. There is no change to recovery workflows.
- Backup and recovery of external users from a security group is supported.

Best Practices

To protect your Microsoft 365 data, Cohesity recommends the following best practices:

- Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling failed requests. There is no need to increase the value for this setting in the Cohesity Protection Policy.
- Use **Auto-Protect** to avoid adding new users to the Protection Group manually.
- Avoid adding the same users in multiple Protection Groups.

Backup Exchange Online Mailboxes

Exchange Online mailboxes are protected using a Protection Group and a Protection Policy. A Protection Group enables you to select and group the objects that you want to back up. It

uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To protect Exchange Online mailboxes:

1. Do one of the following:
    - On the **Dashboards** page, from the drop-down at the top of the page, select **Data Protection** and do one of the following:
        - If there is no existing Protection Group, click **Start Data Protection** > **Microsoft 365** > **Mailbox**.
        - Click **Protect** > **Microsoft 365** > **Mailbox**.
    - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Microsoft 365** > **Mailbox**.
    - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Microsoft 365 source and then select **Protect Mailbox**.

      If you want to back up mailboxes using security groups, click the <**Security Group**> icon to view the available security groups. Select the required security group to protect the mailboxes of the users in the security group.

      > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

2. In the **New Protection** pop-up, do one of the following:
    - Protect the mailbox by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
        i. **Add Objects**: Select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 account. From the **Users** list, select the users you plan to back up.
        ii. **Protection Group**: Specify a name for the Protection Group.

iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.

iv. **Storage Domain**: Select the appropriate storage domain.

> **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

v. Click **Protect**.

- To protect the mailbox by specifying all the options, click **More Options**.

3. From the **Source** drop-down list, select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 source.
4. From the **Users** list, select the users that you plan to backup and click **Save Selection**.

> **Note:** You can use the advanced filters that are available to search and find users based on user attributes. Also, you can click the shield icon to enable Auto Protect. For more information on Auto Protect, see Enable or Disable Auto Protect.

5. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name of the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, Create or Use an Existing Protection Policy.
7. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
8. After selecting a Storage Domain, if you need to change any of the advanced settings on the **New Microsoft 365 Protection Group** page, scroll down and click **Edit** on the right. For more information, see Advanced Settings.
9. Click **Protect**.

   Your Exchange Online Data Protection is now active and running.

Advanced Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection run should start. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Exclusion List	Select the folders that you plan to exclude from the backup or click **Add** to add custom folders that you want to exclude from the backup. **Note:** Exclusion list is not supported for Public Folders.
Indexing	By default, indexing is enabled. **Note:** Indexing is mandatory for granular restore of an Exchange Online mailbox , such as restoring a folder or restoring an email. For more information, see Recover Mailbox .

Advance Settings	Description
Alerts and Priority	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	**Optional**. Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

Create a Protection Policy with the settings listed in the Policy Settings for Exchange Online. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Exchange Online

The following table lists the applicable policy settings for Exchange Online:

Field	Description
DataLock	Add a DataLock for compliance and regulatory requirement. This ensures that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users from modifying or deleting any Snapshots that were generated by the protection runs that use this Policy. The users include those who have the Data Security role in Cohesity as well. > **Note:** Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained, and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer duration than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the protection run fails. The default time between retries is thirty minutes. You can customize the number of retries and how long to wait between each attempt. > **Note:** Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling the failed requests. There is no need to increase the value for this setting in the Protection Policy.

Field	Description
Quiet Times	Quiet Times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	Replicate snapshots to another Cohesity cluster.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with applicable policy settings for Exchange Online. In this case, all the policy settings apart from the settings applicable to Exchange Online are ignored.

To use an existing Protection Policy:

1. Navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for Exchange Online. For information on the applicable policy settings, see Policy Settings for Exchange Online.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage the Exchange Online Backup Configuration

In Cohesity, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirement.

> **Note:** You cannot update the registered Microsoft 365 domain and the assigned Storage Domain in the Protection Group settings.

Add Mailboxes to the Protection Group

You can add new mailboxes to the existing Protection Group for backup. The newly added mailbox is backed up from the next protection run and it will be a full backup for the newly added user.

To add mailboxes to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Mailbox** page, click ✎ in the **Users** section.
5. From the **Users** list, select the mailboxes that you plan to add to the backup and click **Save Selection**.
6. Click **Save**.

Remove Mailboxes from the Protection Group

You can remove mailboxes from the Protection Group that you do not want to back up in the next protection run. When you remove a user from the protection group, incremental backup of the user's mailbox is not performed from the next protection run. The existing backup for the user is preserved until the snapshot expires.

To remove mailboxes from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Mailbox** page, click ✎ in the **Users** section.
5. From the **Users** list, unselect the mailboxes that you want to remove from the Protection Group and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new mailboxes that are added to Microsoft 365.

When the Auto Protect option is enabled:

- All the mailboxes that are added to Microsoft 365 in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the mailboxes that are removed from Microsoft 365 in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the user is preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly onboarded users on Microsoft 365 to the Protection Group.
- Remove existing users that are offboarded on Microsoft 365 from the Protection Group. The existing backups of the users are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Mailbox** page, click ✎ in the **Users** section.
5. From the **User** list, click the icon to enable or disable Auto Protect and click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Protect**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up Exchange Online. It provides you with options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.
- Add replication.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.

3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. In the **Backup** drop-down list, select any one of the following:
        - Backup all Objects in the Protection Group
        - Backup only selected Objects in the Protection Group
    2. In the **Backup Type** drop-down list, select the backup type.
    3. Optional. In the **Replicate** section, select the replication target and the retention duration or click **Add** to add the replication target and the retention duration. Also, select **Replicate only fully successful runs** to replicate the data only from the completed protection runs.
    4. Optional. In the **Archive** section, click **Add** to add the external target to which you want to archive the data and the retention duration. Also, select **Archive only fully successful runs** to archive the data only from the completed protection runs.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups, local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details will not be displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - **Delete Protection Group Only**
   - **Delete Protection Group and Snapshots**

Recover Mailbox

You can recover Exchange Online data using Cohesity at various granularities, such as at the mailbox, folder, and email level. You can recover to the same location, alternate location, or across Microsoft 365 domains. Also, you can download the offline copy of the mailboxes or emails in PST file format.

> **Note:** To download backed-up mailboxes or email(s) in PST format, you need to register a Windows Connector as a physical server source on Cohesity and connect it with your Microsoft 365 source. For more information, Set Up a Windows Connector for Microsoft 365.

Recover Mailboxes

To recover mailboxes, you can search through the backups by using username or using a string pattern in the username for the mailbox that you plan to recover and then start the recover operation.

> **Note:** You can use * character to list all the mailboxes.

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To recover mailboxes:

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Mailboxes** as the recovery source.
2. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, select **Entire Mailbox**.
3. Search the user mailbox to be recovered, select the mailboxes you need, and click **Continue**.
4. Based on your requirements, select any one of the following options:

Options	Description
Original Location	To recover data to the same mailbox from which it was backed up.
New Location	To recover data to an alternate location in a different Microsoft 365 domain. Select a **Registered Source** and the **Target Mailbox**, and specify the **Folder name** to which you plan to recover. **Note:** • Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source. • If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.

Options	Description
Export as PST	To export a backed-up mailbox as a PST file, you must provide a password for the exported PST file. **Note:** • Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, seeRecover Mailbox . • The PST export process might take a few minutes or hours depending on the size of the mailbox.
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

5. Select your **Recovery Options**:
    1. **Continue on Error.** Enable to complete the recovery task even if errors occur when recovering Mailboxes.
    2. **Task Name.** Change the default name of the recovery task.
    3. **Selective Recovery**. Toggle ON the **Recover within specified timeframe** option to filter the Mailbox items to be recovered. Under **Include items** in recovery, select the required items:

        Mails, Calendars, Contacts, Tasks, and Notes.

        You can also recover data of the last specific days prior to the selected snapshot.

        - All upcoming calendar events will be recovered.
        - Any calendar events that were scheduled or received in the selected time frame will be recovered.

- Contacts, Tasks, and Notes will be fully recovered unaffected by the days selected.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

4. **Include Recoverable Items**. Toggle ON to **Recover Recoverable Items**.

   The following folders will be recovered:

   - Deletion
   - Purges
   - DiscoveryHolds
   - SubstrateHolds

6. Click **Recover** to start the recover operation.

Individual recovery task is run for each mailbox. Once the recovery is completed, you can click on the **Show Subtasks** button to view the details in the logs.

Recover Mailbox Items

You can recover the mailbox items such as individual emails, folders, calendar items, contacts, notes, or tasks from the backed-up mailboxes.

Cohesity now supports granular recovery of items in Recoverable Items and Archive Recoverable Items folders. Additionally, all four Microsoft 365 mailbox folder roots, Primary Mailbox, In-Place Archive, Recoverable Items, and Archive Recoverable Items are now displayed in the granular recovery workflow, allowing you to select any root individually to browse and recover its items.

Before you proceed with recovery, you must review the considerations. For more information, see Backup Exchange Online Mailboxes.

Recover Emails

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To recover specific emails:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Mailboxes**.
3. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, select **Search Mailbox Items**.
4. From the **Item Type** drop-down list, select **Emails**.

5. Select the root type filter, and select one or more roots to search for items in that. The roots are Primary Mailbox, In-Place Archive, Recoverable Items, and Archive Recoverable Items. At least one root must be selected. The **Primary Mailbox** root is selected by default.
6. To search for emails, enter the * wildcard character in the search bar, or enter text with a matching subject to find a specific email, and then select the emails to recover from the search results.
   Or click **Advanced Search** and use the following filters to search for specific emails:

Filters	Description
Subject	Subject line in the email.
From	Mail sender email address.
To	Mail recipient email address. Use a comma or space separator to enter multiple addresses.
Date Range	Using the calendar, select a specific date range to search the emails.
Email Type	Select one of the email types: • All Emails • Only emails with attachments • Only emails without attachments
cc	The email address in the Cc: line of the email. Use a comma or space separator to enter multiple addresses.
bcc	The email address in the Bcc: line of the email. Use a comma or space separator to enter multiple addresses.
Search in Folder	Search for the email within the specified folder. For example, Inbox, Drafts, and so on. Use a comma or space separator to enter multiple folder names.

7. Select the items that are to be restored.
8. To use a different recovery point, click the **Edit** icon on the tile. Find the recovery point you need and click **Select Recovery Point**.
9. Click **Next: Recover Options**.
10. From the **Recover To** section, select one of the following options based on your requirements:

Recover To options	Description
Original Location	To recover data to the same mailbox from which it was backed up.
New Location	To recover data to a different Microsoft 365 mailbox. Do the following: • Select the required **Registered Source** from the list. Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source • Select the **Target Mailbox** and specify the **Folder name** to which you want to recover the data. If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.
Download as PST	To download the required email as a PST file. To enhance security, you can set a password for the downloaded PST file. Once the recovery task is completed, the exported PST file is valid for 7 days. Ensure that you download the PST within 7 days. For more information, seeRecover Mailbox .
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

11. Select your **Recovery Options**:
    - **Continue on Error**. Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesity will still attempt to recover the other selected emails.
    - **Task Name**. Change the default name of the recovery task.
12. Click **Recover**.
    The system creates a separate recovery task as needed. A sequential suffix number is added to each recovery task name. You can monitor and track these recovery tasks individually on the **Recoveries** page.

Recover Calendar Items

With this procedure you can recover specific calendar items from a protected Microsoft 365 mailbox.Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To recover specific calendar items:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Mailboxes**.
3. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, select **Search Mailbox Items**.
4. From the **Item Type** drop-down list, select **Calendars**.
5. To search for emails, enter the * wildcard character in the search bar, or enter text with a matching subject to find a specific email, and then select the emails to recover from the search results.
   Or click **Advanced Search** and use the following filters to search for specific emails:

Filters	Description
Subject of Event	Subject line in the calendar invite.
Organizer	The email address of the event organizer.
Invitee	Event recipients' email addresses. Use a comma or space separator to enter multiple addresses.
Invitation Date	Using the calendar, select a specific date range to search the calendar invites.

6. Select the items that are to be restored.
7. To use a different Recovery Point, click the Edit icon on the tile. Find the recovery point you need and click **Select Recovery Point**.
8. Click **Next: Recover Options**.
9. From the **Recover To** section, select one of the following options based on your requirements:

Recover To options	Description
Original Location	To recover data to the same mailbox from which it was backed up.
New Location	To recover data to a different Microsoft 365 mailbox. Do the following: • Select the required **Registered Source** from the list. Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source • Select the **Target Mailbox** and specify the **Folder name** to which you want to recover the data. If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.
Download as PST	To download the required email as a PST file. To enhance security, you can set a password for the downloaded PST file. Once the recovery task is completed, the exported PST file is valid for 7 days. Ensure that you download the PST within 7 days. For more information, seeRecover Mailbox .
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

10. Select your **Recovery Options**:
    - **Continue on Error**. Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesitywill still attempt to recover the other selected emails.
    - **Task Name**. Change the default name of the recovery task.
11. Click **Recover**.
    The system creates a separate recovery task as needed. A sequential suffix number is added to each recovery task name. You can monitor and track these recovery tasks individually on the **Recoveries** page.

Recover Contacts

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To recover specific contacts:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Mailboxes**.
3. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, select **Search Mailbox Items**.
4. From the **Item Type** drop-down list, select **Contacts**.
5. To search for emails, enter the * wildcard character in the search bar, or enter text with a matching subject to find a specific email, and then select the emails to recover from the search results.
   Or click **Advanced Search** and use the following filters to search for specific emails:

Filters	Description
First Name	The first name of the contact.
Last Name	The last name of the contact.
Email Address	The email address of the contact.

6. Select the items that are to be restored.
7. To use a different Recovery Point, click the Edit icon on the tile. Find the recovery point you need and click **Select Recovery Point**.
8. Click **Next: Recover Option**s.
9. From the **Recover To** section, select one of the following options based on your requirements:

Recover To options	Description
Original Location	To recover data to the same mailbox from which it was backed up.

Recover To options	Description
New Location	To recover data to a different Microsoft 365 mailbox. Do the following: • Select the required **Registered Source** from the list. Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source • Select the **Target Mailbox** and specify the **Folder name** to which you want to recover the data. If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.
Download as PST	To download the required email as a PST file. To enhance security, you can set a password for the downloaded PST file. Once the recovery task is completed, the exported PST file is valid for 7 days. Ensure that you download the PST within 7 days. For more information, seeRecover Mailbox .
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

10. Select your **Recovery Options**:
    - **Continue on Error**. Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesity will still attempt to recover the other selected emails.
    - **Task Name**. Change the default name of the recovery task.
11. Click **Recover**.
    The system creates a separate recovery task as needed. A sequential suffix number is added to each recovery task name. You can monitor and track these recovery tasks individually on the **Recoveries** page.

Recover Notes

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To recover specific notes:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Mailboxes**.
3. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, select **Search Mailbox Items**.
4. From the **Item Type** drop-down list, select **Notes**.
5. To search for emails, enter the * wildcard character in the search bar, or enter text with a matching subject to find a specific email, and then select the emails to recover from the search results.
   Or click **Advanced Search** and use the following filters to search for specific emails:

Filters	Description
Subject	The subject of the note.
Creation Date	Using the calendar, select a specific date range to search the notes based on the creation date.
Modification Date	Using the calendar, select a specific date range to search the notes based on the modification date.

6. Select the items that are to be restored.
7. To use a different Recovery Point, click the Edit icon on the tile. Find the recovery point you need and click **Select Recovery Point**.
8. Click **Next: Recover Options**.
9. From the **Recover To** section, select one of the following options based on your requirements:

Recover To options	Description
Original Location	To recover data to the same mailbox from which it was backed up.

Recover To options	Description
New Location	To recover data to a different Microsoft 365 mailbox. Do the following: • Select the required **Registered Source** from the list. Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source • Select the **Target Mailbox** and specify the **Folder name** to which you want to recover the data. If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.
Download as PST	To download the required email as a PST file. To enhance security, you can set a password for the downloaded PST file. Once the recovery task is completed, the exported PST file is valid for 7 days. Ensure that you download the PST within 7 days. For more information, seeRecover Mailbox .
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

10. Select your **Recovery Options**:
    - **Continue on Error**. Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesity will still attempt to recover the other selected emails.
    - **Task Name**. Change the default name of the recovery task.
11. Click **Recover**.
    The system creates a separate recovery task as needed. A sequential suffix number is added to each recovery task name. You can monitor and track these recovery tasks individually on the **Recoveries** page.

Recover Tasks

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To recover specific tasks:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Mailboxes**.
3. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, select **Search Mailbox Items**.
4. From the **Item Type** drop-down list, select **Tasks**.
5. To search for emails, enter the * wildcard character in the search bar, or enter text with a matching subject to find a specific email, and then select the emails to recover from the search results.
   Or click **Advanced Search** and use the following filters to search for specific emails:

Filters	Description
Subject	The subject of the task.
Creation Date	Using the calendar, select a specific date range to search the tasks based on their creation date.
Due Date	Using the calendar, select a specific date range to search the tasks based on their due date.
Status	The status of the task.

6. Select the items that are to be restored.
7. To use a different Recovery Point, click the Edit icon on the tile. Find the recovery point you need and click **Select Recovery Point**.
8. Click **Next: Recover Options**.
9. From the **Recover To** section, select one of the following options based on your requirements:

Recover To options	Description
Original Location	To recover data to the same mailbox from which it was backed up.

Recover To options	Description
New Location	To recover data to a different Microsoft 365 mailbox. Do the following: • Select the required **Registered Source** from the list. Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source • Select the **Target Mailbox** and specify the **Folder name** to which you want to recover the data. If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.
Download as PST	To download the required email as a PST file. To enhance security, you can set a password for the downloaded PST file. Once the recovery task is completed, the exported PST file is valid for 7 days. Ensure that you download the PST within 7 days. For more information, seeRecover Mailbox .
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

10. Select your **Recovery Options**:
    - **Continue on Error**. Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesity will still attempt to recover the other selected emails.
    - **Task Name**. Change the default name of the recovery task.
11. Click **Recover**.
    The system creates a separate recovery task as needed. A sequential suffix number is added to each recovery task name. You can monitor and track these recovery tasks individually on the **Recoveries** page.

Recover Items in Recoverable and Archive Recoverable Items Folder

Cloud Protection Service now supports granular recovery of items in Recoverable Items and Archive Recoverable Items folders. Additionally, all four Microsoft 365 mailbox folder roots, Primary Mailbox, In-Place Archive, Recoverable Items, and Archive Recoverable Items are

now displayed in the granular recovery workflow, allowing you to select any root individually to browse and recover its items.

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

To recover specific items from the Recoverable Items folder:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Mailboxes**.
3. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, click **Mailbox Contents**.
4. From the **Mailbox Items** drop-down list, select **Emails**.
5. Select the root type filter (default: **Primary Mailbox**), and select one or more roots to search for items.
   At least one root must be selected.

To search for emails, enter the * wildcard character in the search bar, or enter text with a matching subject to find a specific email, and then select the emails to recover from the search results.
Or click **Advanced Search** and use the following filters to search for specific emails:

Filters	Description
Subject	Subject line in the email.
From	Mail sender email address.
To	Mail recipient email address. Use a comma or space separator to enter multiple addresses.
Date Range	Using the calendar, select a specific date range to search the emails.
Email Type	Select one of the email types: • All Emails • Only emails with attachments • Only emails without attachments
cc	The email address in the Cc: line of the email. Use a comma or space separator to enter multiple addresses.

Filters	Description
bcc	The email address in the Bcc: line of the email. Use a comma or space separator to enter multiple addresses.
Search in Folder	Search for the email within the specified folder. For example, Inbox, Drafts, and so on. Use a comma or space separator to enter multiple folder names.

1. Select the items that are to be restored.
2. To use a different recovery point, click the **Edit** icon on the tile. Find the recovery point you need and click **Select Recovery Point**.
3. Click **Next: Recover Options**.
4. From the **Recover To** section, select one of the following options based on your requirements:

Recover To options	Description
Original Location	To recover data to the same mailbox from which it was backed up.
New Location	To recover data to a different Microsoft 365 mailbox. Do the following: • Select the required **Registered Source** from the list. Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source • Select the **Target Mailbox** and specify the **Folder name** to which you want to recover the data. If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.
Download as PST	To download the required email as a PST file. To enhance security, you can set a password for the downloaded PST file. Once the recovery task is completed, the exported PST file is valid for 7 days. Ensure that you download the PST within 7 days. For more information, seeRecover Mailbox .

Recover To options	Description
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

5. Select your **Recovery Options**:
   - **Continue on Error**. Enable to recover even if errors occur when recovering Mailboxes. For example, if one of the emails cannot be recovered, Cohesity will still attempt to recover the other selected emails.
   - **Task Name**. Change the default name of the recovery task.
6. Click **Recover**.
   The system creates a separate recovery task as needed. A sequential suffix number is added to each recovery task name. You can monitor and track these recovery tasks individually on the **Recoveries** page.

Recover Emails and Folders

To recover emails and folders, you need to search through the backups for the email or folders using the following search filters:

Search Filters	Description
From	Mail sender email address
To	Mail recipient email address
Subject	Subject line for the email
Folder Name	Folder name which needs to be recovered
Start Date	Time interval
End Date	

Search Filters	Description
Only show Emails with Attachments	Search only the emails with attachments
Only show Folders	Search only the folders
Cc	The user email address that was added as a cc in the email
Bcc	The user email address that was added as bcc in the email
Microsoft 365 Server	From the drop-down list, select the Microsoft 365 server.
Mailbox	From the drop-down list, select the mailboxes from which you want to recover the emails or folders.
Protection Group	From the drop-down list, select the Protection Group.

Once the search yields the desired result, select the mail or folder you plan to recover and start the recover operation.

To recover emails and folders:

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

1. Navigate to **Protection** > **Recoveries**. Click the plus icon (**+**) on the top-right of the page and then select **Microsoft 365** as the recovery source.
2. On the **New Recovery Microsoft 365 - Mailbox** page, in the **Recovery Type**, select **Emails & Email Folders**.
3. Search the emails or folders using the search filter, click **Search**.
4. Select the mailbox from the results, and click **Continue**.
5. Based on your requirements, select any one of the following options:

Options	Description
Original Location	To recover data to the same mailbox from which it was backed up.

Options	Description
New Location	To recover data to an alternate location in a different Microsoft 365 domain. Select a **Registered Source** and the **Target Mailbox**, and specify the **Folder name** to which you plan to recover. **Note:** • Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more details, see Register and Manage Microsoft 365 Source. • If a folder with the specified name does not exist, Cohesity creates the folder and recovers the data to it.
Export as PST	To export the backed-up email(s) or mailbox folder(s) as a PST file, you must provide a password for the exported PST file. **Note:** • Once the recovery task is completed, the exported PST file is valid for 72 hours. Ensure that you download the PST within 72 hours. For more information, see Download the Exported Email or Mailbox Folder PST File. • The PST export process might take a few minutes or hours depending on the size of the mailbox.
On-prem Exchange Server	To recover data to an on-premises Exchange Server created via API, select the required on-premises Exchange Server from the drop-down list. Optionally, enter an alternate SMTP address for recovering mailbox content. If no alternate SMTP address is specified, the source user's SMTP address is used by default. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

6. Select your Recovery Options:
   i. **Continue on Error.** Enable to complete the recovery task even if errors occur when recovering Mailboxes.
   ii. **Task Name.** Change the default name of the recovery task.
   iii. **Selective Recovery.** Toggle ON the **Recover within specified timeframe** option to filter the Mailbox items to be recovered.

   Under **Include items in recovery**, select the required items:

- Mails
- Calendars
- Contacts, Tasks, and Notes.

  You can also **recover data of the last specific days prior to the selected snapshot**.

> **Note:**
> - All upcoming calendar events will be recovered.
> - Any calendar events that were scheduled or received in the selected time frame will be recovered.
> - Contacts, Tasks, and Notes will be fully recovered unaffected by the days selected.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

7. Click **Start Recovery** to start the recover operation.

Download the Exported Mailbox PST File

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

After the recovery task is completed, within 72 hours you can download the exported PST file of the mailbox that you choose to recover.

To download the PST file:

1. Navigate to **Protection** > **Recoveries**.
2. Locate and click on the recovery task from which you want to download the exported PST file.
3. Click **Download All**.

The PST file is downloaded to your local system.

> **Note:** If the PST file is protected with a password, you must contact the admin user and obtain the password to open the downloaded PST file.

To download PST files for individual mailboxes:

1. Hover over the required object.
2. Click the download icon .

The PST file for the mailbox will be downloaded.

To view the logs for the mailbox,

1. Hover over the required object.
2. Click the log icon .

> **Note:**
> - Ensure you use a download manager when downloading a large number of mailboxes.
> - Set a default download folder in your browser settings to prevent being prompted for a location each time a ZIP file is downloaded.
> - The ZIP file name includes the mailbox name and recovery task ID by default. For example, Adele Vance_2025-09-24T125237Z_PSTs.zip

Download the Exported Email or Mailbox Folder PST File

Before you proceed with recovery, you must review the considerations. For more information, see Plan and Prepare for Mailbox Protection.

After the recovery task is completed, within 72 hours you can download the exported PST file of the email(s) that you choose to recover.

To download the PST file:

1. Navigate to **Data Protection > Recoveries**.
2. Locate and click on the recovery task from which you want to download the exported PST file.
3. Click **Download Files**.

The PST file is downloaded to your local system.

Download Private Chats

Before you proceed with recovery, you must review the considerations. For more information, see Backup Exchange Online Mailboxes.

To download the Private Chats of a user:

1. Navigate to **Sources**.
2. Click on the required **Source** and click the **Mailbox** tab.

3. Select **Show All** > **Protected** and use the other filters, search box, and views at the top to narrow your search.
4. Click the **Actions** menu next to the object and click **Download Private Chats**.
5. In the **Download Private Chats** page, select the required snapshot, provide the task name, and click **Recover**.

   > **Note:** Attachments in the Private Chats will not be downloaded.

6. Click the View Progress button in the pop-up message or click the **Activity** menu.
7. Once the recovery is successful, click Download. The Private Chats will be downloaded.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Archive Mailbox Data

Cohesity enables you to archive the Exchange Online Mailboxes data to a designated external target, such as tape or cloud storage. For more details about external targets, see Manage External Targets.

It provides a policy-based archival solution for Exchange Online Mailboxes data. In the Protection Policy, you can define the archive schedule, select the external target, and specify the retention period of the data on the external target.

> **Note:** CloudRetrieve is not supported.

To archive Exchange Online Mailboxes data:

1. Register an external target. For more details, see Register an External Target.
2. Edit the existing Protection Policy that you had created for backup or if you are creating a new Protection Policy for backup. Configure the Archival settings as per your requirement. For more information, see **Step 18** in Create or Edit a Standard Policy.
3. Run the Protection Group that you have created for Exchange Online Mailboxes based on your requirement or the Archival settings will take effect in the next scheduled Protection Group run.

Exchange Online Public Folders

Cohesity enables data protection of Microsoft Exchange Online. Cohesity provides a policy-based data protection solution for protecting Microsoft Exchange Online Public Folders on Microsoft 365.

> **Note:** Starting October 2026, Cohesity will discontinue support for Public Folder protection, following Microsoft's retirement of APIs for managing Public Folders.

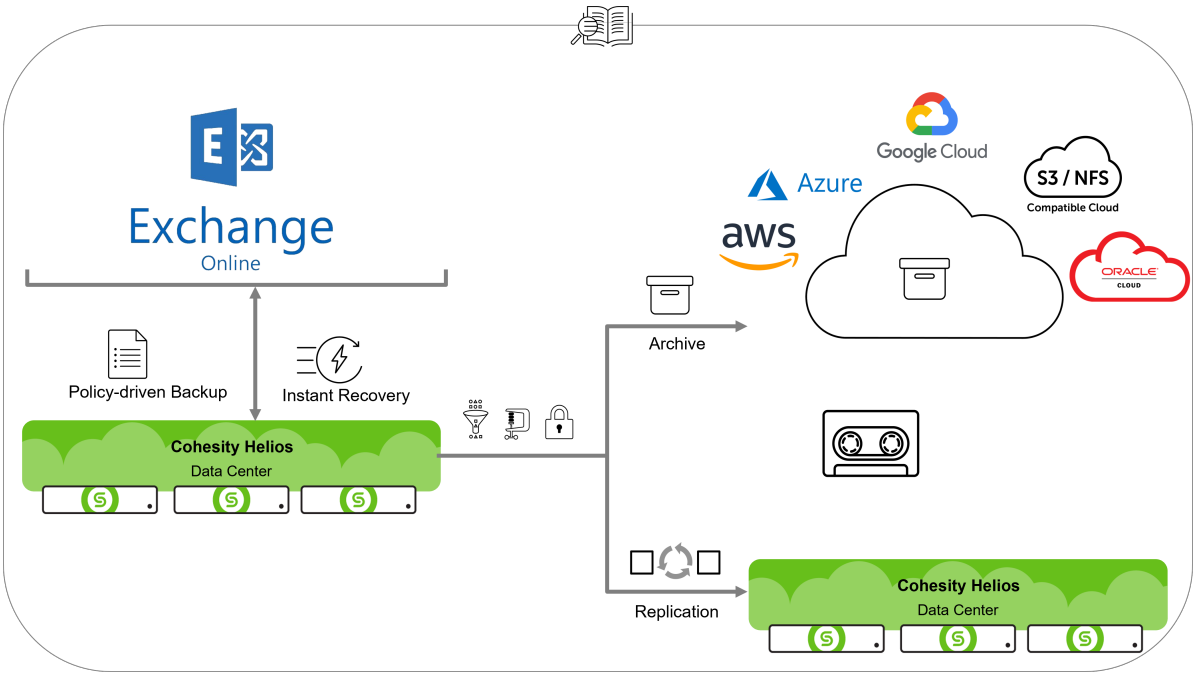

Plan and Prepare for Public Folders Protection

Before you register your Microsoft 365 domain with Cohesity to protect Microsoft 365 data, ensure:

- You have completed all the Prerequisites.
- You have reviewed and understood the Considerations and Best Practices.
- An Exchange Online license with a mailbox provisioned.
- A default Public folder mailbox configured (preferably Primary).

Considerations

Review and understand the following considerations before you protect your Exchange Online public folder data:

- For public folder backups, if you move a subfolder from one root folder to another root folder, a full backup of the subfolder will be performed in the next run.
- Exclusion list is not supported for Public Folders Protection Group.
- Suppose a newly added folder at any location in the folder hierarchy does not have the required permissions mentioned in required Public Folder Root Permissions for Microsoft 365 User Account. The folder, including its subfolders, is not discoverable and hence will not be backed up by Cohesity.
- Microsoft does not support OAuth for Public Folders and backups will fail if you have enabled OAuth for Exchange Online backup.
- Restoring the public folder or content of the public folder from the backup to an alternate location in a different Microsoft 365 domain is not supported.
- Public Folders are not supported in Government Community Cloud (GCC) High.
- Microsoft 365 items with only special characters will not be listed in granular search. The item name must contain at least one alphanumeric character to be listed in the search result.
- Backup and recovery of files with sensitivity labels is supported.

  For files with sensitivity labels,

  - Only a user with access via the label can open the file after downloading, provided the label is not permanently deleted in SharePoint.
  - When recovering to original location, if the file is already present in the destination, the recovery fails. If the file is not present in the destination, a new file will be created.
  - The files downloaded from SharePoint as encrypted. You can download the same file from the backup storage.

Best Practices

To protect your Microsoft 365 data, Cohesity recommends the following best practices:

- Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling failed requests. There is no need to increase the value for this setting in the Cohesity Protection Policy.
- Use **Auto-Protect** to avoid adding new users to the Protection Group manually.
- Avoid adding the same users in multiple Protection Groups.

Register and Manage Microsoft 365 Exchange Online Public Folders

To start protecting Microsoft 365 applications, you need to register the Microsoft 365 domain as a source in Cohesity.

Cohesity uses the Microsoft Graph API for object discovery, backup, and recovery in Microsoft 365. Cohesity uses an Azure application created and registered on the Azure

portal with the necessary permissions to use the Graph API. You can either let Cohesity create the Azure application or you can manually enter Azure application details while registering your Microsoft 365 domain as a source in Cohesity.

Express Registration for Microsoft 365 Sources

Before you register Office 365 as a source on Cohesity, ensure that you:

- Review the requirements, considerations, and best practices.
- Perform the prerequisite tasks.

To register your Microsoft 365 domain:

1. Navigate to **Sources** and select **Register** > **Microsoft 365**.
2. Choose **Public Folders** as the **Microsoft 365 Applications** to discover.
3. In the **Account Credentials** section, enter the Microsoft 365 **Username** and **Password**.

   > **Note:** Toggle ***on*** the **Enable OAuth** option. Ensure that you added the required permission.

4. Adding multiple Microsoft 365 service accounts using the **Additional Service Accounts** section requires OAuth disabled, as Microsoft has enabled OAuth by default, ignore adding multiple Microsoft 365 service accounts.
5. In the **Azure Applications** section, enter the number of Azure applications you want to create based on your requirements and click **Create**.

   > **Note:** By default, two Azure applications will be created. To better manage Microsoft 365 throttling, Cohesity recommends at least two Azure apps.

6. In the **Add Azure Application** form, copy the device code and click the **Microsoft Azure App** link to open the Microsoft Azure App authorization service in a new tab.

   > **Note:** If you prefer to create your Azure apps manually, see Manual Registration for Microsoft 365 Sources topic.

7. Paste the copied code in the **Microsoft Azure App authorization** service and click **Next**.
8. Log in to Microsoft Azure, enter the **Username** and **Password** of your Microsoft 365 account and click **Sign in**.

> **Note:** Ensure that your Microsoft 365 account has a global administrator access.

9. Follow the instructions to complete the authorization on the Microsoft Azure portal. Then, wait for Microsoft Azure Authorization to complete.
10. Optional. Toggle on **Use Windows Connector** and specify the hostname (FQDN) or server IP of the Windows Connector. The **Use Windows Connector** option is applicable only if you want to protect SharePoint Online data or to export and download mailboxes or email(s) in PST format. For more information, see Set Up a Windows Connector for Microsoft 365.
11. Click **Register**.

Manual Registration for Microsoft 365 Sources

Before you register Office 365 as a source on Cohesity, ensure that you:

- Review the requirements, considerations, and best practices.
- Perform the prerequisite tasks.

To register your Microsoft 365 domain:

1. Navigate to **Sources** and select **Register** > **Microsoft 365**.
2. Choose **Public Folders** as the **Microsoft 365 Applications** to discover.
3. In the **Account Credentials** section, enter the Microsoft 365 **Username** and **Password**.

> **Note:** Toggle ***on*** the **Enable OAuth** option. Ensure that you added the required permission.

4. Adding multiple Microsoft 365 service accounts using the **Additional Service Accounts** section requires OAuth disabled, as Microsoft has enabled OAuth by default, ignore adding multiple Microsoft 365 service accounts.
5. In the **Azure Applications** section, enter the number of Azure applications you want to create based on your requirements and click **Create**.

> **Note:** By default, two Azure applications will be created. To better manage Microsoft 365 throttling, Cohesity recommends at least two Azure apps.

6. In the **Add Azure Application** form, click the **You can also add Azure App manually** link.

- For **Certificate** authentication, the **App ID** and the **Associated Certificate** of your custom Azure app will be populated. You can select the **Associated Certificate** from the list of certificates present in your certificate store or Microsoft 365 Certificate-Based Authentication.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- For **App Secret** authentication, the **App ID** and **App Secret Key** of your custom Azure app will be populated.

> **Note:** You can add multiple Azure apps for a Microsoft 365 source to load balance your backup and restore operations. Click **+** to add multiple Azure apps. When you do, ensure that you provide the valid **App ID** and **App Secret Key**.

To migrate the existing apps from App Secret to Certificate authentication method, see Microsoft 365 Certificate-Based Authentication.

7. Optional. Toggle on **Use Windows Connector** and specify the hostname (FQDN) or server IP of the Windows Connector. The **Use Windows Connector** option is applicable only if you want to protect SharePoint Online data or to export and download mailboxes or email(s) in PST format. For more information, see Set Up a Windows Connector for Microsoft 365.
8. Click **Register**.

Manage Microsoft 365 Source Registration

After registering your Microsoft 365 domain as a source, you can:

- Update the Microsoft 365 source configuration.
- Refresh the source details.
- Unregister the Microsoft 365 domain from Cohesity.

Update the Microsoft 365 Source Configuration

After registering your Microsoft 365 source on Cohesity, you might have changed the Microsoft 365 domain configuration by:

- Changing the credentials of Microsoft 365 user account credentials.
- Updating the app secret by adding more permissions to the custom app.
- Creating a new app ID.

You can update the Microsoft 365 details provided during the registration process with the latest Microsoft 365 configuration details.

To edit the Microsoft 365 source configuration:.

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Edit**.
3. On the **Update Microsoft 365 Source** page, update the required configurations.
4. Click **Register**.

Refresh the Microsoft 365 Source

You can refresh the Microsoft 365 domain configuration and fetch the latest changes on the Microsoft 365 domain.

To refresh the Microsoft 365 Source Configuration:

1. Navigate to **Data Protection** > **Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Refresh**.

Unregister the Microsoft 365 Domain

If you plan to stop taking the backup of your Microsoft 365 domain, you can unregister the Microsoft 365 source from Cohesity.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the Microsoft 365 source.

To unregister the Microsoft 365 domain:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Microsoft 365 source and select **Unregister**.

Backup Exchange Online Public Folders

Exchange Online public folders are protected using a Protection Group and a Protection Policy. A Protection Group enables you to select and group the objects that you want to back up. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule, and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. You can select which standard or custom policy to use when creating a Protection Group. For more information, see Manage Protection Policy.

To protect Exchange Online public folders:

1. Do one of the following:
   - On the **Dashboards** page, from the drop-down at the top of the page, select **Data Protection** and do one of the following:
     - If there is no existing Protection Group, click **Start Data Protection** > **Microsoft 365** > **Public Folder**.
     - Click **Protect** > **Microsoft 365** > **Public Folder**.
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Microsoft 365** > **Public Folder**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Microsoft 365 source and then select **Protect Public Folders**.
2. In the **New Protection** pop-up, do one of the following:
   - Protect the public folder by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 account. In the **Public Folders** section, select the public folder(s) from the Root Public Folder List that you plan to back up.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. Click **Protect**.
   - To protect the public folder by specifying all the options, click **More Options**.
3. From the **Source** drop-down list, select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 source.
4. In the **Public Folders** section, under **Objects**, select the public folder(s) from the **Root Public Folder List** that you plan to backup and click **Save Selection**.

> **Note:** You can use the advanced filters that are available to find the public folders or search for the public folder using the folder name. Also, you can click the shield icon to enable Auto Protect. For more information on Auto Protect, see Enable or Disable Auto Protect.

5. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name of the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, Create or Use an Existing Protection Policy.
7. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
8. After selecting a Storage Domain, if you need to change any of the advanced settings on the **New Office365 Protection Group** page, scroll down and click **Edit** on the right. For more information, see Advanced Settings.
9. Click **Protect**.

   Your Exchange Online Data Protection is now active and running.

Advanced Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection run should start. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Exclusion List	Select the folders that you plan to exclude from the backup or click **Add** to add custom folders that you want to exclude from the backup. **Note:** Exclusion list is not supported for Public Folders.
Indexing	By default, indexing is enabled. **Note:** Indexing is mandatory for granular restore of an Exchange Online mailbox , such as restoring a folder or restoring an email. For more information, see Recover Mailbox .
Alerts and Priority	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.

Advance Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	**Optional**. Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

Create a Protection Policy with the settings listed in the Policy Settings for Exchange Online. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Exchange Online

The following table lists the applicable policy settings for Exchange Online:

Field	Description
DataLock	Add a DataLock for compliance and regulatory requirement. This ensures that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users from modifying or deleting any Snapshots that were generated by the protection runs that use this Policy. The users include those who have the Data Security role in Cohesity as well. > **Note:** Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained, and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer duration than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the protection run fails. The default time between retries is thirty minutes. You can customize the number of retries and how long to wait between each attempt. > **Note:** Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling the failed requests. There is no need to increase the value for this setting in the Protection Policy.
Quiet Times	Quiet Times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	Replicate snapshots to another Cohesity cluster.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with applicable policy settings for Exchange Online. In this case, all the policy settings apart from the settings applicable to Exchange Online are ignored.

To use an existing Protection Policy:

1. Navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for Exchange Online. For information on the applicable policy settings, see Policy Settings for Exchange Online.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage Public Folder Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered Microsoft 365 domain and the assigned Storage Domain in the Protection Group settings.

Add Public Folders to the Protection Group

You can add new public folders to the existing Protection Group for backup. The newly added public folder is backed up from the next protection run and it will be a full backup for the newly added public folder.

To add public folders to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Public Folders** section.
5. In the **Objects** section, select the public folders that you want to add to the backup and click **Save Selection**.
6. Click **Save**.

Remove Public Folders From the Protection Group

You can remove public folder(s) from the Protection Group that you do not want to backup in the next protection run. When you remove a public folder from the protection group, incremental backup of the contents in the public folder is not performed from the next protection run. The existing backup for the public folder is preserved until the snapshot expires.

To remove public folders from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Public Folders** section.
5. In the **Objects** section, select the public folders that you want to remove from the backup and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new public folder(s) that are added to Microsoft 365.

When the Auto Protect option is enabled:

- All the public folders that are added to Microsoft 365 in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the public folders that are removed from Microsoft 365 in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backups for the public folder are preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly onboarded public folder on Microsoft 365 to the Protection Group.
- Remove existing public folders that are offboarded on Microsoft 365 from the Protection Group. The existing backups of the public folders are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, click ✎ in the **Public Folders** section.
5. In the **Objects** section, click the icon to enable or disable Auto Protect and click **Save Selection**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Protect**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up Exchange Online. It provides you with options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.
- Add replication.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. In the **Backup** drop-down list, select any one of the following:
        - Backup all Objects in the Protection Group
        - Backup only selected Objects in the Protection Group
    2. In the **Backup Type** drop-down list, select the backup type.
    3. Optional. In the **Replicate** section, select the replication target and the retention duration or click **Add** to add the replication target and the retention duration.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups, local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details will not be displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Exchange Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - **Delete Protection Group Only**
   - **Delete Protection Group and Snapshots**

Recover Public Folders

You can restore Exchange Online public folders data using Cohesity at various granularities, such as at the public folder and folder contents level. You can restore to the same location or to an alternate location.

Restore Public Folders

To restore public folders, you can search for the public folder(s) through the backups by the public folder name or substring in the name of the public folder.

> **Note:** You can use the * character to list all the public folders.

To restore public folder(s):

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Public Folder** as the recovery source.
2. In the **Recovery Type** section, select **Entire Public Folders**.
3. Search the public folder(s) using the partial name of the root public folder or use * character to list all the public folders. Also, you use the available advanced filters to narrow down on the public folder(s) that you plan to recover. Select the public folder (s) you need, and click **Next:Recover Options**.
4. In the **Recover To** section, based on your requirements, select any one of the following options:

Options	Description
Original Location	To restore the public folder to the same location from which it was backed up.

Options	Description
New Location	**Alternate location in the same Microsoft 365 domain**: Choose the **Root Public Folder** to which data needs to be restored and also define the **Path** in the public folder where you plan to restore the data.

5. In the **Recovery Options** section, perform the following:
    1. Toggle on **Continue on Error** to continue the recovery process even if errors occur while recovering the public folder(s).
    2. In the **Task Name** field, enter a name for the recovery task.
6. Click **Start Recovery** to start the restore operation.

Restore Public Folder Contents

To restore the contents of a public folder, you need to search for the contents of the public folder through the backups using the following item types and search filters:

Item Type	Search Criteria
Event	Subject
Contacts	Display Name
Emails & Posts	Subject
Subfolders	Folder Name
Tasks	Subject
Journals	Subject
Notes	Subject

Once the search yields the desired result, select the mail or folder you plan to restore and start the restore operation.

To restore emails and folders:

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Public Folder** as the recovery source.
2. In the **Recovery Type** section, select **Folder Contents**.
3. Search for the public folder contents using **Item Type** and search filters. Also, you use the available advanced filters to narrow down on the public folder content from

the search results. Select the public folder content that you need, and click **Next:Recover Options**.

> **Note:** You cannot select the items across Protection Group, Storage Domain and Office365 domain.

4. In the **Recover To** section, based on your requirements, select any one of the following options:

Options	Description
Original Location	To restore the public folder content to the same location from which it was backed up.
New Location	**Alternate location in the same Microsoft 365 domain**: Choose the **Root Public Folder** to which data needs to be restored and also define the **Path** in the public folder where you plan to restore the public folder content.

5. In the **Recovery Options** section, perform the following:
    1. Toggle on **Continue on Error** to continue the recovery process even if errors occur while recovering the public folder(s).
    2. In the **Task Name** field, enter a name for the recovery task.
6. Click **Start Recovery** to start the restore operation.

OneDrive for Business

OneDrive for Business is a SaaS application that is bundled in your Microsoft 365 subscription service. It is an intelligent files app for Microsoft 365 connecting you to all your files so you can share and work together from anywhere while protecting your work. It enables you to easily store, access, and discover your individual and shared work files in Microsoft 365. Using the policy-based data protection solution from Cohesity, you can protect OneDrive for Business data on Microsoft 365.

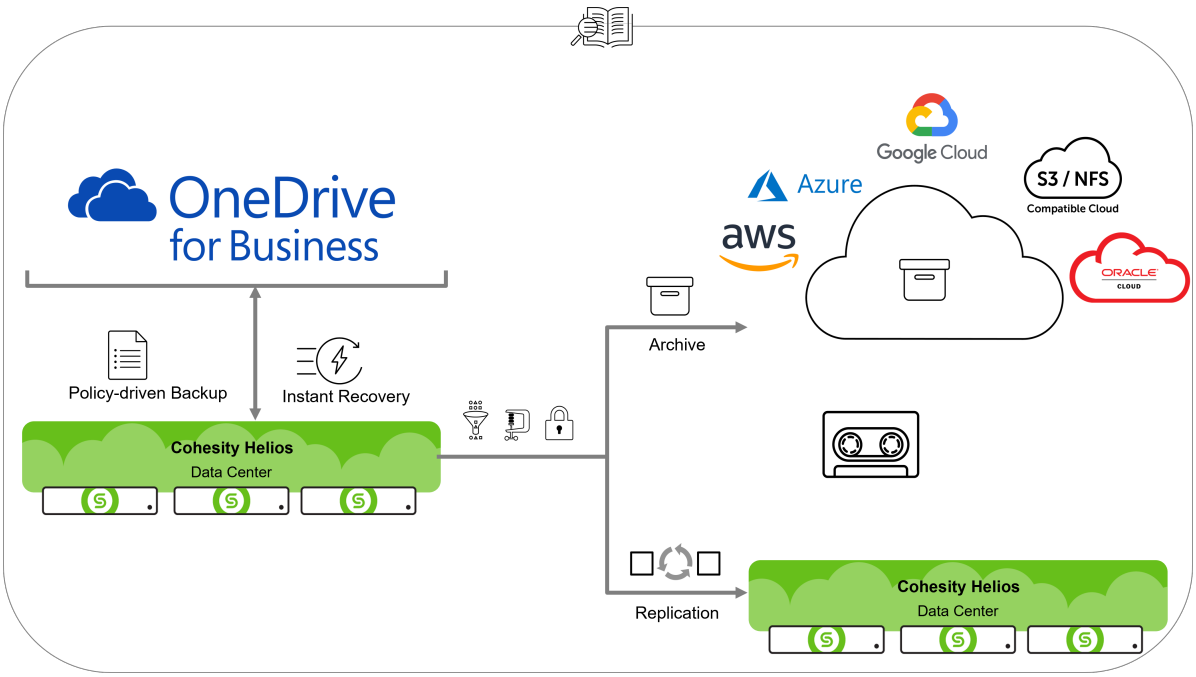

Plan and Prepare for OneDrive Protection

Before you register your Microsoft 365 domain with Cohesity to protect Microsoft 365 data, ensure:

- You have completed all the Prerequisites.
- You have reviewed and understood the Considerations and Best Practices.

Considerations

- SharePoint does not support file or folder names that contain unsupported special characters. Backup, restore, or download of such items is not supported. See, Restrictions and limitations in OneDrive and SharePoint.
- Backup and recovery of OneNote files in OneDrive is supported. Consider the following:
  - Certificate-based authentication must be enabled. For more information, see Microsoft 365 Certificate-Based Authentication.
  - The first OneNote backup after enabling this feature may take longer than usual incremental backups, as it will include all OneNote files that were previously skipped.
- Backup and recovery of the Recycle Bin in OneDrive is not supported.
- Backup of the mailbox and the OneDrive for an individual user happens sequentially. However, backup of mailboxes and OneDrive belonging to different users may happen in parallel.

- The backup or recovery of Mailboxes, OneDrives, SharePoint Online Sites, and MS Teams are executed in parallel as long as there are enough resources on the Cohesity cluster. However, the backup or recovery operations are serialized for the Mailboxes and OneDrives that belong to the same user, and the operations execute in the first-in-first-out order.

  For example, if the backup of a user's mailbox is in progress, then the backup of the user's OneDrive starts only after the backup of the mailbox completes.
- The user list containing users with protected mailboxes shows as protected when creating a Protection Group for OneDrive and vice-versa.
- OneDrive protection is supported only for users present on Azure AD or Hybrid setup with users present on Azure AD and not for users on the on-premises Active Directory.
- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.
- OneDrive does not support item permissions from applications and devices.
- Microsoft 365 items with only special characters will not be listed in granular search. The item name must contain at least one alphanumeric character to be listed in the search result.
- Permissions are recovered only at file level and not folder level.
- If more than 50 thousand files are present, the permissions are recovered only for 50 thousand files (there is no option to identify these files in advance).
- Sharing links permissions are not recovered.
- Cohesity supports protecting 100k Microsoft 365 objects.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  This includes the following:

  - Objects on individual applications, for example, 100k Teams.
  - Objects on combined applications, for example, 50k Teams, 30k Mailboxes, and 20k Groups.

  Validation messages are displayed when backing up more than 7500 objects and recovering more than 5000 objects.
- Backup and recovery of Preservation Hold Library (PHL) is supported. Consider the following points:

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

> **Note:** PHL is a hidden, system-managed library that OneDrive and SharePoint automatically create when a retention policy, retention label, or legal/eDiscovery hold is applied. It stores preserved copies of files that are edited or deleted while under retention.

  - **Shared permissions are not restored**: Recovering PHL data recovers the file content only. Any sharing permissions it had must be reapplied manually after recovery.
  - **No granular recovery**: PHL data can only be recovered as part of a full OneDrive recovery; individual PHL files or folders cannot be recovered.
  - **PHL data is not searchable**: PHL data is not indexed by standard Microsoft Search and cannot be browsed or searched directly.
  - **Recovery time and space impact**: Recovering PHL data may increase recovery time substantially. Recovery will fail if sufficient space is not available at the target location.
- Backup and recovery of files with sensitivity labels is supported. For files with sensitivity labels:
  - Only a user with access via the label can open the file after downloading, provided the label is not permanently deleted in SharePoint.
  - When recovering to original location, if the file is already present in the destination, the recovery fails. If the file is not present in the destination, a new file will be created.
  - The files downloaded from SharePoint as encrypted. You can download the same file from the backup storage.
- The following metadata fields are restored when SharePoint data is restored:
  - Created By (author)
  - Modified By (editor)
  - Created Date
  - Modified Date

    During cross-domain restore operations, the Created By (author) and Modified By (editor) fields are not restored because the corresponding users from the source domain will be unavailable in the target domain.
    The Created By metadata is not restored if the user referenced in the Created

By field is deleted. The Modified By and Created By metadata are not restored if the user referenced in the Modified By field is deleted.

- Backup and recovery of Exchange Online mailboxes using security groups is supported. You must review the following considerations before protecting mailboxes using a security group:

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  - Each protection group can contain only one security group. Mixing multiple security groups within the same protection group is not supported.
  - The member count for a security group includes only users present in the Entity Hierarchy, regardless of whether they have a valid mailbox or OneDrive license. If a user's license is removed, they remain in the count until the next source refresh filters them out. Members without a valid license will fail individually during backup, but the protection job proceeds for the remaining members.
  - Security group membership is fetched from Microsoft Entra ID. Newly added members are included in the next scheduled refresh; members removed from the group or deleted from Entra ID are automatically excluded after the next refresh. If security group optimization is enabled, members are not fetched during the entity hierarchy refresh and are instead fetched during the backup run.
  - Security group protection is supported only for the members of the security group, not the owners of the group.
  - Both mail-enabled and mail-disabled security groups are supported.
  - Members of nested security groups are automatically included in protection. The member count displayed reflects the unique (deduplicated) count across all levels of nesting.
  - If the selected security group has no members, the protection job fails. Jobs resume normally once members are added to the group and the next entity hierarchy refresh completes.
  - Remove the protection group manually if the security group is deleted from Entra ID.
  - If a user is a member of two security groups that are each assigned to a separate protection group, that user's mailbox or OneDrive is backed up by both protection jobs.
  - The security group itself is manually selected for protection; individual members within the group are treated as auto-protected.

- Recovery of mailboxes or OneDrives protected via a security group follows the same workflow as recovery of individually protected objects. There is no change to recovery workflows.
- Backup and recovery of external users from a security group is supported.

Best Practices

To protect your OneDrive for Business data, Cohesity recommends the following best practices:

- Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling failed requests. There is no need to increase the value for this setting in the Cohesity Protection Policy.
- Use **Auto Protect** to avoid adding new users to the Protection Group manually.
- Avoid adding the same users in multiple Protection Groups.

Backup OneDrive for Business

OneDrive for Business is protected using a Protection Group and a Protection Policy. A Protection Group enables you to select and group the objects that you want to back up. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a group of settings that define the type of backup, backup schedule and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. You can select which standard or custom policy to use when creating a Protection Group. For more details, see Manage Protection Policy.

> **Note:**
> - As part of SPO backup, the Preservation Hold Library (PHL) is also backed up. The PHL backup feature is not enabled by default. You must explicitly enable the PHL backup option when configuring protection for the SPO site. Similarly, while recovering, you must explicitly enable the PHL recovery option to recover PHL data.
>
> > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To protect OneDrive for Business:

1. Do one of the following:
   - On the **Dashboards page**, from the drop-down at the top of the page, select **Data Protection** and do one of the following:
     - If there is no existing Protection Group, click **Start Data Protection** > **Microsoft 365** > **OneDrive**.
     - Click **Protect** > **Microsoft 365** > **OneDrive**.
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Microsoft 365** > **OneDrive**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Microsoft 365 source and then select **Protect OneDrive**.
2. In the **New Protection** pop-up, do one of the following:
   - Protect OneDrive by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 account. From the **Users** list, select the users you plan to back up.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy:** Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. Click **Protect**.
   - To protect OneDrive by specifying all the options, click **More Options**.
3. From the **Source** drop-down list, select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 source.
4. From the **Users** list, select the users that you plan to backup and click **Save Selection**.

> **Note:** You can use the advanced filters that are available to search and find users based on user attributes. Also, you can click the shield icon to enable Auto Protect. For more information on Auto Protect, see Enable or Disable Auto Protect.

5. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name of the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Use an Existing Protection Policy.
7. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
8. After selecting a Storage Domain, if you need to change any of the advanced settings on the **New Office365 Protection Group** page, scroll down and click **Edit** on the right. For more information, see Advanced Settings.
9. Click **Protect**.

   Your OneDrive for Business Data Protection is now active and running.

Advanced Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection run should start. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. QoS policies are designed for workloads such as backups, which keep a lot of IOs outstanding. We recommend one of the following: • **Backup Target HDD (Default)**: This policy assumes higher latency is OK and the workload will get high throughput if it issues a lot of writes in parallel. The data lands on hard disks if it is sequential and on SSDs if it is random writes. • **Backup Target SSD**: With this policy, both sequential and random IOs land on SSDs, but the IOs are not treated as high-priority IOs and we assume the workload keeps lots of IOs outstanding. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Exclusion List	Select the paths that you plan to exclude from the backup.
Protect Preservation Hold Library	Enable the **Protect Preservation Hold Library** option to protect the Preservation Hold Library, which is used to store files needed for compliance reasons. Consider the following: • **Shared permissions are not restored**: Recovering PHL data recovers the file content only. Any sharing permissions it had must be reapplied manually after recovery. • **No granular recovery**: PHL data can only be recovered as part of a full OneDrive recovery; individual PHL files or folders cannot be recovered. • **PHL data is not searchable**: PHL data is not indexed by standard Microsoft Search and cannot be browsed or searched directly. • **Recovery time and space impact**: Recovering PHL data may increase recovery time substantially. Recovery will fail if sufficient space is not available at the target location. > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature. For more information, see Retention for OneDrive.
Indexing	By default, indexing is enabled. > **Note:** Indexing is mandatory to search for files or folders in a OneDrive.

Advance Settings	Description
Alerts and Priority	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the protection runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	**Optional**. Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

A Protection Policy is a reusable set of settings. It enables you to define how and when OneDrive data is protected and replicated. A standard policy provides various settings that you can configure for a Protection Group.

Create a Protection Policy with the settings listed in the Policy Settings for OneDrive. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for OneDrive

The following table lists the applicable policy settings for OneDrive:

Field	Description
DataLock	Add a DataLock for compliance and regulatory requirement. This ensures that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users from modifying or deleting any Snapshots that were generated by the protection runs that use this Policy. The users include those who have the Data Security role in Cohesity as well. > **Note:** Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained, and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer duration than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the protection run fails. The default time between retries is thirty minutes. You can customize the number of retries and how long to wait between each attempt. > **Note:** Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling the failed requests. There is no need to increase the value for this setting in the Protection Policy.

Field	Description
Quiet Times	Quiet Times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	Replicate snapshots to another Cohesity cluster.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with the applicable policy settings for OneDrive. In this case, all the policy settings apart from the settings applicable to OneDrive are ignored.

To use the Protection Policy:

1. Navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for OneDrive. For information on the applicable policy settings, see Policy Settings for OneDrive.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage the OneDrive Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered Microsoft 365 domain and the assigned Storage Domain in the Protection Group settings.

Add OneDrives to the Protection Group

You can add new OneDrives to the Protection Group for backup. The newly added OneDrive is backed up from the next protection run and it will be a full backup for the newly added user.

To add OneDrives to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - OneDrive** page, click ✎ in the **Users** section.
5. From the **User** list, select the OneDrives that you plan to add to the backup and click **Save Selection**.
6. Click **Save**.

Remove OneDrives from the Protection Group

You can remove OneDrive from the Protection Group that you do not want to back up in the next protection run. When you remove a user from the protection group, incremental backup of the OneDrive is not performed from the next protection run. The existing backup for the user is preserved until the snapshot expires.

To remove OneDrives from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - OneDrive** page, click ✎ in the **Users** section.
5. From the **User** list, unselect the OneDrives that you want to remove from the Protection Group and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new OneDrives that are added to Microsoft 365.

When the Auto Protect option is enabled:

- All the OneDrives that are added to Microsoft 365 in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the OneDrives that are removed from Microsoft 365 in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the user is preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly onboarded users on Microsoft 365 to the Protection Group.
- Remove existing users that are offboarded on Microsoft 365 from the Protection Group. The existing backups of the users are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - OneDrive** page, click ✎ in the **Users** section.
5. From the **User** list, click the **shield** icon to enable or disable Auto Protect.
6. Click **Protect**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up OneDrive. It provides you with options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.
- Add replication.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.

4. In the pop-up window, perform the following:
    1. In the **Backup** drop-down list, select any one of the following:
        - Backup all Objects in the Protection Group
        - Backup only selected Objects in the Protection Group
    2. In the **Backup Type** drop-down list, select the backup type.
    3. Optional. In the **Replicate** section, select the replication target and the retention duration or click **Add** to add the replication target and the retention duration. Also, select **Replicate only fully successful runs** to replicate the data only from the completed protection runs.
    4. Optional. In the **Archive** section, click **Add** to add the external target to which you want to archive the data and the retention duration. Also, select **Archive only fully successful runs** to archive the data only from the completed protection runs.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Download Error Report for a Protection Run

You can now download a list of errors encountered during the OneDrive protection run. Once the protection run is finished and if any files are skipped or have encountered errors, the protection run status is shown as **Warning**. You can download a report that lists the files that are skipped or have encountered errors.

> **Note:** The files skipped or encountered errors are re-tried for the backup in the following protection run.

To download the report:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select a OneDrive Protection Group with an icon in the **Status** column.

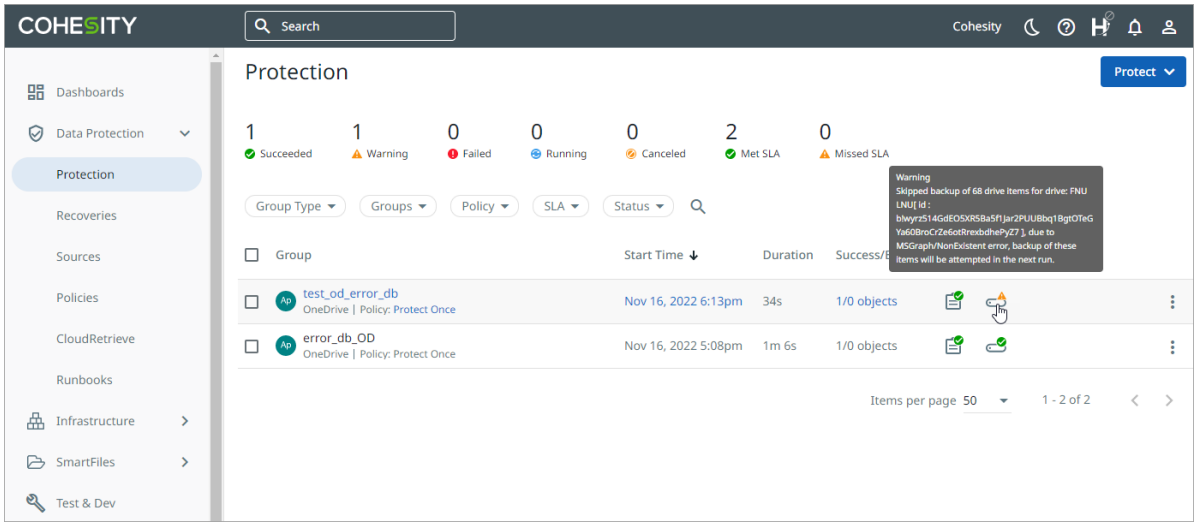

3. In the **Group Details** page, select the protection run with the icon in the **Status** column.

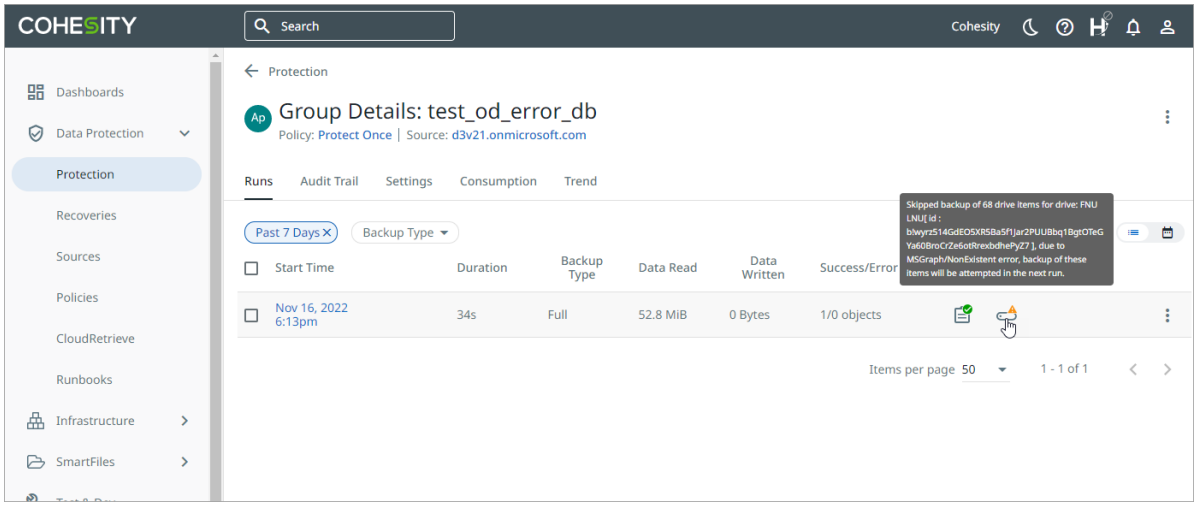

4. On the **Run Details** page, select **view more** to view the list of errors for the protection run and click **Download Error Logs**.

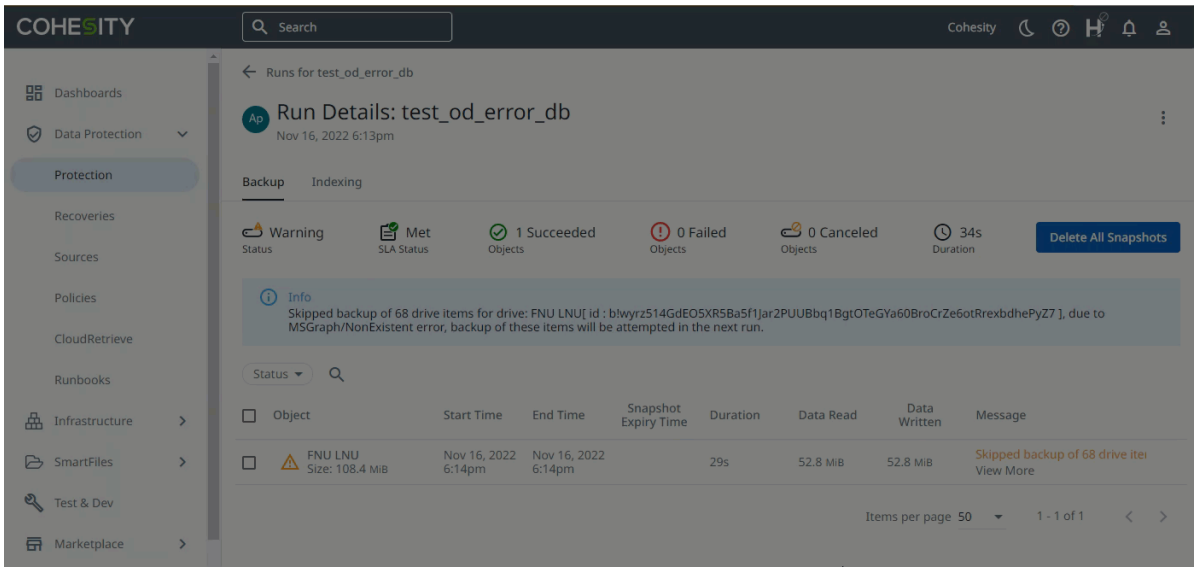

This will download all the error logs in a CSV file format.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:

- Delete Protection Group Only
- Delete Protection Group and Snapshots

Recover OneDrive

You can restore OneDrive for Business data at various granularities, such as at the drive, folder, and file levels. You can restore to the same OneDrive or to a different OneDrive, in the same or a different Microsoft 365 domain.

Recover OneDrive

To recover OneDrive, you need to search through the backups by username or Protection Group name of the OneDrive that you plan to restore and then start the restore operation.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To recover OneDrive:

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **OneDrive** as the recovery source.
2. On the **New Recovery Microsoft 365 - OneDrive** page, in the **Recovery Type**, select **Entire OneDrive**.
3. Search using the user or protection group name, select the user OneDrive you need to restore, and click **Next: Recover Options**.
4. Under **Recover To**, based on your requirements, select one of the following options:

Options	Description
Original Location	Restore data to the same OneDrive from which it was backed up. **Note:** If the files and folders being restored already exist in the folder, the file is restored as the latest version.

Options	Description
New Location	**An alternate folder within the same drive**: Choose the same **Registered Source**, select the **Target OneDrive** as the source user drive and keep the default Folder name or provide the **Folder** name to which the file should be restored.
	Alternate location in the same Microsoft 365 domain: Choose the same **Registered Source**, select the **Target OneDrive** and provide the **Folder** name to which the file should be restored.
	Alternate location in a different Microsoft 365 domain: Choose the target **Registered Source**, select the **Target OneDrive** and provide the **Folder** name to which the file should be restored. **Note:** Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more information, see Register and Manage Microsoft 365 Source.

5. Select your **Recovery Options**:
   - **Continue on Error**. Enable to recover even if errors occur when recovering OneDrive(s). For example, if one of the OneDrives cannot be recovered, Cohesity will still attempt to recover the other selected OneDrive.
   - **Task Name**. Change the default name of the recovery task.
6. Click **Recover**.

Restore Files and Folders

To recover files or folders, you can either search for the files/folders or browse through the backups using the user OneDrive name and then select the file or folder that you plan to restore.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To restore files and folders:

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** as the recovery source.
2. On the New Recovery Microsoft 365 - OneDrive page, in the Recovery Type, select any one of the following tabs to search for the file or folder:
   - **Browse OneDrive and Recover**. You can browse the individual user OneDrive to navigate and select the files or folders to be restored.
   - **Search Files and Recover**. You can use the global search to find the files and folders that need to be restored.
3. If you plan to browse and recover, do the following:

1. In the **Recovery Type** section, select **Browse OneDrive and Recover**.
2. Search for the username and click the username or click **Browse** to browse the drive for the user.
3. Select the file or folder you plan to restore. Do any one of the following based on your requirements:
   - To recover the file(s) or folder(s), click **Next**.
   - To download the file(s) or folder(s), click **Download Files**.

     A new recovery task is created to download the file(s) or folder(s). When the task completes, from the Recoveries page, click the task name and then click Download Files to download the generated zip file.

     > **Note:**
     > - You cannot download empty folders.
     > - A recovery task is not created if you download a single file. Instead, the file is directly downloaded from the browser.
4. Click **Next: Recover Options**.

4. If you plan to search and recover, do the following:
   1. In the **Recovery Type** section, select **Search Files and Recover**.
   2. Search for the file or folder name you plan to recover or Click **Advanced Search** and select **Both**, **Files**, or **Folder** and search based on the available filters and click **Apply**.
   3. Select the file or folder you plan to restore and Click **Next: Recover Options**.
5. Under **Recover To**, based on your requirements, select one of the following options:

Options	Description
Original Location	Restore data to the same OneDrive from which it was backed up. **Note:** If the file being restored already exists in the folder, the file is restored as the latest version.

Options	Description
Restore OneDrive items to an alternate location	**An alternate folder within the same drive**: Choose the same **Registered Source**, select the **Target OneDrive** as the source user drive and keep the default **Folder** name or provide the **Folder** name to which the file should be restored.
	Alternate location in the same Microsoft 365 domain: Choose the same **Registered Source**, select the **Target OneDrive** and provide the **Folder** name to which the file should be restored.
	Alternate location in a different Microsoft 365 domain: Choose the target **Registered Source**, select the **Target OneDrive** and provide the **Folder** name to which the file should be restored. > **Note:** Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more information, see Register and Manage Microsoft 365 Source.

6. Select the **Recovery Options** as required:

Options	Description
Continue on Error	Enable to recover even if errors occur when recovering OneDrive(s). For example, if one of the OneDrives cannot be recovered, Cohesity will still attempt to recover the other selected OneDrive.

Options	Description
Include Preservation Hold Library	Enable this option to recover data from the Preservation Hold Library (PHL). Consider the following points: • **Shared permissions are not restored**: Recovering PHL data recovers the file content only. Any sharing permissions it had must be reapplied manually after recovery. • **No granular recovery**: PHL data can only be recovered as part of a full OneDrive recovery; individual PHL files or folders cannot be recovered. • **PHL data is not searchable**: PHL data is not indexed by standard Microsoft Search and cannot be browsed or searched directly. • **Recovery time and space impact**: Recovering PHL data may increase recovery time substantially. Recovery will fail if sufficient space is not available at the target location. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.
Task Name	Change the default name of the recovery task.

7. Click **Recover**.

Archive OneDrive Data

Cohesity enables you to archive the OneDrive for Business data to a designated external target, such as tape or cloud storage. For more details about external targets, see Manage External Targets .

It provides a policy-based archival solution for OneDrive for Business data. In the Protection Policy, you can define the archive schedule, select the external target, and specify the retention period of the data on the external target.

> **Note:** CloudRetrieve is not supported.

To archive OneDrive for Business data:

1. Register an external target. For more details, see Register an External Target.
2. Edit the existing Protection Policy that you had created for backup or if you are creating a new Protection Policy for backup. Configure the Archival settings as per your requirement. For more information, see **Step 13** in Create or Edit a Standard Policy.

3. Run the Protection Group that you have created for OneDrive for Business based on your requirement or the Archival settings will take effect in the next scheduled Protection Group run.

SharePoint Online

SharePoint Online is a SaaS application bundled with the Microsoft 365 service. It provides an extensive range of collaborative and creative capabilities enabling organizations to share, manage, and access information from almost any device.

Using the policy-based data protection solution from Cohesity, you can back up and recover the SharePoint Online site templates. Thus enabling you to backup and recover the SharePoint Online sites or subsites and their contents such as document libraries, lists, view, and so on.

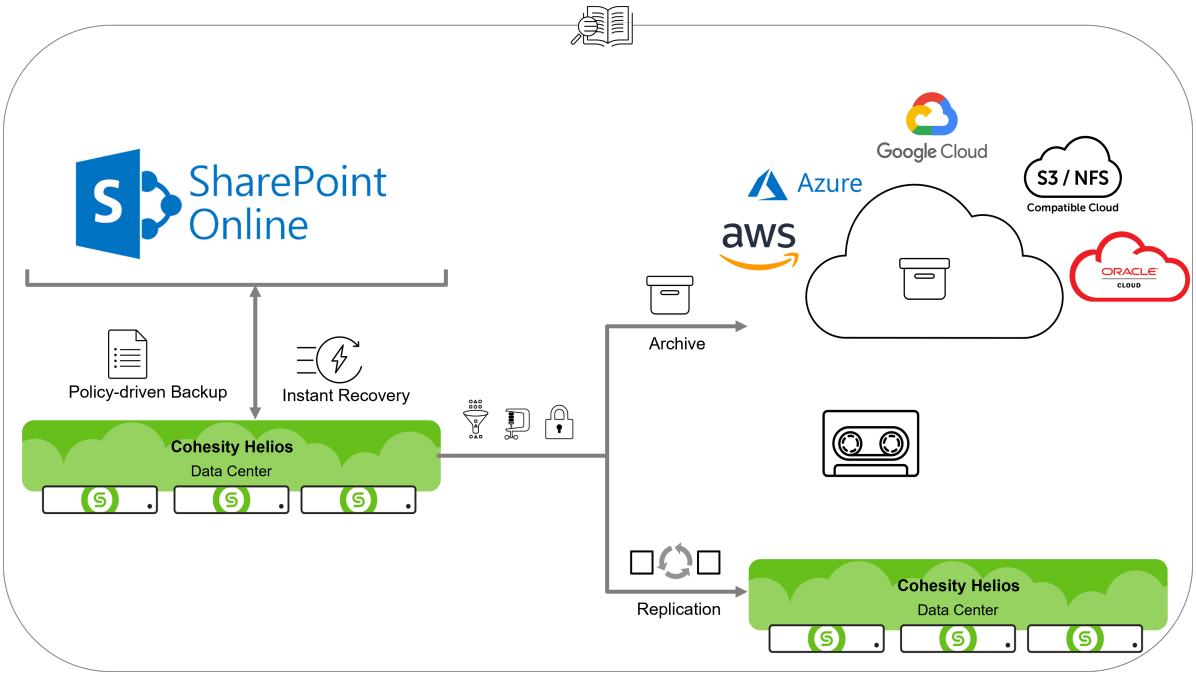

Site Tagging

SharePoint Online sites are containers that are used to store and organize information in SharePoint. Sites can be standalone or can be associated with Groups and Teams referred to as Groups sites and Teams sites respectively. The Site Tagging option categorizes the sites and tags them to their respective containers.

> **Note:** Any Site that is tagged as a Group or Team will not be visible in the Sites section and these sites will be protected through the corresponding Group or Team protection.

> **Important:** Azure ACS reached full retirement across all Microsoft 365 tenants on April 2, 2026. Authentication requests relying on ACS are no longer supported. If your Azure application is currently using ACS-based authentication, Cohesity recommends migrating to Certificate-Based Authentication (CBA) to ensure continued support for SharePoint backup and recovery. The Azure application itself remains in use only the authentication method needs to be updated. For more information, see Microsoft 365 Certificate-Based Authentication.

Plan and Prepare for SharePoint Protection

Before you register your Microsoft 365 domain with Cohesity to protect Microsoft 365 data, ensure:

- You have completed all the Prerequisites.
- You have reviewed and understood the Considerations and Best Practices.

Considerations

> **Important:** Azure ACS reached full retirement across all Microsoft 365 tenants on April 2, 2026. Authentication requests relying on ACS are no longer supported. If your Azure application is currently using ACS-based authentication, Cohesity recommends migrating to Certificate-Based Authentication (CBA) to ensure continued support for SharePoint backup and recovery. The Azure application itself remains in use only the authentication method needs to be updated. For more information, see Microsoft 365 Certificate-Based Authentication.

- Only the list schema is backed up as part of the SharePoint online backup. The list items are not part of the SharePoint online backup.
- Backup and recovery of SharePoint group permissions are not supported.
- Backup and recovery of OneNote files in site drives is supported. Consider the following:
  - Certificate-based authentication must be enabled. For more information, see Microsoft 365 Certificate-Based Authentication.
  - The first OneNote backup after enabling this feature may take longer than usual incremental backups, as it will include all OneNote files that were previously skipped.

- After registration, the size of all discovered SharePoint sites is displayed. However, the size of old or inactive SharePoint sites is not displayed. The size becomes available only after these sites are updated and a subsequent refresh completes.
- After registration, the size of all discovered SharePoint sites is displayed. However, the size of old or inactive SharePoint sites is not displayed. The size becomes available only after these sites are updated and a subsequent refresh completes.
- Document libraries enabled with the ForceCheckout option are not restored.
- Restore of sites with Out of the box (OOTB) modern theme or composed look is not supported.
- Backup and restore of sites or subsites URLs with non-ANSI characters are not supported.
- Restore of a site collection is not supported if the site URL has changed after the backup.
- In a document library, the documents in the initial version v1.0 and checked-out are not backed up.
- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.
- SharePoint sites that are archived or locked will not be backed up and the source discovery may complete early due to these locked sites. Contact your Cohesity account team to configure the cluster to ignore the locked sites during discovery.
- The following List template types are not supported:
    - Custom Workflow Process
    - Workflow History
    - Project Tasks
    - Public Workflows External List
    - MicroFeed (MySiteMicroBlog)
    - Tasks (HierarchyTasksList)
    - Sharing Links
    - Access Requests
- Microsoft 365 items with only special characters will not be listed in granular search. The item name must contain at least one alphanumeric character to be listed in the search result.
- Recovery of Content Type Hub SharePoint Online sites (https://{domain_name}.sharepoint.com/sites/contentTypeHub) is not supported.
- From the recovery workflow, you cannot download an empty folder.

- The system discovers and protects the SharePoint Online sites created in the central storage location of your Microsoft 365 tenant.
- SharePoint Online sites created in satellite storage locations of your Microsoft 365 tenant can be discovered and protected. You can recover to the same tenant and same location or an alternate tenant and default location. For more details, see Multi-Geo Capabilities in SharePoint Online.

  > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

- Cohesity currently does not support Geo-Stretched Microsoft 365 tenants.
- Custom scripts setting is not supported in SharePoint.
- Granular recovery of Team and Group sites is not supported when Site Tagging is enabled.
- During SharePoint granular recovery, you cannot browse or download files from the document libraries that contain a slash (/) in its name. As a workaround, you can perform a full recovery.
- Backup and recovery of Preservation Hold Library (PHL) is supported. Consider the following points:

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  > **Note:** PHL is a hidden, system-managed library that OneDrive and SharePoint automatically create when a retention policy, retention label, or legal/eDiscovery hold is applied. It stores preserved copies of files that are edited or deleted while under retention.

  - **Shared permissions are not restored**: Recovering PHL data recovers the file content only. Any sharing permissions it had must be reapplied manually after recovery.
  - **No granular recovery**: PHL data can only be recovered as part of a full SharePoint recovery; individual PHL files or folders cannot be recovered.
  - **PHL data is not searchable**: PHL data is not indexed by standard Microsoft Search and cannot be browsed or searched directly.
  - **Recovery time and space impact**: Recovering PHL data may increase

recovery time substantially. Recovery will fail if sufficient space is not available at the target location.

- Microsoft 365 Backup Storage service is supported for SharePoint Online.

  > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

- For SharePoint recoveries larger than 1TB, you must adjust the quotas to allow for up to 25TB on the SharePoint recovery for it to function properly. Without this adjustment, the recovery process may fail due to size constraints.
- During the recovery process:
  - The original sharing link is not restored. Instead, a new sharing link with similar settings is created.
  - Passwords associated with sharing links, if set, are not restored.
- Recovery of docx, pptx, xlsx in system docslib is not supported.
- Teams and Groups Sites cannot be recovered directly. The site must be pre-created if backed up from the Sites tab.
- Site tagging can be enabled by default. Sites tagged to Teams or Groups can be directly protected during SharePoint protection if Teams or Groups workflow is not enabled. This ensures that no Teams Sites or Groups Sites are missed during recovery.

  > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

- Recovery always creates a new version. The actual creation time and owner details of the file are not recovered. The current time is displayed as the creation/modification time and the SharePoint app is displayed as the owner/modifier.
- Backup and recovery of custom headers (columns) in SharePoint Online document libraries and lists is supported.

  > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

- Backup and recovery of files with sensitivity labels is supported. For files with sensitivity labels,

    - Only a user with access via the label can open the file after downloading, provided the label is not permanently deleted in SharePoint.
    - When recovering to original location, if the file is already present in the destination, the recovery fails. If the file is not present in the destination, a new file will be created.
    - The files downloaded from SharePoint as encrypted. You can download the same file from the backup storage.
- The following metadata fields are restored when SharePoint data is restored:
    - Created By (author)
    - Modified By (editor)
    - Created Date
    - Modified Date
    - During cross-domain restore operations, the Created By (author) and Modified By (editor) fields are not restored because the corresponding users from the source domain will be unavailable in the target domain.
      The Created By metadata is not restored if the user referenced in the Created By field is deleted. The Modified By and Created By metadata are not restored if the user referenced in the Modified By field is deleted.
- Backup and recovery of system document libraries and templates is supported.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- SharePoint does not support file or folder names that contain unsupported special characters. Backup, restore, or download of such items is not supported. See, Restrictions and limitations in OneDrive and SharePoint.
- Backup and recovery of SharePoint site pages and web parts is supported. You must consider the following points before you proceed with backup and recovery of site pages and web parts:

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

    - Full site page recovery to an alternate location behavior
        - For a full site recovery to an alternate location, links and IDs referenced in recovered site pages and their components are adjusted on a best effort basis.

    - For a full site recovery to an alternate location, reference links and IDs in recovered site pages and their components are recovered on a best effort basis. Broken references may still occur and are unsupported.
    - During recovery of translated site page to an alternate location, these site pages may be recovered twice and may not be correctly linked to their master page.
- Unsupported items for site pages recovery to an alternate location
  When recovering site pages to a new or alternate location, the following items are not recovered:
    - Page analytics
    - Page favorite status
    - Pin to Top status
    - Comments
    - Likes
    - Save for Later actions
- Site page granular recovery behavior
  These considerations apply to both full site recovery and granular recovery, and to both original and alternate location recoveries.
    - Granular recovery to an alternate location always recovers only the content. Referenced images and files are not recovered. Links and IDs referenced in the site page are not changed; as a result, links and web parts may not work.
    - Granular recovery recovers only the latest version of the site pages when recovering to the original location or alternate location.
    - Granular recovery of individual site pages is supported only to an existing site pages document library. Recovery to a different document library is not supported.
- Unsupported site pages features
  Recovered site pages may have layout and functionality inconsistencies for the following:
    - Flexible Sections
    - Picture as Section Background
    - Page Title Formatting
    - Empty Sections
- Site pages library prerequisite
  If the site pages library is deleted, it must be manually created before initiating

recovery. Disable the site pages feature from **Site Settings** > **Features**. Enable the site pages feature. The site pages library will be recreated. Recovery can then be initiated.

- Unsupported web parts
  Backup and recovery of the following web parts are not supported:
  - AI Agent
  - Power Apps
  - Power BI
  - Viva learning
  - Viva engage
  - Any third-party web parts
  - List and list properties web parts
  - Microsoft forms web parts
- Download behavior
  A download is a direct download of backed-up content in its original format. The file cannot be opened independently in a browser.

- Backup and recovery of SharePoint wiki pages is supported. Consider the following points before you proceed with backup and recovery of SharePoint wiki pages.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  - Full site recovery to original location behavior
    - When you recover a full site to the original location, wiki pages, referenced images, and files that belong to the same site are recovered. Paths are also preserved.
  - Full site recovery to alternate location behavior
    - When you recover a full site to an alternate location, wiki pages, referenced images, and files that belong to the same site are recovered. However, paths are not updated. You must manually correct the paths.
  - Granular recovery behavior
    - When you perform a granular recovery, only the wiki page content is recovered and published. Referenced images, files, and their paths are not recovered. This applies to both original and alternative location recovery.
    - Granular recovery of individual site pages is supported only to an existing site pages document library. Recovery of a site page to any different

document library is not supported. This is applicable to alternative location.

- Download behavior
  - All backed-up content is downloaded in its original format. The file cannot be opened independently in a browser.
- Common considerations
  The following considerations are applicable to all recovery scenarios:
  - When content is referenced from outside the recovered site, that content is not recovered.
  - Backup and recovery of classic app/web parts within wiki pages and web part pages is not supported. When such pages are recovered only content will be recovered.

Best Practices

To protect your SharePoint Online data, Cohesity recommends the following best practices:

- In the Protection Policy, set **Retry Options** to **2** or **3**, with a small interval between the retry attempts.
- Use **Auto Protect** to automatically include the newly added sites on SharePoint Online to the Protection Group.
- Avoid adding the same sites in multiple Protection Groups.

Backup SharePoint Online

SharePoint Online is protected using a Protection Group and a Protection Policy. A Protection Group enables you to select the objects that you want to back up. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a reusable list of settings that define the type of backup, backup schedule, and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. You can select which standard or custom policy to use when creating a Protection Group. For more details, see Manage Protection Policy.

Note:

- As part of OneDrive backup, the Preservation Hold Library (PHL) is also backed up. The PHL backup feature is not enabled by default. You must explicitly enable the PHL backup option when configuring protection for the SPO site. Similarly, while recovering, you must explicitly enable the PHL recovery option to recover PHL data.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- SharePoint Online sites associated with Microsoft 365 Groups or modern Microsoft 365 Teams are protected as part of their dedicated workflows. For more information, Groups and Teams.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To protect SharePoint Online sites:

1. Do one of the following:
   - On the **Dashboards** page, from the drop-down at the top of the page, select **Data Protection** and do one of the following:
     - If there is no existing Protection Group, click **Start Data Protection** > **Microsoft 365** > **Sites**.
     - Click **Protect** > **Microsoft 365** > **Sites**.
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Microsoft 365** > **Sites**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Microsoft 365 source and then select **Protect SharePoint Online**.
2. In the **New Protection** pop-up, do one of the following:
   - Protect the sites by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 account. From the **Sites** list, select the sites or subsites that you plan to back up.
     ii. **Protection Group**: Specify a name for the Protection Group.

iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.

iv. **Storage Domain**: Select the appropriate storage domain.

> **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

v. Click **Protect**.

- To protect sites by specifying all the options, click **More Options**.

3. From the **Source** drop-down list, select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 source.
4. From the **Sites** list, select the sites or subsites that you plan to backup and click Save Selection or click the icon to enable Auto Protect at the domain level or site level. For more information on Auto Protect, see Enable or Disable Auto Protect.
5. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name of the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, Create or Use an Existing Protection Policy.
7. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
8. After selecting a Storage Domain, if you need to change any of the advanced settings on the **Microsoft 365 - SharePoint Online** page, scroll down and click **Edit** on the right. For more information, see Advanced Settings.
9. Click **Protect**.

   Your SharePoint Online Data Protection is now active and running.

Advanced Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection run should start. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Exclusion List	Select the folders that you plan to exclude from the backup or click **Add** to add custom folders that you want to exclude from the backup. > **Note:** Exclusion list is not supported for Public Folders.

Advance Settings	Description
Protect Preservation Hold Library	Enable the **Protect Preservation Hold Library** option to protect the Preservation Hold Library, which is used to store files needed for compliance reasons. Consider the following: • **Shared permissions are not restored**: Recovering PHL data recovers the file content only. Any sharing permissions it had must be reapplied manually after recovery. • **No granular recovery**: PHL data can only be recovered as part of a full SharePoint site recovery; individual PHL files or folders cannot be recovered. • **PHL data is not searchable**: PHL data is not indexed by standard Microsoft Search and cannot be browsed or searched directly. • **Recovery time and space impact**: Recovering PHL data may increase recovery time substantially. Recovery will fail if sufficient space is not available at the target location. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.
Indexing	By default, indexing is enabled. **Note:** Indexing is mandatory for granular restore of an Exchange Online mailbox , such as restoring a folder or restoring an email. For more information, see Recover Mailbox .
Alerts and Priority	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.

Advance Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	**Optional**. Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

Create a Protection Policy with the settings listed in the Policy Settings for SharePoint Online. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for SharePoint Online

The following table lists the applicable policy settings for SharePoint Online:

Field	Description
DataLock	Add a DataLock for compliance and regulatory requirement. This ensures that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users from modifying or deleting any Snapshots that were generated by the protection runs that use this Policy. The users include those who have the Data Security role in Cohesity as well. > **Note:** Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained, and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer duration than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the protection run fails. The default time between retries is thirty minutes. You can customize the number of retries and how long to wait between each attempt. > **Note:** Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling the failed requests. There is no need to increase the value for this setting in the Protection Policy.
Quiet Times	Quiet Times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	Replicate snapshots to another Cohesity cluster.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with applicable policy settings for SharePoint Online. In this case, all the policy settings apart from the settings applicable to SharePoint Online are ignored.

To use an existing Protection Policy:

1. Navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for SharePoint Online. For information on the applicable policy settings, see Policy Settings for SharePoint Online.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage the SharePoint Online Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered Microsoft 365 domain and the assigned Storage Domain in the Protection Group settings.

Add Sites or Subsites to the Protection Group

You can add a new SharePoint Online site(s) or subsite(s) to the Protection Group for backup. The newly added sites or subsites are backed up from the next protection run, and it will be a full backup for the recently added user.

To add Sites or Subsites to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - SharePoint Online** page, click ✎ in the **Users** section.
5. From the **Sites** list, select the sites or subsites that you plan to add to the backup and click **Save Selection**.

   > **Note:** If you select a site with subsites, automatically, the subsites are selected for the protection run.

6. Click **Save**.

Remove Sites or Subsites from the Protection Group

You can remove the SharePoint Online site(s) or subsite(s) from the Protection Group that you do not want to back up in the next protection run. When you remove a site from the protection group, the site's incremental backup is not performed from the future protection run. The existing backup for the site is preserved until the snapshot expires.

To remove Sites or Subsites from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - SharePoint Online** page, click ✎ in the **Users** section.
5. From the **Sites** list, unselect the sites or subsites that you plan to remove from the backup and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new sites that are added to Microsoft 365.

When the Auto Protect option is enabled:

- All the sites that are added to SharePoint Online in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the sites that are removed from SharePoint Online in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the site is preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly onboarded sites on SharePoint Online to the Protection Group.
- Remove existing sites that are offboarded on SharePoint Online from the Protection Group. The existing backups of the sites are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - SharePoint Online** page, click ✎ in the **Users** section.
5. From the **Sites** list, click the 🛡 icon in the domain level or site level to enable or disable Auto Protect and click **Save Selection**.
6. Click **Protect**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up SharePoint Online sites. It provides you with options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.
- Add replication.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SharePoint Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. In the **Backup** drop-down list, select any one of the following:
        - Backup all Objects in the Protection Group
        - Backup only selected Objects in the Protection Group
    2. In the **Backup Type** drop-down list, select the backup type.
    3. Optional. In the **Replicate** section, select the replication target and the retention duration.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SharePoint Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Download Error Report for a Protection Run

You can now download a list of errors encountered during the SharePoint Site protection run. Once the protection run is finished and if any files are skipped or have encountered errors, the protection run status is shown as **Warning**. You can download a report that lists the files that are skipped or have encountered errors.

> **Note:** The files skipped or encountered errors are re-tried for the backup in the following protection run.

To download the report:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select a SharePoint Site Protection Group with an icon in the **Status** column.

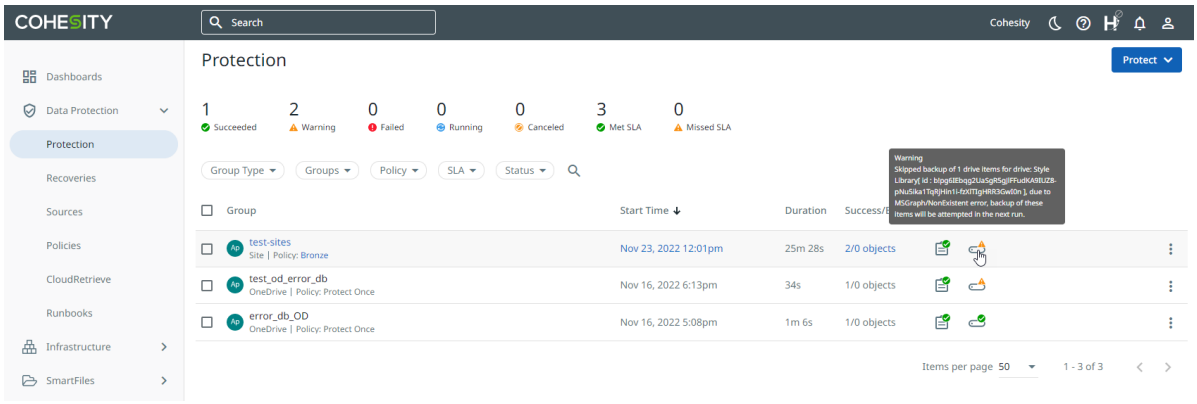

3. In the **Group Details** page, select the protection run with the icon in the **Status** column.

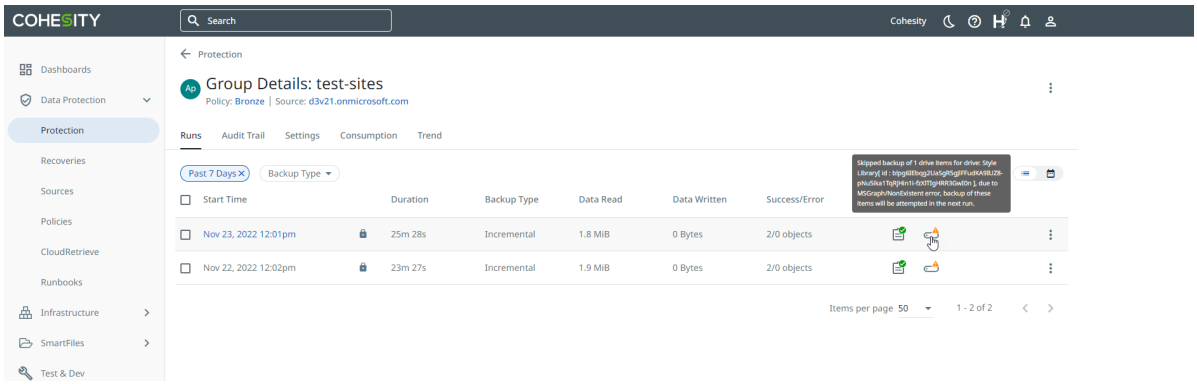

4. On the **Run Details** page, select view more to view the list of errors for the protection run and click **Download Error Logs**.

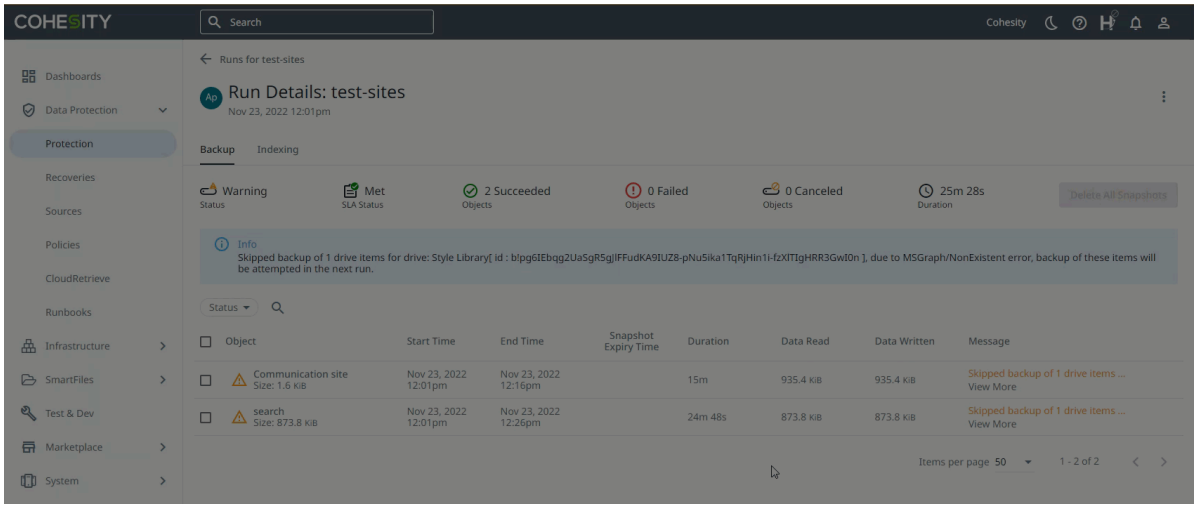

This will download all the error logs in a CSV file format.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group and Snapshots, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for SharePoint Online.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Protection Group Only
   - Delete Protection Group and Snapshots

Recover SharePoint

You can restore the document library on a SharePoint Online Site at various granularities, such as the document library, folder, and file levels. You can restore to the same SharePoint location or a different SharePoint location, in the same or a different Microsoft 365 domain.

Recover Sites

To restore sites, you need to search through the backups by the site name or Protection Group name of the Site that you plan to restore and then start the restore operation.

You can also restore the subsites of a site collection. Search for site collection (root site) using the site collection name or its relative site URL, for example, `/sites/myroot`. If the site collection is part of a single backup snapshot, you can select the subsites in the site collection, and the subsites are restored in their original order.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To restore sites:

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Sites** as the recovery source.
2. On the **New Recovery Microsoft 365 - SharePoint Online** page, under the **Recovery Type** section, click **Entire Site**.
3. Search using the site name or protection group name, select the site(s) that you plan to restore, and click **Next: Recover Options**.

   > **Note:** Cohesity displays the Microsoft 365 site title (Site Name) registered during source discovery. This title corresponds to the site display name in Microsoft 365 and may not match the site name segment in the SharePoint URL. For example, a site whose URL ends in /sites/ProjectAlpha may appear in Cohesity under a different display title configured in Microsoft 365. When locating objects, use the Site Name shown in the Cohesity UI rather than assuming it matches the URL path.

4. Under **Recover To**, based on your requirements, select one of the following options:

Options	Description
Original Location	Click **Recover** to restore data to the same SharePoint site from which it was backed up.
New Location	**Alternate site in the same Microsoft 365 domain**: Choose the **Site** to which the document libraries need to be restored.
	Alternate site in a different Microsoft 365 domain: Choose the **Microsoft 365 Server** to which the site needs to be restored. > **Note:** Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more information, see Register and Manage Microsoft 365 Source.

5. Select the **Recovery Options** as required:

Options	Description
Continue on Error	Enable to recover even if errors occur when recovering OneDrive(s). For example, if one of the OneDrives cannot be recovered, Cohesity will still attempt to recover the other selected OneDrive.
Include Preservation Hold Library	Enable this option to recover data from the Preservation Hold Library (PHL). Consider the following points: • **Shared permissions are not restored**: Recovering PHL data recovers the file content only. Any sharing permissions it had must be reapplied manually after recovery. • **No granular recovery**: PHL data can only be recovered as part of a full SharePoint site recovery; individual PHL files or folders cannot be recovered. • **PHL data is not searchable**: PHL data is not indexed by standard Microsoft Search and cannot be browsed or searched directly. • **Recovery time and space impact**: Recovering PHL data may increase recovery time substantially. Recovery will fail if sufficient space is not available at the target location. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.
Task Name	Change the default name of the recovery task.

6. Click **Recover** to start the restore operation.

Restore Document Libraries and Files

To restore document libraries or files, you can either search for the files/document libraries or browse through the backups using the site name and then select the file or document library that you plan to restore.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

> **Note:** When you restore a deleted SharePoint site, the associated document libraries are automatically restored as well.

To restore files and document library:

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Sites** as the recovery source.

2. On the **New Recovery Microsoft 365 - SharePoint Online** page, under the **Recovery Type** section, select any one of the following to search for the file or folder:
    - **Browse Site and Recover**. You can browse the individual site to navigate and select the files/document library to be restored.
    - **Search Files and Recover**. You can use the global search to find the files and document libraries that need to be restored.
3. If you plan to browse and recover, do the following:
    1. In the **Recovery Type** section, select **Browse Site and Recover**.
    2. Search for the site name and click the site name to browse the site.
    3. Select the file or document library you plan to restore. Do any one of the following based on your requirements:
        - To recover the file(s) or document library(s), click **Next**.
        - To download the file(s) or document library(s), click **Download Files**.

        A new recovery task is created to download the file(s) or document library (s). When the task is completed, from the **Recoveries** page, click the task name and then click **Download Files** to download the generated zip file.

        > **Note:**
        > - You cannot download empty folders.
        > - A recovery task is not created, if you download a single file. Instead the file is directly downloaded from the browser.

    4. Click **Next: Recover Options**.
4. If you plan to search and recover, do the following:
    1. In the **Recovery Type** section, select **Search Files and Recover**.
    2. Search for the file or document libraries you plan to recover or click **Advanced Search** and select **Both**, **Files**, or **Folder** and search based on the available filters and click **Apply**.
    3. Click **Next: Recover Options**.
5. Under **Recover To**, based on your requirements, select one of the following options:

Options	Description
Original Location	Click **Recover** to restore data to the same SharePoint site from which it was backed up. > **Note:** If the file being restored already exists in the document library, the file is restored as the latest version.
New Location	**An alternate document library within the same site**: Choose the target **Document Library** that is the same as the source site and keep the existing document library name or provide the document library name to which the file should be restored.
	Alternate location in a different Microsoft 365 domain: Choose the **Microsoft 365 Server** to which data needs to be restored and then specify the **Document Library** to which the data needs to be restored. > **Note:** Both the Microsoft 365 domains must be registered as sources with the same Cohesity cluster. For more information, see Register and Manage Microsoft 365 Source.

6. Click **Recover** to start the restore operation.

Groups

Microsoft 365 groups are used for collaboration between users, both inside and outside your company. With each Microsoft 365 group, members get a group email and shared workspace for conversations, files, and calendar events, and a planner.

Using the policy-based data protection solution from Cohesity, you can backup and recover Unified Groups data in Microsoft 365.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

The following are the data backed by Cohesity for Microsoft 365 Groups:

- Shared Outlook Inbox
- Shared Calendars
- Shared SharePoint Document Library

> **Important:** Mail-enabled security groups, security groups, and distribution lists are not supported.

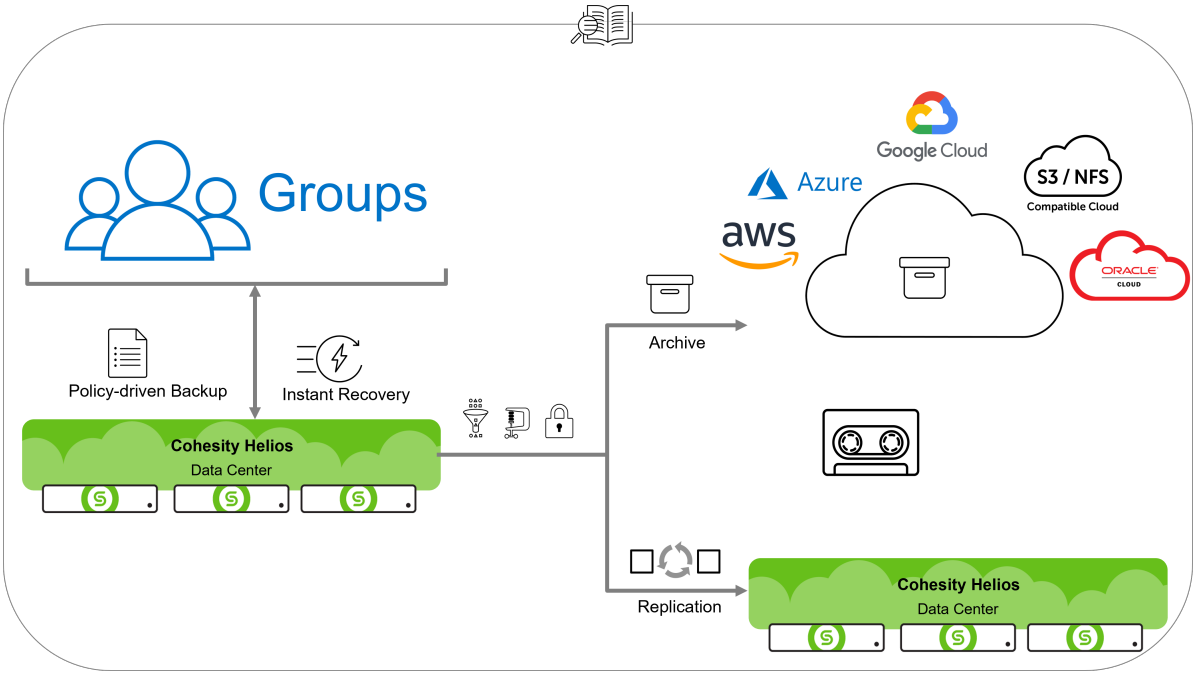

Plan and Prepare for Groups Protection

Before you register your Microsoft 365 domain with Cohesity to protect Microsoft 365 data, ensure:

- You have completed all the Prerequisites.
- You have reviewed and understood the Considerations and Best Practices.

Considerations

- Restoring to an alternate Microsoft 365 domain is not supported.
- Backup and recovery of OneNote files in Groups is supported. Consider the following:
  - Certificate-based authentication must be enabled. For more information, see Microsoft 365 Certificate-Based Authentication.
  - The first OneNote backup after enabling this feature may take longer than usual

incremental backups, as it will include all OneNote files that were previously skipped.

- Using a Group recovery job, you can restore only one Group.
- Recovering system document libraries (libraries that are present in the Group site by default - Site Assets, Style Library, Documents, and Form Templates) is not supported. You can recover only the non-system document libraries in a Group site.

  > **Note:** Files associated with Private Channels in Teams saved under the Documents library can be protected and recovered.

- Cohesity does not support Groups data protection in a multitenant environment.
- Groups must have owners for the backups to function properly.
- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.
- Microsoft 365 items with only special characters will not be listed in granular search. The item name must contain at least one alphanumeric character to be listed in the search result.
- Backup and recovery of files with sensitivity labels is supported.

  For files with sensitivity labels,

  - Only a user with access via the label can open the file after downloading, provided the label is not permanently deleted in SharePoint.
  - When recovering to original location, if the file is already present in the destination, the recovery fails. If the file is not present in the destination, a new file will be created.
  - The files downloaded from SharePoint are encrypted. You can download the same file from the backup storage.
- The following metadata fields are restored when SharePoint data is restored:
  - Created By (author)
  - Modified By (editor)
  - Created Date
  - Modified Date

    During cross-domain restore operations, the Created By (author) and Modified By (editor) fields are not restored because the corresponding users from the source domain will be unavailable in the target domain.
    The Created By metadata is not restored if the user referenced in the Created

By field is deleted. The Modified By and Created By metadata are not restored if the user referenced in the Modified By field is deleted.

- Backup and recovery of system document libraries and templates is supported.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- SharePoint does not support file or folder names that contain unsupported special characters. Backup, restore, or download of such items is not supported. See, Restrictions and limitations in OneDrive and SharePoint.
- Backup and recovery of subsites is supported. Consider the following points before you proceed for backup and recovery of subsites.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  - When you search for a site using its display name, only items from that specific site are displayed. However, if you search using the URL of the site, items from that site and all its subsites are displayed.

Best Practices

To protect your SharePoint Online data, Cohesity recommends the following best practices:

- In the Protection Policy, set **Retry Options** to **2** or **3**, with a small interval between the retry attempts.
- Use **Auto Protect** to automatically include the newly added sites on SharePoint Online to the Protection Group.
- Avoid adding the same sites in multiple Protection Groups.

Backup Microsoft 365 Groups

Microsoft 365 Groups data is protected using a Protection Group and a Protection Policy. A Protection Group enables you to select the objects that you want to back up. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a reusable list of settings that define the type of backup, backup schedule, and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. You can select which standard or custom policy to use when creating a Protection Group. For more details, see Manage Protection Policy.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To protect Microsoft 365 Groups:

1. Do one of the following:
   - On the **Dashboards** page, from the drop-down at the top of the page, select **Data Protection** and do one of the following:
     - If there is no existing Protection Group, click **Start Data Protection** > **Microsoft 365** > **Group**.
     - Click **Protect** > **Microsoft 365** > **Group**.
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Microsoft 365** > **Group**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Microsoft 365 source and then select **Protect Group**.
2. In the **New Protection** pop-up, do one of the following:
   - Protect Group by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 account. From the **Group** list, select the group(s) that you plan to back up.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. Click **Protect**.
   - To protect Teams by specifying all the options, click **More Options**.
3. From the **Source** drop-down list, select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 source.
4. From the **Groups** list, select the group(s) that you plan to back up or click the icon to enable Auto Protect at the domain level. For more information on Auto Protect, see Enable or Disable Auto Protect.
5. In the **Protection Group** section, do one of the following:

- Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name of the Protection Group.
- Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.

6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, Create or Use an Existing Protection Policy.
7. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
8. After selecting a Storage Domain, if you need to change any of the advanced settings on the **Microsoft 365 - Groups** page, scroll down and click **Edit** on the right. For more information, see Advanced Settings.
9. Click **Protect**.

   Your Microsoft 365 Groups Data Protection is now active and running.

Advanced Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection run should start. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Indexing	By default, indexing is enabled.
Alerts and Priority	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs

Advance Settings	Description
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	**Optional**. Enter a brief description about the Protection Group.

Create or Use an Existing Protection Policy

Create a Protection Policy with the settings listed in the Policy Settings for Microsoft 365 Groups. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Microsoft 365 Groups

The following table lists the applicable policy settings for Microsoft 365 Groups:

Field	Description
DataLock	Add a DataLock for compliance and regulatory requirement. This ensures that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users from modifying or deleting any Snapshots that were generated by the protection runs that use this Policy. The users include those who have the Data Security role in Cohesity as well. > **Note:** Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained, and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer duration than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the protection run fails. The default time between retries is thirty minutes. You can customize the number of retries and how long to wait between each attempt. > **Note:** Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling the failed requests. There is no need to increase the value for this setting in the Protection Policy.
Quiet Times	Quiet Times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.
Replication	Replicate snapshots to another Cohesity cluster.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with applicable policy settings for Microsoft 365 Groups. In this case, all the policy settings apart from the settings applicable to Microsoft 365 Groups are ignored.

To use an existing Protection Policy:

1. Navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for Microsoft 365 Groups. For information on the applicable policy settings, see Policy Settings for Microsoft 365 Groups.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage Groups Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered Microsoft 365 domain and the assigned Storage Domain in the Protection Group settings.

Add Microsoft 365 Groups to the Protection Group

You can add a new Microsoft 365 Groups to the Protection Group for backup. The newly added Microsoft 365 Groups are backed up from the next protection run, and it will be a full backup for the recently added user.

To add Microsoft 365 Groups to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Groups** page, click ✎ .
5. From the **Groups** list, select the Microsoft 365 Groups that you plan to add to the backup and click **Save Selection**.
6. Click **Save**.

Remove Microsoft 365 Groups from the Protection Group

You can remove the Microsoft 365 Groups from the Protection Group that you do not want to back up in the next protection run. When you remove a site from the protection group, the site's incremental backup is not performed from the future protection run. The existing backup for the site is preserved until the snapshot expires.

To remove Microsoft 365 Groups from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Groups** page, click ✎ .
5. From the **Groups** list, unselect the Microsoft 365 Groups that you plan to remove from the backup and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new sites that are added to Microsoft 365.

When the Auto Protect option is enabled:

- All the groups that are added to Microsoft 365 in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the Groups that are removed from Microsoft 365 in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backups for the user are preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly onboarded groups on Microsoft 365 to the Protection Group.
- Remove existing groups that are offboarded on Microsoft 365 from the Protection Group. The existing backups of the users are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Groups** page, click ✎ .
5. From the **Groups** list, click the 🛡 icon in the domain level to enable or disable Auto Protect and click **Save Selection**.
6. Click **Protect**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to

your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up Microsoft 365 Groups. It provides you options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.
- Add replication.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Microsoft 365 Groups.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. In the **Backup** drop-down list, select any one of the following:
        - Backup all Objects in the Protection Group
        - Backup only selected Objects in the Protection Group
    2. In the **Backup Type** drop-down list, select the backup type.
    3. Optional. In the **Replicate** section, select the replication target and the retention duration.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Microsoft 365 Groups.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered

To resume a paused protection run:

1. Navigate to > **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Download Error Report for a Protection Run

You can now download a list of errors encountered during the Groups protection run. Once the protection run is finished and if any files are skipped or have encountered errors, the protection run status is shown as **Warning**. You can download a report that lists the files that are skipped or have encountered errors.

> **Note:** The files skipped or encountered errors are re-tried for the backup in the following protection run.

To download the report:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select a OneDrive Protection Group with an icon in the **Status** column.

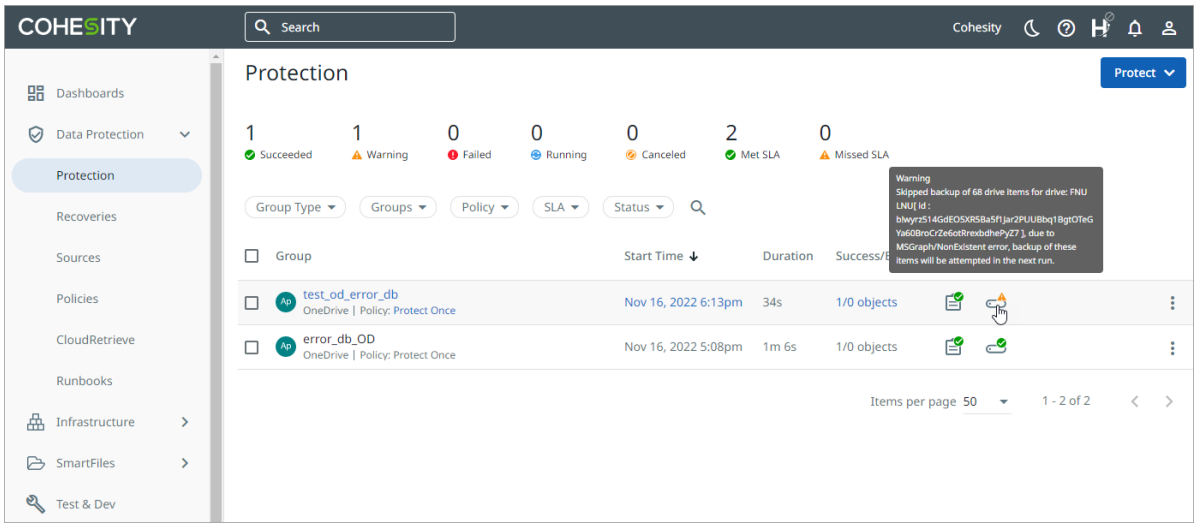

3. In the **Group Details** page, select the protection run with the icon in the **Status** column.

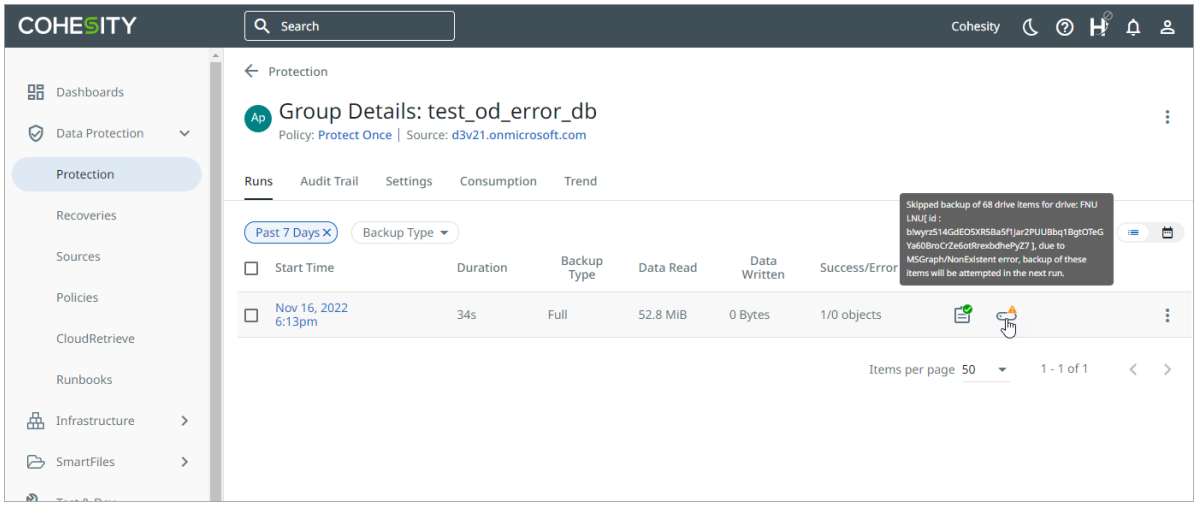

4. On the **Run Details** page, select **view more** to view the list of errors for the protection run and click **Download Error Logs**.

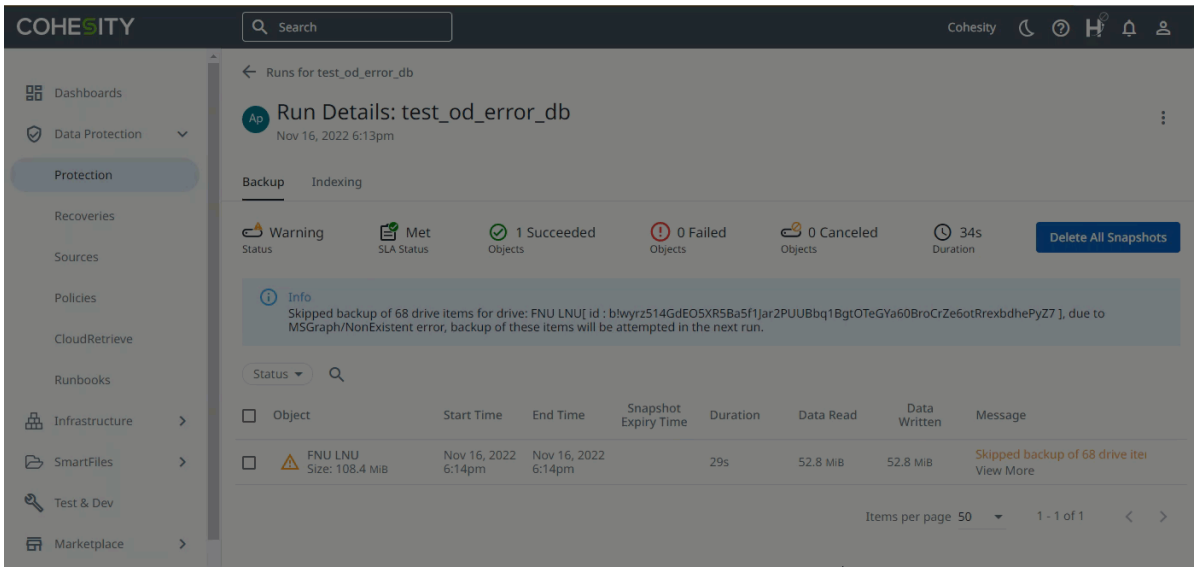

This will download all the error logs in a CSV file format.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group and Snapshots, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity cluster.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Microsoft 365 Groups.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:

- Delete Protection Group Only
- Delete Protection Group and Snapshots

Recover Groups

You can restore Microsoft 365 Group data using Cohesity. In the same Microsoft 365 domain, you can restore data to the same Microsoft 365 Group, to an alternate Microsoft 365 Group, or to a new Microsoft 365 Group.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

Points to note:

- If you're restoring a Microsoft 365 Group that does exist in the Microsoft 365 domain, Cohesity creates a new Microsoft 365 Group with the metadata and data from the backup snapshot.
- If you restore to an existing Microsoft 365 Group, the group resources in the existing Microsoft 365 Group are overwritten or appended with the restored data. The following table details the group resources that are overwritten or appended:

Restore Behaviour	Group Resource Type
Appended	members
	owners
	mails (data)
Overwritten	hideFromAddressLists
	hideFromOutlookClients
	displayName
	visibility
	securityEnabled
	description
	theme

To restore Microsoft 365 Groups, you can search through the backups by group name or using a string pattern in the group name that you plan to restore and then start the restore operation.

> **Note:** You use * character to list all the groups.

You can restore the Microsoft 365 Group data to:

- The same Microsoft 365 Group
- A different Microsoft 365 Group in the same Microsoft 365 domain
- A new Microsoft 365 Group in the same Microsoft 365 domain.

To restore Microsoft 365 Groups:

1. Navigate to **Data Protection > Recoveries**. Click **Recover** on the top-right of the page and then select **Microsoft 365** > **Groups** as the recovery source.
2. Search and select the Microsoft 365 Group you need, and click **Next:Recover Options**.

   > **Note:** Cohesity displays the Microsoft 365 site title (Site Name) registered during source discovery. This title corresponds to the site display name in Microsoft 365 and may not match the site name segment in the SharePoint URL. For example, a site whose URL ends in /sites/ProjectAlpha may appear in Cohesity under a different display title configured in Microsoft 365. When locating objects, use the Site Name shown in the Cohesity UI rather than assuming it matches the URL path.

3. Under **Recover To**, select any one of the following options:

Options	Description
Original Location	**Recover to the same Microsoft 365 Group**: To restore data to the same group from which it was backed up.
New Location	**Alternate Microsoft 365 Group in the same Microsoft 365 domain:**: Enter the **Group Name** to which the data needs to be restored. > **Note:** If the **Group Name** (display name or nickname) matches an existing group, the data is restored to the existing group. If not, Cohesity creates a new group and restores the data in it.

4. Under **Recovery Options**, perform the following:
    - Enable the **Continue on Error** option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
    - In the **Task Name** field, enter the name of the recovery task, if you want to change the default recovery task name.
5. Click **Start Recovery** to start the restore operation.

Recover Groups Items

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

After you protect the Groups in your domain, you can recover the following Group items:

- Group Mailbox items, such as Group emails, mailbox folders, and calendar items.
- Group Site items, such as, document libraries.

> **Important:** Before you can recover a Group's items, you need to set up Microsoft 365 Groups protection with Indexing enabled.

Recover Groups Mailbox Items

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To recover a Groups Mailbox items:

1. In **DataProtect**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name and select the **Group** tab.
3. Above the tree, select **Show All** > **Protected**.
4. Find the Group you need and click the **Recover Groups Items** button on that row to open the **New Recovery** form.
5. On the **New Recovery Microsoft 365 - Groups** page, in the **Recovery Type**, select **Search Group's Mailbox Items**.
6. From the **Item Type** drop-down, do the following:
    - Select **Emails & Folders** and use the global search to find the emails and folders that needs to be restored.

- Select **Calendars** and use the global search to find the calendar item that needs to be restored.

Click **Next: Recover Options** to return to the form.

7. Under **Recover To**, select **Original Location** or **New Location**.
   - If you choose **New Location**, select a **Registered Source** and the **Target MS Group**.
   - Click **Download as PST** to export the backed-up email(s), mailbox folder(s), or calendar items as a PST file. To enhance security, enable **Password Protection** under **Recovery Options** and set a password for the exported PST file.

     A new recovery task is created to download the PST file. When the task is completed, from the Activity page, click the task name and then click **Download Files** to download the generated zip file.
8. Select your Recovery Options:
   - **Continue on Error**. Enable to recover even if errors occur when recovering Groups. For example, if one of the Groups cannot be recovered, Cohesity will still attempt to recover the other selected Groups.
   - **Task Name**. Change the default name of the recovery task.
   - **Password Protection**. Enable to add a password for the exported PST file.
9. Click **Start Recovery**.

Recover Groups Site Items

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To recover a Groups Site items:

1. In **DataProtect**, go to **Sources** to set up your recovery task.
2. Click on the **Source** name and select the **Group** tab.
3. Above the tree, select **Show All** > **Protected**.
4. Find the Group you need and click the **Recover Groups Items** button on that row to open the **New Recovery** form.
5. On the **New Recovery Microsoft 365 - Groups** page, in the **Recovery Type**, select **Search Group's Site Items**.
6. Search and select the Group Site items you need, and click **Next: Recover Options**.
7. Under **Recover To**, select **Original Location** or **New Location**.
   - If you choose **New Location**, select a **Registered Source** and the **Target MS Group**. Then, enter the name of the **Document Library**.

- Click **Download** to download the Group Site item.

  A new recovery task is created to download the Group Site item(s). When the task is completed, from the **Activity** page, click the task name and then click **Download Files** to download the generated zip file.

8. Select your Recovery Options:
   - **Continue on Error**. Enable to recover even if errors occur when recovering Groups. For example, if one of the Groups cannot be recovered, Cohesity will still attempt to recover the other selected Groups.
   - **Restore Original Owner Member**. Enable to restore the original owner member as part of the recovery.
   - **Additional Group Owner**. Enable to add additional group owners to the recovered Group.
   - **Task Name**. Change the default name of the recovery task.
9. Click **Start Recovery**.

Teams

Microsoft Teams is a collaboration solution provided by Microsoft that is bundled with the Microsoft 365 service. For more information, see Microsoft documentation. Using the policy-based data protection solution from Cohesity, you can backup and recover Teams data in Microsoft 365.

The following are the data backed by Cohesity for Microsoft 365 Teams:

- Standard (public) channels
  - Files associated with these channels
  - Conversations/post associated with these channels
- Private channels
  - Files associated with these channels
  - Conversations/post associated with these channels
- Group mailboxes associated with the Teams

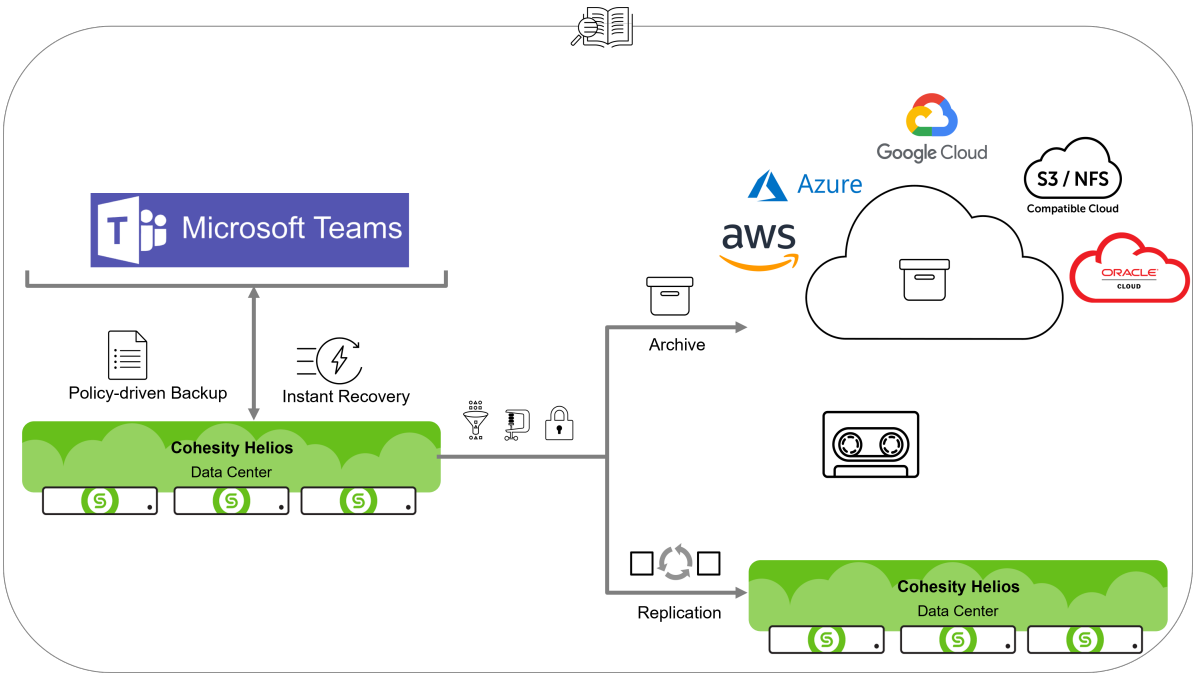

Plan and Prepare for Teams Protection

Before you register your Microsoft 365 domain with Cohesity to protect Microsoft 365 data, ensure:

- You have completed all the Prerequisites.
- You have reviewed and understood the Considerations and Best Practices.

Considerations

- Microsoft Teams data recovery capabilities:
  - Recovery of Teams data to the original Microsoft 365 domain.
  - Recovery of Teams data to an alternate Microsoft 365 domain.
  - Recovery of Teams data to the same Microsoft 365 domain but to a different channel.
  - Recovery of multiple Teams to the original Microsoft 365 domain.
  - Recovery of data to an existing public or private channel. You can also create a new public or private channel during the recovery operation.
  - Recover data to an existing or a new shared channel. You can also create a new shared channel during the recovery operation.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- A Team must have at least one owner with a license and login-enabled for that owner.
- Alternate restore of multiple Teams is not supported. But, you can recover multiple Teams to the original Microsoft 365 domain.
- Restoring deleted private channels or private channels in a deleted team to the original location is not supported.
- Private channels are not restored as a part of the alternate restore.
- Recovery of the following Team data is not supported:
  - Channel Tabs
  - Team Apps
  - Team Site Templates
  - Team Site Lists
  - Public or Private Channel conversations or posts
  - Files that belong to system document libraries.
- Downloading of conversations/posts in channels is supported.
- Backup and download of the following is not supported:
  - Self-message (messages sent to self)
  - Saved or pinned Posts property in the conversation
  - Meeting recordings metadata
  - Emojis and stickers
- Recovery of system document libraries that are included by default in a Team site, such as Site Assets, Style Library, Documents, and Form Templates is not supported. Only non-system document libraries within a Team site can be recovered. Additionally, restoring Team site templates is not supported. The original appearance and configuration of the Team site will not be restored; instead, the new Team site will be created using Microsoft's default template.
- Backup and recovery of OneNote files in Teams is supported. Consider the following:
  - Certificate-based authentication must be enabled. For more information, see Microsoft 365 Certificate-Based Authentication.
  - The first OneNote backup after enabling this feature may take longer than usual incremental backups, as it will include all OneNote files that were previously skipped.

- If the private channel members are deleted from the channel but are still part of the Team or changed from owner to member, the private channel members are restored. However, if a member has changed to the owner, the change is retained after the restoration.
- The entities protected for Teams include the SharePoint sites associated with the Team.
- New Teams Channels created as part of the restore are created with the suffix: `_coh<uniqueint>`. Where `_coh` is appended to the channel name, and `uniqueint` is a random integer value.
- A new Team is created when you restore the deleted Teams to their Original Location. If the restore fails, you need to clean up the newly created team manually.
- Restoration fails if the group properties are not updated appropriately during restoration. In this case, you need to retry the restore.
- For alternate search and restore of Teams files, the new Teams Channels created as target Team Channels will be public channels.
- Cohesity does not support Teams data protection in a multitenant environment.
- Suppose folders such as Feeds, Sync Issues, Legacy Archive Journals, Outbound, Managed Folders, Files, Yammer Root, Clutter, MeContact, and Archive, are not already present. In that case, these folders are skipped during restore.
- Backup and recovery of channel tabs are not supported.
- File/folder permissions with no role will not be recovered since Microsoft does not allow adding such permissions via Graph APIs. For example, no role will be assigned when a Restricted View permission is added to a user from Advanced Setting.
- Microsoft 365 items with only special characters will not be listed in granular search. The item name must contain at least one alphanumeric character to be listed in the search result.
- Recovery of Private Channels is not supported in GCC High.
- Protection of Teams Chat is not chargeable.

  Microsoft Teams Export APIs used to retrieve Teams Chat data for backup were previously billed on a per-message basis. As of August 25, 2025, Microsoft no longer meters or bills these APIs, and no billing configuration is required.

  For more information, see Microsoft's Metered APIs and Services in Microsoft Graph FAQ.
- Backup and recovery of files with sensitivity labels is supported.

  For files with sensitivity labels,

  - Only a user with access via the label can open the file after downloading, provided the label is not permanently deleted in SharePoint.
  - When recovering to original location, if the file is already present in the destination, the recovery fails. If the file is not present in the destination, a new file will be created.
  - The files downloaded from SharePoint are encrypted. You can download the same file from the backup storage.
- The following metadata fields are restored when SharePoint data is restored:
  - Created By (author)
  - Modified By (editor)
  - Created Date
  - Modified Date

    During cross-domain restore operations, the Created By (author) and Modified By (editor) fields are not restored because the corresponding users from the source domain will be unavailable in the target domain.
    The Created By metadata is not restored if the user referenced in the Created By field is deleted. The Modified By and Created By metadata are not restored if the user referenced in the Modified By field is deleted.
- Teams meeting recordings saved the default location can be excluded from backup. This exclusion is enabled by default for new protections. For existing protections, you need to edit the protection, select the **Teams Meeting Recordings** checkbox, and then save the changes. Recordings will be excluded starting with the next backup run. Backup and recovery of Teams meeting recordings metadata is not supported.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- Backup and recovery of system document libraries and templates is supported.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- SharePoint does not support file or folder names that contain unsupported special characters. Backup, restore, or download of such items is not supported. See, Restrictions and limitations in OneDrive and SharePoint.
- Backup and download for Teams Chat are supported; however, restore is not supported currently.

- Backup and recovery of subsites is supported. Consider the following points before you proceed for backup and recovery of subsites.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  - When you search for a site using its display name, only items from that specific site are displayed. However, if you search using the URL of the site, items from that site and all its subsites are displayed.
- Backup and recovery of shared channels is supported. You must review the following considerations before you proceed with backup and recovery of shared channels.

  > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  - Supported recovery scenarios
    - Full recovery of shared channels is supported.
    - Granular recovery, that is recovery of items within shared channels, is supported.
  - Supported recovery destinations
    - Recovery to the original location is supported.
    - Recovery to an alternative location (different Teams and channels) in the same tenant is supported.
    - Recovery to an alternative location (different Teams and channels) in another tenant is supported, that is, cross-tenant recovery is supported.
  - Backup and recovery of the following items in shared channels is not supported:
    - Sensitivity labels
  - Recovery of member teams associated with shared channels is not supported.
  - Once enabled, shared channels are automatically included in the next Teams backup, no job changes are required.
  - Shared channels are indexed and searchable, with support to download or restore content and messages to original or alternate locations.
  - Backup and recovery is supported even when a shared channel has no owner.
- Backup and recovery of Microsoft Teams tabs is supported. You must review the following considerations before you proceed with the backup and recovery of Teams tabs.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

- Supported tab types
  - Backup and recovery is supported only for configurable channel tabs and group chat tabs.
  - Personal tabs added only for one user are not supported.
- Recovery limitations
  - Only the last backed-up version of a tab is restored.
- Restore scope
  - Tabs are restored only as part of their associated Teams channel or group chat.
  - Granular restore of individual tabs or tab components is not supported.
  - Tabs cannot be restored independently outside their Team or channel.
- Permissions and access
  - Restoring tabs does not modify user permissions on the associated content.
  - Restoring PDF and OneNote tabs in newly created tenants requires Certificate-Based Authentication (CBA).
  - In existing tenants, the registered applications must have SharePoint Administrator permissions.

Best Practices

To protect your Microsoft 365 Teams data, Cohesity recommends the following best practices:

- In the Protection Policy, set **Retry Options** to **2** or **3**, with a small interval between the retry attempts.
- Use **Auto Protect** to automatically include the newly added sites on Microsoft 365 Teams to the Protection Group.
- Avoid adding the same sites in multiple Protection Groups.

Backup Microsoft 365 Teams

Microsoft 365 Teams data is protected using a Protection Group and a Protection Policy. A Protection Group enables you to select the objects that you want to back up. It uses the schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a reusable list of settings that define the type of backup, backup schedule, and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. You can select which standard or custom policy to use when creating a Protection Group. For more details, see Manage Protection Policy.

Before you proceed with backup, you must review the considerations. For more information, see Considerations.

To protect Microsoft 365 Teams:

1. Do one of the following:
   - On the **Dashboards** page, from the drop-down at the top of the page, select **Data Protection** and do one of the following:
     - If there is no existing Protection Group, click **Start Data Protection** > **Microsoft 365** > **Teams**.
     - Click **Protect** > **Microsoft 365** > **Teams**.
   - Select **Data Protection > Protection**. Click **Protect** on the top-right of the page and then select **Microsoft 365** > **Teams**.
   - Select **Data Protection > Sources**. Click the actions menu (⋮) next to the Microsoft 365 source and then select **Protect Teams**.
2. In the **New Protection** pop-up, do one of the following:
   - Protect Teams by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 account. From the **Teams** list, select the Team(s) that you plan to back up.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. Click **Protect**.
   - To protect Teams by specifying all the options, click **More Options**.
3. From the **Source** drop-down list, select the registered Microsoft 365 source or click **Register Source** to register a Microsoft 365 source.

4. From the **Teams** list, select the Team(s) that you plan to back up or click the icon to enable Auto Protect at the domain level. For more information on Auto Protect, see Enable or Disable Auto Protect.
5. In the **Protection Group** section, do one of the following:
   - Select **New Group** to create a new Protection Group. In the **Protection Name** field, enter a name of the Protection Group.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
6. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, Create or Use an Existing Protection Policy.
7. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
8. After selecting a Storage Domain, if you need to change any of the advanced settings on the **Microsoft 365 - Teams** page, scroll down and click **Edit** on the right. For more information, see Advanced Settings.
9. Click **Protect**.

   Your Microsoft 365 Teams Data Protection is now active and running.

Advanced Settings

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the protection run should start. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
End Time	**Optional**. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Indexing	By default, indexing is enabled. **Note:** Only the channels of the Microsoft 365 Teams are indexed.
Alerts and Priority	**Optional**. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.

Advance Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	**Optional**. Enter a brief description about the Protection Group.

Advance Settings	Description	
Exclusions	**Teams Meeting Recordings** exclusions	Teams meeting recordings saved the default location can be excluded from backup. This exclusion is enabled by default for new protections. For existing protections, you need to edit the protection, select the **Teams Meeting Recordings** checkbox, and then save the changes. Recordings will be excluded starting with the next backup run. Backup and recovery of Teams meeting recordings metadata is not supported. **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature. **Note: Teams meeting recording**: A video or audio recording of a Teams meeting. It is stored in OneDrive for non-channel meetings, or in SharePoint for channel meetings, and lets participants review the meeting later.

Create or Use an Existing Protection Policy

Create a Protection Policy with the settings listed in the Policy Settings for Microsoft 365 Teams. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Microsoft 365 Teams

The following table lists the applicable policy settings for Microsoft 365 Teams:

Field	Description
DataLock	Add a DataLock for compliance and regulatory requirements. This ensures that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users from modifying or deleting any Snapshots that were generated by the protection runs that use this Policy. The users include those who have the Data Security role in Cohesity as well. > **Note:** Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Scheduled Backup	Defines when backups are captured, how long they are retained, and the type of backup to capture.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer duration than defined by the protection schedule.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, Cohesity attempts to capture Snapshots three times before the protection run fails. The default time between retries is thirty minutes. You can customize the number of retries and how long to wait between each attempt. > **Note:** Set **Retry Options** to **1**, as the Microsoft 365 connector already has the retry mechanism handling the failed requests. There is no need to increase the value for this setting in the Protection Policy.
Quiet Times	Quiet Times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet time period.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with applicable policy settings for Microsoft 365 Teams. In this case, all the policy settings apart from the settings applicable to Microsoft 365 Teams are ignored.

To use an existing Protection Policy:

1. Navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for Microsoft 365 Teams. For information on the applicable policy settings, see Policy Settings for Microsoft 365 Teams.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage Teams Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered Microsoft 365 domain and the assigned Storage Domain in the Protection Group settings.

Add Microsoft 365 Teams to the Protection Group

You can add a new Microsoft 365 Teams to the Protection Group for backup. The newly added Microsoft 365 Teams are backed up from the next protection run, and it will be a full backup for the recently added user.

To add Microsoft 365 Teams to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Teams** page, click ✎ .
5. From the **Teams** list, select the Microsoft 365 Teams that you plan to add to the backup and click **Save Selection**.
6. Click **Save**.

Remove Microsoft 365 Teams from the Protection Group

You can remove the Microsoft 365 Teams from the Protection Group that you do not want to back up in the next protection run. When you remove a site from the protection group, the site's incremental backup is not performed from the future protection run. The existing backup for the site is preserved until the snapshot expires.

To remove Microsoft 365 Teams from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Teams** page, click ✎ .
5. From the **Teams** list, unselect the Microsoft 365 Teams that you plan to remove from the backup and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the new sites that are added to Microsoft 365.

When the Auto Protect option is enabled:

- All the Teams that are added to Microsoft 365 in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the Teams that are removed from Microsoft 365 in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backups for the user are preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly onboarded Teams on Microsoft 365 to the Protection Group.
- Remove existing Teams that are offboarded on Microsoft 365 from the Protection Group. The existing backups of the users are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Microsoft 365 - Teams** page, click ✎ .
5. From the **Groups** list, click the A icon in the domain level to enable or disable Auto Protect and click **Save Selection**.
6. Click **Protect**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to

your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for OneDrive.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up Microsoft 365 Teams. It provides you with options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Microsoft 365 Teams.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. In the **Backup** drop-down list, select any one of the following:
        - Backup all Objects in the Protection Group
        - Backup only selected Objects in the Protection Group
    2. In the **Backup Type** drop-down list, select the backup type.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Microsoft 365 Teams.

3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Download Error Report for a Protection Run

You can now download a list of errors encountered during the Teams protection run. Once the protection run is finished and if any files are skipped or have encountered errors, the protection run status is shown as **Warning**. You can download a report that lists the files that are skipped or have encountered errors.

> **Note:** The files skipped or encountered errors are re-tried for the backup in the following protection run.

To download the report:

1. Navigate to **Data Protection** > **Protection**.
2. On the **Protection** page, select a OneDrive Protection Group with an icon in the **Status** column.

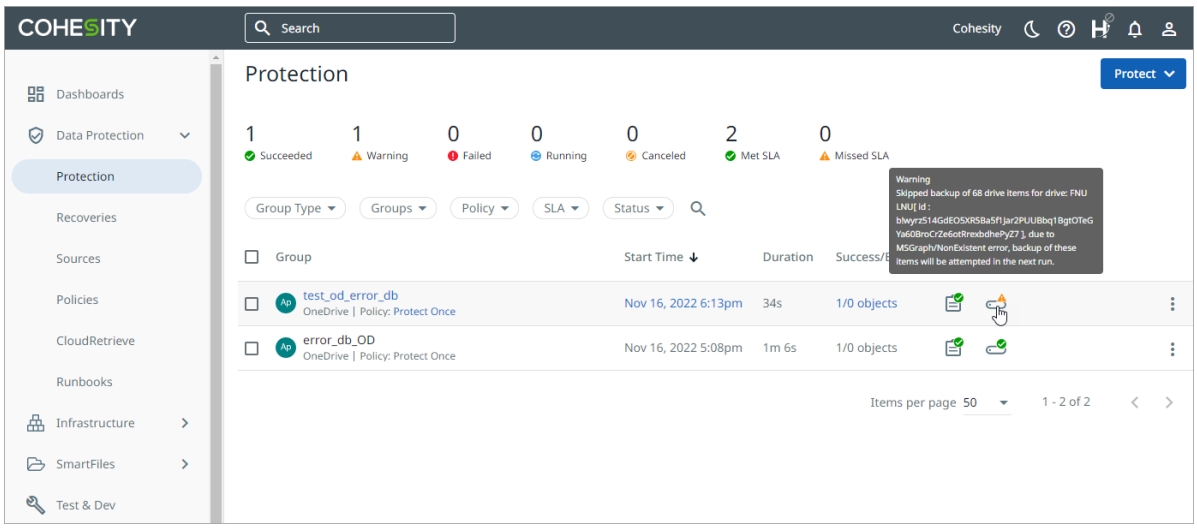

3. In the **Group Details** page, select the protection run with the icon in the **Status** column.

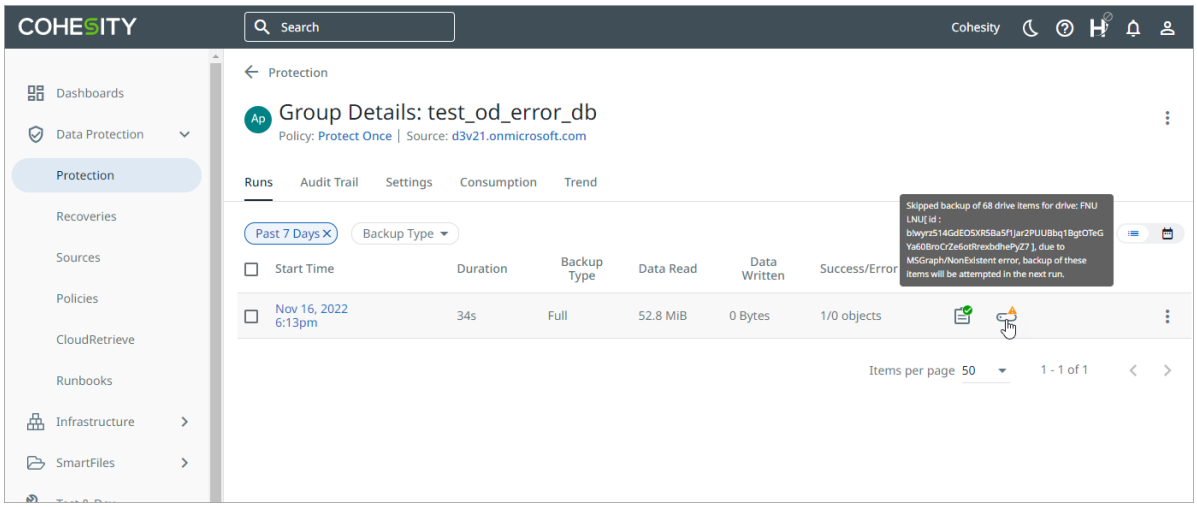

4. On the **Run Details** page, select **view more** to view the list of errors for the protection run and click **Download Error Logs**.

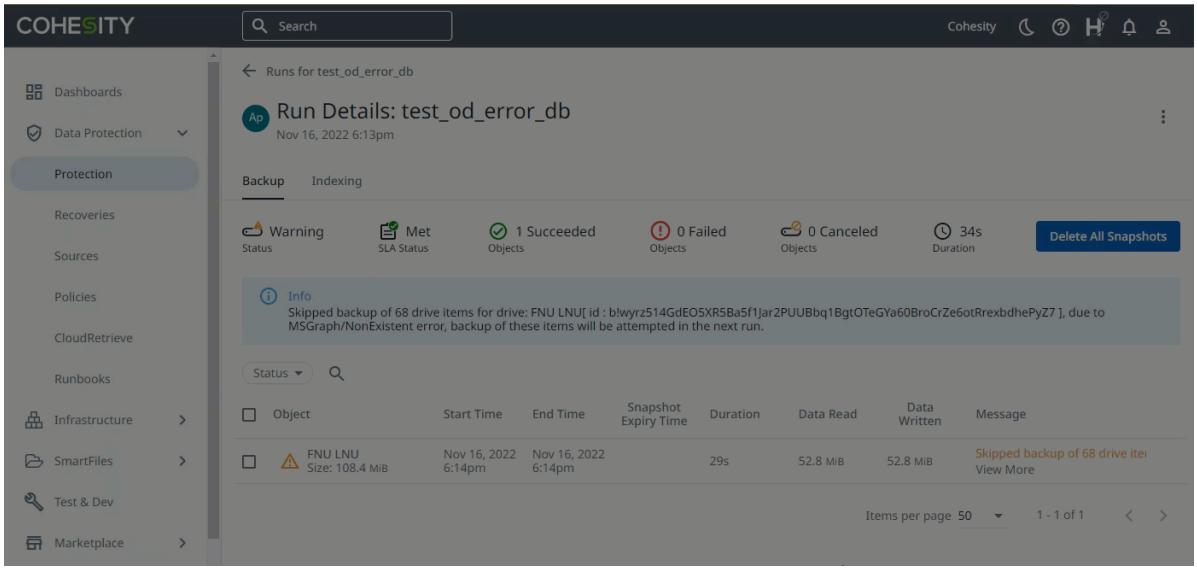

This will download all the error logs in a CSV file format.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

> **Note:** After deleting the Protection Group and Snapshots, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Cluster.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Microsoft 365 Teams.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Protection Group Only
   - Delete Protection Group and Snapshots

Recover Teams

Cohesity enables you to restore Team data from the backup. You can restore a Team's data to the original or different Microsoft 365 domain. Additionally, you can restore multiple Teams to the original Microsoft 365 domain.

Recover Team

To restore a team, you can search through the backups by team name or using a string pattern in the team name that you plan to restore and then start the restore operation.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To restore a team:

1. Navigate to **Data Protection** > **Recoveries**. Next, click **Recover** on the top-right of the page and select **Microsoft 365** > **Teams** as the recovery source.
2. On the **New Recovery Microsoft 365 - Teams** page, select **Entire Team**.
3. Search and select the team you need, and then click **Next: Recover Options**.

> **Note:** Cohesity displays the Microsoft 365 site title (Site Name) registered during source discovery. This title corresponds to the site display name in Microsoft 365 and may not match the site name segment in the SharePoint URL. For example, a site whose URL ends in /sites/ProjectAlpha may appear in Cohesity under a different display title configured in Microsoft 365. When locating objects, use the Site Name shown in the Cohesity UI rather than assuming it matches the URL path.

> **Note:** You can use * character to list all the teams.

4. Under **Recover To**, select any one of the following:

Options	Description
Original Location	To restore to the original Microsoft 365 domain. **Note:** In case the team is deleted from the Microsoft 365 domain after the backup. Then, Cohesity creates a team with the same name and recovers the team's data to the newly created team.
New Location	To restore to a different team on the original Microsoft 365 domain or a different Microsoft 365 domain. For detailed instructions, see Sample Recovery Use Cases.

5. Under **Recovery Options**, perform the following:
   - By default, the **Restore Original Owner Members** option is enabled. Disable the option if you do not want to restore the original owners and channel members to the newly created team.

     > **Note:** The **Restore Original Owner Members** option is not applicable if you are restoring the Team's data to a different Microsoft 365 domain.

   - From the **Additional Team Owner** drop-down list, select the target team owner that needs to be added to the original team owners.

- Enable the **Continue on Error** option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled, and the recovery operation will fail if one of the objects encounters an error.
- In the **Task Name** field, enter the name of the recovery task if you want to change the default recovery task name.

6. Click **Recover** to start the restore operation.

Sample Recovery Use Cases

Use Case	Instructions
Restore the Team's data to the original Microsoft 365 domain.	In the **Recover To** section, select **Original Location**.
Restore a Team's data to the original Microsoft 365 domain and add a new owner.	Perform the following: 1. In the **Recover To** section, select **Original Location**. 2. Under the **Recovery Options** section, disable **Restore Original Owner Members** and select an **Additional Team Owner** from the **Team Owner** drop-down list.
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain.	Perform the following: 1. In the **Recover To** section, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source** and select the **Target Team** to which you plan to restore the data.
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain. **Note:** If you are restoring to a different Team, the channels are restored with suffix: `<_coh><some_int_value>`. Where `_coh` is appended to the channel name and `some_int_value` is a random integer value.	Perform the following: 1. In the **Recover To** section, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source** and select the **Target Team** to which you plan to restore the data.

Use Case	Instructions
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain and add a new owner.	Perform the following: 1. In the **Recover To** section, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source** and select the **Target Team** to which you plan to restore the data. 3. Under the **Recovery Options** section, select an **Additional Team Owner** from the **Team Owner** drop-down list.
Restore the Team's data to a new Team on the original Microsoft 365 domain. **Note:** If you are restoring to a new Team, the channels are restored with suffix: `<_coh><some_int_value>`. Where `_coh` is appended to the channel name and `some_int_value` is a random integer value.	Perform the following: 1. In the **Recover To** section, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source.** 3. In the **Target Team** drop-down, click **Create New Team** and specify the **Team Name**.
Restore the Team's data to a new Team on the original Microsoft 365 domain and add a new owner.	Perform the following: 1. In the **Recover To** section, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source**. 3. In the **Target Team** drop-down, click **Create New Team** and specify the **Team Name**. 4. Under the **Recovery Options** section, select an **Additional Team Owner** from the **Team Owner** drop-down list.
Restore the Team's data to a different Team that exists on a different Microsoft 365 domain.	Perform the following: 1. In the **Recover To** section, select **New Location**. 2. Select the target Microsoft 365 domain from the **Registered Source** drop-down. 3. Select the **Target Team** to which you plan to restore the data.

Use Case	Instructions
Restore the Team's data to a new Team on a different Microsoft 365 domain. Also, add a new owner.	Perform the following: 1. In the **Recover To** section, select **New Location**. 2. Select the target Microsoft 365 domain from the **Registered Source** drop-down. 3. In the **Target Team** drop-down, click **Create New Team** and specify the **Team Name**. 4. Under the **Recovery Options** section, select an **Additional Team Owner** from the **Team Owner** drop-down list.

Recover Multiple Teams

You can restore multiple Teams to the original Microsoft 365 domain. To restore teams, you can search through the backups by team name or using a string pattern in the team name that you plan to restore and then start the restore operation.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To restore multiple Teams:

1. Navigate to **Protection** > **Recoveries**. Next, click **Recover** on the top-right of the page and select **Microsoft 365** > **Teams** as the recovery source.
2. On the **New Recovery Microsoft 365** - **Teams** page, select **Entire Team**.
3. Search and select the teams you need, and then click **Next: Recover Options**.

   > **Note:** You use * character to list all the teams.

4. Under **Recover To**, by default, the **Original Location** is selected.

   > **Note:** In case the team is deleted from the Microsoft 365 domain after the backup. Then, Cohesity creates a team with the same name and recovers the team's data to the newly created team on the Microsoft 365 domain.

5. Under **Recovery Options**, perform the following:
   - By default, the **Restore Original Owner Members** option is enabled. Disable the option if you do not want to restore the original owners and channel members to the newly created team.

- From the **Additional Team Owner** drop-down list, select the target team owner that needs to be added to the original team owners.
- Enable the **Continue on Error** option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled, and the recovery operation will fail if one of the objects encounters an error.
- In the **Task Name** field, enter the name of the recovery task if you want to change the default recovery task name.

6. Click **Start Recovery** to start the restore operation.

Search and Restore Teams Files

You can restore Teams files to the original or different Microsoft 365 domain. To restore Teams files, you can search through the backups by folder name or file name, or use a string pattern in the folder name or file name that you plan to restore and then start the restore operation.

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To search and restore Teams files:

1. In Cohesity DataPlatform, navigate to **Data Protection** > **Recoveries**. Next, click **Recover** on the top-right of the page and select **Microsoft 365** > **Teams** as the recovery source.
2. On the **New Recovery Microsoft 365 - Teams** page, select **Search Files and Recover**.
3. Use the '*****' wildcard character and/or enter the text (such as '***.xlsx**' or '***.jpg**) to search for the folders or files with a matching folder name or file name in the **Search** bar. Select the folders or files to recover from the search results.

   Or

   Click **Advanced Search** and select **Both**, **Files**, or **Folder** and search based on the available filters and click **Apply**.
4. To use a different **Recovery Point** for a folder or file, click the **Edit** icon on the tile for that folder or file. Find the recovery point you need and click **Select Recovery Point**.
5. Click **Next: Recover Options**.
6. Under **Recover To**, select any one of the following:

Options	Description
Original Location	To restore to the original Microsoft 365 domain. > **Note:** In case the team is deleted from the Microsoft 365 domain after the backup. Then, Cohesity creates a team with the same name and recovers the team's data to the newly created team.
New Location	To restore to a different team on the original Microsoft 365 domain or a different Microsoft 365 domain. Select a **Registered Source** and the **Target Team** or click **Create New Team** to create a new Team. Then, select the **Target Channel** or click **Create New Channel** to create a new channel in the selected Team. For sample recovery use cases, see Sample Teams Content Recovery Use Cases. > **Note:** The files are recovered to a folder in the Target Channel. The folder name will be `Recover-M365-Teams-<dateTime>`, for example, `Recover-M365-Teams-2022-01-19T06_45_29Z`. The folder contains a team folder with the team name, and the team folder includes channel folders with channel names, and the recovered files or folders are stored in these channel folders.
Download	To download the Teams content. A new recovery task is created to download the content. When the task is completed, from the Activity page, click the task name and then click Download Files to download the generated zip file. > **Note:** When you download multiple files from different snapshots, multiple recovery jobs are run and the files are downloaded separately.

7. Under **Recovery Options**, perform the following:
   - By default, the **Restore Original Owner Members** option is enabled. Disable the option if you do not want to restore the original owners and channel members to the newly created team.
   - From the **Additional Team Owner** drop-down list, select the target team owner that needs to be added to the original team owners.
   - Enable the **Continue on Error** option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled, and the recovery operation will fail if one of the objects encounters an error.

- In the **Task Name** field, enter the name of the recovery task if you want to change the default recovery task name.

8. Click **Recover** to start the restore operation.

Recover Teams Shared Channels

Before you proceed with recovery, you must review the considerations. For more information, see Considerations.

To recover Teams shared channels

1. Navigate to **Data Protection** > **Recoveries**.
2. Next, click **Recover** on the top-right of the page and select **Microsoft 365** > **Teams** as the recovery source. On the **New Recovery Microsoft 365 - Teams** page, select **Teams Content**.
3. From the **Teams Content** drop-down, select **Channels**.
4. Search for the required shared channel.
5. For full recovery, you can select the shared channel you want to recover, or for granular recoveyr, you can select the individual item you want to recover.
6. Click **Next: Recover Options**.
7. Under **Recover To**, select any one of the following:
   - **Original Location**

     To recover to the original location. If the team is deleted from the Microsoft 365 domain after the backup, Cohesity creates a team with the same name and recovers the team's data to the newly created team.
   - **New Location:**

     To recover to a different shared channel in the original Microsoft 365 domain or a different Microsoft 365 domain. From the **Registered Source** drop-down, select the required tenant. You can also initiate registering a new tenant by clicking the **Registered Source** option.

     From **Target Teams**, select the required team where you want to recover the selected channel/item. From **Target Channel**, select the required channel, or create a new one.

     To create a new channel, click **Create New**, click **Shared Channel**, enter a name for the new channel, enter the name of the owner, and then click **Create**.

     > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

8. Under **Recovery Options**, do the following:

- By default, the **Restore Original Owner Members** option is enabled. Disable this option if you do not want to recover the original owners and channel members to the newly created team.
- The **Restore Original Owner Members** option is not applicable if you are recovering the Team's data to a different Microsoft 365 domain.
- From the **Additional Team Owner** drop-down list, select the target team owner that needs to be added to the original team owners.
- Enable the **Continue on Error** option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled, and the recovery operation fails if one of the objects encounters an error.
- In the **Task Name** field, enter the name of the recovery task if you want to change the default recovery task name.

9. Click **Recover** to start the recovery operation.

Sample Teams Content Recovery Use Cases

Use Case	Instructions
Restore the Team's data to the original Microsoft 365 domain.	Under **Recover To**, select **Original Location**.
Restore a Team's data to the original Microsoft 365 domain and add a new owner.	Perform the following: 1. Under **Recover To**, select **Original Location**. 2. In **Recovery Options**, disable **Restore Original Owner Members** and select an **Additional Team Owner** from the **Team Owner** drop-down list.
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source** and select the **Target Team** to which you plan to restore the data. 3. Select the **Target Channel** or click **Create New Channel** to create a new target channel to which you plan to restore the channel data.

Use Case	Instructions
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain. Also, restore the channel data to an existing channel.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source** and select the **Target Team** to which you plan to restore the data. 3. Select the **Target Channel** to which you plan to restore the channel data.
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain. Also, restore the channel data to a new channel.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source** and select the **Target Team** to which you plan to restore the data. 3. From the **Target Channel** drop-down, click **Create New Channel** to create a new channel in the selected Team.
Restore the Team's data to a different Team that exists on the original Microsoft 365 domain and add a new owner.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source** and select the **Target Team** to which you plan to restore the data. 3. Select the **Target Channel** or click **Create New Channel** to create a new target channel to which you plan to restore the channel data. 4. Under **Recovery Options**, select an **Additional Team Owner** from the **Team Owner** drop-down list.

Use Case	Instructions
Restore the Team's data to a new Team on the original Microsoft 365 domain.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source**, click **Create New Team** and specify the Team Name. 3. Select the **Target Channel** or click **Create New Channel** to create a new target channel to which you plan to restore the channel data.
Restore the Team's data to a new Team on the original Microsoft 365 domain and add a new owner.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the original Microsoft 365 domain as the **Registered Source**, click **Create New Team** and specify the **Team Name**. 3. Select the **Target Channel** or click **Create New Channel** to create a new target channel to which you plan to restore the channel data. 4. Under **Recovery Options**, select an **Additional Team Owner** from the **Team Owner** drop-down list.
Restore the Team's data to a different Team that exists on a different Microsoft 365 domain.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the target Microsoft 365 domain from the **Registered Source** drop-down. 3. Select the **Target Team** to which you plan to restore the data. 4. Select the **Target Channel** or click **Create New Channel** to create a new target channel to which you plan to restore the channel data.

Use Case	Instructions
Restore the Team's data to a new Team on a different Microsoft 365 domain. Also, add a new owner.	Perform the following: 1. Under **Recover To**, select **New Location**. 2. Select the target Microsoft 365 domain from the **Registered Source** drop-down. 3. Click **Create New Team** and specify the **Team Name**. 4. From the **Target Channel** drop-down, click **Create New Channel** to create a new channel in the selected Team. 5. Under **Recovery Options**, select an **Additional Team Owner** from the **Team Owner** drop-down list.

Teams Metadata Restore Scenarios

The following table lists the restore behavior for Teams metadata in various Teams data restore scenarios:

Teams Metadata Type	Restore Scenario	Restore Behavior
Team Name	Restore to the original location and the team is not deleted in the source.	Restored
	Restore to the original location and the team is deleted in the source.	Restored
	Alternate restore to an existing Team.	Not restored
	Alternate restore to a new Team.	Not restored
Team Description	Restore to the original location and the team is not deleted in the source.	Restored
	Restore to the original location and the team is deleted in the source.	Restored
	Alternate restore to an existing Team.	Not restored
	Alternate restore to a new Team.	Restored

Teams Metadata Type	Restore Scenario	Restore Behavior
Team Settings	Restore to the original location and the team is not deleted in the source.	Restored
	Restore to the original location and the team is deleted in the source.	Restored
	Alternate restore to an existing Team.	Not restored
	Alternate restore to a new Team.	Restored
Team Nickname	Restore to the original location and the team is not deleted in the source.	Restored
	Restore to the original location and the team is deleted in the source.	Restored
	Alternate restore to an existing Team.	Not restored
	Alternate restore to a new Team.	Restored
Team Privacy Settings	Restore to the original location and the team is not deleted in the source.	Restored
	Restore to the original location and the team is deleted in the source.	Restored
	Alternate restore to an existing Team.	Not restored
	Alternate restore to a new Team.	Restored

Download 1:1 Chats and Group Chats

To download 1:1 chats and group chats:

1. Navigate to **Recoveries**.
2. Click **Recover** at the right-left of page.
3. Click **Microsoft 365** > **Mailboxes**.
4. On the **Microsoft 365 – Mailbox** page, click **Entire Mailbox**.
5. Set the filter criteria as required.
6. Select the users for whose chat you want to restore.
7. Click **Next: Recover Option**.
8. Click **Download Private Chats**.
9. Click **Recover**.
   A separate download job will run for each user. You will be redirected to **Recoveries** page to monitor the progress. Once the download job completes successfully, click the job to view the subtasks and also click the **Download** button at the top-right to

browse to the downloaded chat file. The system creates a separate file for each user and includes all their chats in that file.

Download Channels Posts

To download channel posts:

1. Navigate to **Recoveries**.
2. Click **Recover** at the right-left of page.
3. Click **Microsoft 365** > **Teams**.
4. On the recovery page, click **Teams**.
5. Select the Team that you want to restore.
6. Click **Next: Recover Option**.
7. Click **Download Posts**.
8. Click **Recover**.

Download Specific Teams Channel

To download specific Teams channel:

1. Navigate to **Recoveries**.
2. Click **Recover** at the right-left of page.
3. Click **Microsoft 365** > **Teams**.
4. On the recovery page, click **Team Content**.
5. From the **Teams Content** drop-down list, select **Channels**.
6. Set the filter criteria as required.
7. Select the Team that you want to restore.
8. Click **Next: Recover Option**.
9. Click **Download Posts**.
10. Rename the task name for your convenience.
11. Click **Recover**.
    You will be redirected to **Recoveries** page to monitor the progress.

Replicate Microsoft 365 Data

Cohesity enables you to replicate Microsoft 365 data to another Cohesity cluster for disaster recovery use cases. For more details about replication, see Replication and Remote Access Setup.

Cohesity provides a policy-based replication solution for Microsoft 365 data. In the Protection Policy, you can specify the remote Cohesity cluster, define conditions for

replication, specify the retention period of the data on the remote Cohesity cluster, and other parameters.

To replicate Microsoft 365 data to another Cohesity cluster:

1. Configure replication to another Cohesity cluster. For more details, see Replication and Remote Access Setup.
2. Edit the existing Protection Policy that you had created for backup or create a new Protection Policy for backup. Configure the replication settings as per your requirement. For more information, see Create or Edit a Standard Policy.
3. Run the Protection Group that you have created for Microsoft Exchange Online or OneDrive based on your requirement or the replication settings will take effect in the next scheduled Protection Group run.

Chapter 16: Physical Servers

After installing the Cohesity Agent on a physical Server and registering it with the Cohesity cluster, the Cohesity cluster can use a block-based or file-based Protection Group to back up the physical Server. Protection Groups can run repetitively on a schedule.

To learn more about the relationship between Protection Groups and Policies, see Manage Protection Policy. For instructions on how to back up Physical Servers, see Manage Protection Groups.

For example, you could create a block-based Protection Group with a protection schedule that backs up all the critical Servers every hour as represented by the *Critical* Protection Group in the following figure.

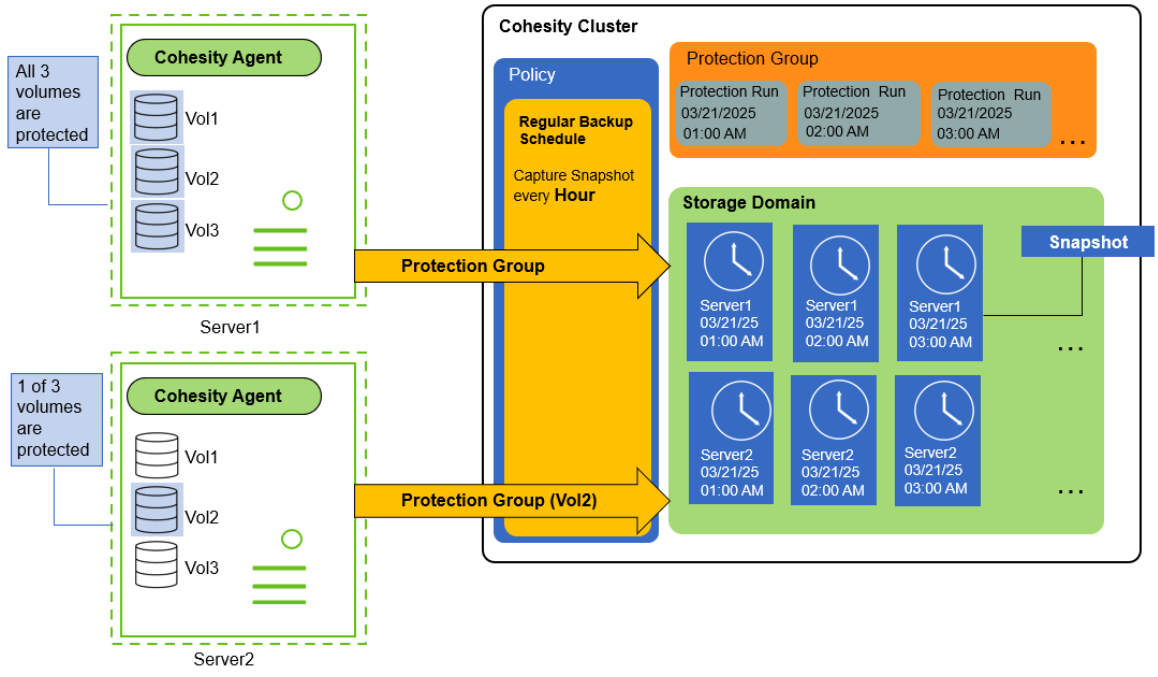

During a Protection Group Run, Snapshots of the Servers are saved on the Cohesity cluster with a date and time stamp.

For instructions on how to back up Physical Servers, see Add or Edit a Protection Group for Physical Servers.

Plan and Prepare for Physical Server Protection

This topic provides information on minimum permissions, prerequisites and considerations to protect physical servers with the Cohesity cluster.

Minimum Permissions

Ensure you have the required minimum permissions to register physical servers as a source with the Cohesity, and to perform backup and recovery. For more information, see Ensure Adequate Privileges for Cohesity on the Source.

Minimum Agent Requirements

Operating System	Minimum CPU	Minimum memory	Free Disk Space	Software Requirement
Linux	1 core	300 MB	400 MB	nfs-utils, rsync, lsof, wget
Debian/Ubuntu	1 core	300 MB	400 MB	nfs-utils, rsync, lsof, wget
SLES/openSUSE	1 core	300 MB	400 MB	nfs-utils, rsync, lsof, wget, libcap-progs
RHEL/CentOS/OEL	1 core	300 MB	400 MB	nfs-utils, rsync, lsof, wget
Solaris	0.5 core	1.5 GB	100 MB	Java 1.8, sudo
AIX	0.25 core	1 GB	100 MB	Java 1.8, sudo
HP-UX	1 core	1 GB	100 MB	Java 1.8, sudo

Support Matrix

Before you register your physical servers with the Cohesity cluster, ensure that Cohesity supports the physical server versions. For more information, see Supported Software

Ports Used for Communication

Ensure all necessary ports are open in the firewall between the Cohesity cluster and physical servers. For more information, see Manage Firewall Ports.

Prerequisites and Considerations

Physical Server Protection

Prerequisites

- The Cohesity Agent running on the hosts should be Cohesity version 6.2 onwards.
- If physical servers are behind NAT, add the NAT server's IP address to the Cohesity cluster's subnets specified in the allowlist.

Physical Server - Generic Considerations

- Making a disk offline and then online triggers a full backup for Simple, Striped and Spanned volumes. Mirrored volumes are not affected.
- Recovering a file or instantly mounting volumes from an AD-joined Cohesity cluster to a server that does not belong to that AD domain is not supported.
- Instant volume mounting does not currently support using drive letters on the target physical Server. Volumes are under C:\CohesityMounts\... for Windows and /opt/cohesity/agent/mount_paths/vol_mounts/….. for Linux systems.
- You cannot add different types of physical servers (Windows, Linux) in a single Protection Group as the supported features of Cohesity Data Cloud are different for different physical servers.
- Cohesity supports recovery only from physical server to physical server. Recovering physical server to VM and vice-versa is not supported.
- Cohesity supports recovering files and folders from Linux to Linux, AIX to AIX, and Solaris to Solaris only. Recovering files and folders from one operating system to another is not supported. For example, recovering files and folders from Linux to AIX or Solaris is not supported.
- After creating a directive file-based Protection Group for a physical server, you can simultaneously trigger more than one protection run of the Protection Group to backup only a specific set of files and folders in each run. However, you cannot simultaneously trigger multiple protection runs to back up the same file or folder.
- When browsing files or folders for recovery, Cohesity does not display files or folders with an absolute path length exceeding approximately 3,500 characters. To recover these files or folders, select the parent directory.
- Due to variations across operating systems and file systems, sparse files may not be detected accurately. In such scenarios, they’ll be backed up as regular files, with

relevant updates in pulse logs.

- When sparse files are detected and processed during backups and recoveries, the amount of data processed during backups and recoveries will reflect the non-zero content. This is likely to differ from previous backups or recoveries where sparse files were processed as normal files.

Windows - Considerations

- When performing a recovery for a Windows physical server, if deduplication functionality is enabled for the backed-up data (files, folders, volumes, or mount volumes), the target machine should have deduplication functionality enabled by default, or the recovery jobs would fail.
- Recovery of deduplicated volume from archival is not supported for Windows physical servers.
- For file and folder recoveries to Windows machines, if the machine is apart of an AD Domain then the Cohesity Cluster must be joined to the same AD domain. The only alternative is to disable **SMB Security** in **Cohesity Views Global Settings** and allow Guest login on the Windows machine.
- Windows file-based backup does not support the following:
  - Backup and recovery of encrypted files (EFS) and alternate data streams (ADS).
  - Backup of symbolic links within a symbolic link pointing to the NAS target.
- A subsequent Full (No CBT) backup would also occur after a machine restart unless it is an intentional restart of a Windows physical server with the Capture Incremental After Server Restart option toggled on.
- You cannot add different types of physical servers (Windows, Linux) in a single Protection Group as the supported features of Cohesity Data Cloud are different for different physical servers.
- Source-side deduplication is not supported in Windows 2008 R2 servers.
- By default, VSS writers are not used for Windows file-based protection. This can be controlled using the crash-consistent backup settings. For more information, see Protecting a Physical Server (File-based).
- To avoid antivirus scans on the Cohesity folder saved on the `C:` drive, place it under the exclusion of the antivirus software.
- For Windows block-based backups, there are additional consideration, see Windows Block-Based iSCSI FC Volume Backups.

Linux

Prerequisites

- To use the Linux PowerPC, Linux on z/Architecture or Linux on ARM (aarch64) agent, the following prerequisites are required:

- Java Runtime Environment (JRE) 8+ (JRE 8, 11, 17)
- Minimum free space of 102400KB is required in /var and /usr directories
- Minimum memory of 2GB is recommended

Considerations

- For block-based backups, the Cohesity Linux agent uses LVM snapshots to back up volumes. LVM does not support any consistency across volumes, so snapshots of different volumes, even in the same backup run, will have slightly different point in time snapshot timestamps.
- Linux block-based backup does not support CBT (Changed Block Tracking). Therefore, every incremental backup is a full backup.
- For block-based backup of the Linux server, Cohesity does not support thin-pool configurations, which allows snapshots to be created outside the volume group.
- For block-based backup of the Linux server, Cohesity supports snapshot creation only inside the same volume group. If there is no space within the same volume group, block-based backup of the Linux server will fail.
- For SLES 11 SP4, the Linux agent is supported for the root user only.
- Linux PowerPC supports only file-based backups.
- Linux PowerPC supports only RPM-based installation.
- Linux on PowerPC, z/Architecture, and ARM (aarch64) supports only file-based backups.
- Linux on PowerPC, z/Architecture, and ARM (aarch64) supports only RPM-based installation.

AIX

Prerequisites

- To use the AIX agent (Java based), Java Runtime Environment (JRE) 8 needs to be installed on the source system.
- AIX agent (Java based) being installed should be a higher version than the existing version of GoAgent.
- Install 64-bit version 8 of IBM-supported AIX compatible Java.
- 1.5 GB of free RAM (at run time) to account for JRE overhead.
- Sudo package must be installed.
- Latest version of Bash must be installed. For more information, see AIX Toolbox for Linux Applications.
- Java Agent must be installed as the Root user.

Considerations

- Only file-based Protection Groups are supported.
- The mount point used for Oracle DB cloning is not persistent across reboots. To resolve this issue, manually update the **/etc/filesystem** file using any editor and update the mount path information.
- IVM (Instant Volume Mount) is not supported.
- The AIX physical server does not support the following:
  - Source-side deduplication
  - AIX clustering
  - AIX version earlier than 6.1
  - Power HA
  - JFS2-based snapshots
  - Veritas Cluster Server
  - BMR
  - Workload partitions (WPARs).

Red Hat High Availability (Pacemaker) Cluster

Considerations

- Pacemaker version 2.1 or higher with PCS is supported.
- Only resources of class ocf, provider heartbeat, type IpAddr2, and Filesystem are supported.
- The Cohesity agent must be installed on all nodes belonging to the cluster. If any additional node is added to the cluster at a later stage, the agent should be installed on it as well.
- It is expected that Filesystem resources are active on the same node where the IpAddr2 resource is running. Cohesity itself assumes no coupling between IpAddr2 and Filesystem resources.
- Backup and recovery are supported only if the agent is installed as the root user.
- If failover happens during the backup run, there may be a few files missed in the run. Explicit failover detection and handling can be enabled by contacting Cohesity Support.

Veritas Cluster Server (VCS)

Prerequisites

- One node in the VCS cluster must be active and the other nodes must be passive.
- Each node in the VCS cluster must have its own service group and each service requires access to all the shared volumes.

- The same version of the Cohesity Linux Agent must be installed on all the nodes in the VCS cluster.
- The following commands must exist on all the nodes in the VCS cluster:
  - whereis hares
  - whereis hagrp
  - whereis hasys
  - whereis haclus
  - whereis lltconfig

  If the command output returns the path of the command, the command is present on the node.

Considerations

- Only file-based backup and recovery of the VCS cluster running on a physical server is supported.
- Only active-passive VCS cluster configuration is supported (active-active configuration is not supported).
- If you are upgrading Cohesity Linux Agents of the VCS cluster from the Cohesity cluster using the VIP, then only the Agent on the node to which the VIP is pointing will be upgraded. Other nodes can be upgraded by one of the following methods:
  - Add individual nodes to the Cohesity cluster as a physical server and upgrade the Cohesity Linux Agent.
  - Do failover of VIP to other nodes and then upgrade the Cohesity Linux Agent.
  - Install fresh Agents on other nodes.
- To protect a VCS cluster shared disk, use VIPs to register the VCS cluster. To protect local disks, register the VCS nodes with their Private IPs.
- If a failover happens during backup, then the backup of some files might be skipped. These skipped files will be backed up in the next Protection Group run.
- If the failover takes more than 10 minutes, then the Protection Group run might fail.
- If a failover happens during the restore, the maximum delay permitted is 10 minutes. If the failover is more than 10 minutes, the restore will fail.
- If you mount a volume to a VIP using the Instant Volume Mount capability, and perform a failover followed by a teardown, then the teardown might fail, and the node's VIP will be orphaned.

Solaris

Prerequisites

- Java Agent must be deployed on the server.
- Physical source should be installed with Java version 8.

Considerations

- Block-based, BMR, Source side dedupe and volume recovery are not supported for Solaris.
- Solaris Zones in Solaris x86_64 are not supported.

HP-UX

Considerations

- Cohesity supports only file-based backup and restores on HP-UX servers.
- Cohesity does not support block-based, BMR, and source side dedupe backups on HP-UX servers.

Register or Edit a Physical Server

Before registering a physical server with the Cohesity cluster, you must download and install the Cohesity agent on each server you plan to register.

Overview

1. Download the Cohesity agent to the physical server.
2. Install the agent on the physical server.

   Instructions for silent installation on Windows servers are available here: Silently Install the Cohesity Agent on Windows Physical Servers.
3. Register the physical server with the Cohesity cluster.

Download and Install the Agent

Installing the agent on a physical server allows it to be registered as a source with the Cohesity cluster.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**.
3. Follow the installation instructions for the corresponding agent operating system:

   Windows Agent

   1. Ensure the agent has been downloaded to the server you want to protect.
   2. As an administrator with local system privileges on that server, run the executable with **Run as administrator** and complete the installation wizard.

      Install the **Volume CBT (Changed Block Tracker)** component to perform incremental backups. Installing this component requires a reboot. Without a

reboot, only full backups can be performed. Not installing this component means only full backups can be performed.

> **Note:**
> The Cohesity Volume CBT driver creates the **Cohesity_Bitmap.cty** file at the root of each volume. There are multiple ways to backup Windows physical server but the **Cohesity_Bitmap.cty** file will be available on the system only if the Volume CBT driver is installed. This file is used to maintain the CBT data for incremental block-based backups across Windows server reboots.
>
> Do not move or delete the "**.cty**" file.

The **File System CBT (Changed Block Tracker)** component is required for MS SQL file-based protection.

3. The agent starts automatically.

> **Note:** This installation does not require a restart. You can wait until a normal maintenance window (such as a Patch Tuesday) to restart. However, until you do, some Protection Groups that depend on Volume CBT will run in full backup mode.

Linux Agent

The Cohesity Linux agent is now available with different installer packages, providing support on multiple Linux distributions. The following table lists the installer packages and Linux distributions on which the installer package is supported.

Installer package	Linux distribution
(Default) RPM	RHEL and its derivative
SUSE RPM	SUSE
Script Installer	All supported Linux operating systems

Linux agent has dependencies on the following packages, which must be installed on the Linux server:

Command/Package	Linux distribution				
	RHEL	SUSE	CentOS	Ubuntu	Debian
rsync	rsync	rsync	rsync	rsync	rsync
mount	mount	mount	mount	mount	mount
lsof	lsof	lsof	lsof	lsof	lsof
umount	umount	umount	umount	umount	umount
cp	cp	cp	cp	cp	cp
chown	chown	chown	chown	chown	chown
chmod	chmod	chmod	chmod	chmod	chmod
mkdir	mkdir	mkdir	mkdir	mkdir	mkdir
rm	rm	rm	rm	rm	rm
tee	tee	tee	tee	tee	tee
hostname	hostname	hostname	hostname	hostname	hostname
stat	stat	stat	stat	stat	stat
blkid	blkid	blkid	blkid	blkid	blkid
ls	ls	ls	ls	ls	ls
losetup	losetup	losetup	losetup	losetup	losetup
dmsetup	dmsetup	dmsetup	dmsetup	dmsetup	dmsetup
timeout	timeout	timeout	timeout	timeout	timeout
lvs	lvs	lvs	lvs	lvs	lvs
vgs	vgs	vgs	vgs	vgs	vgs

Command/Package	Linux distribution				
	RHEL	SUSE	CentOS	Ubuntu	Debian
lvcreate	lvcreate	lvcreate	lvcreate	lvcreate	lvcreate
lvremove	lvremove	lvremove	lvremove	lvremove	lvremove
lvchange	lvchange	lvchange	lvchange	lvchange	lvchange
libpcap-progs	null	libpcap-progs	null	null	null
nfs-utils **Note:** Required for Instant Volume Mount, file-folder recovery from block-based backup and VMware backup.	nfs-utils	nfs-utils	nfs-utils	nfs-utils	nfs-utils
wget	wget	wget	wget	wget	wget

Installing the Linux agent as an LDAP user is supported.

Install RPM, Debian or SUSE RPM Installer Package

The following table explains the action to be performed for installing the RPM, Debian or SUSE installer package:

> **Note:** In the following commands, replace the Cohesity agent version with the desired agent version.

Action	How to
Download the Agent installer	From the Download Agents window, based on your Linux distribution, select **RPM**, **Debian**, or **SUSE RPM** and download it to the server you want to protect.
Navigate to the downloaded directory	As the root user with local system privileges on that server, change the directory to the location of the installer package.
Install the Custom certificates	**LINUX** If you are using a Cohesity certificate then run the following commands to install the user certificate followed by agent installation: `export ENFORCE_USE_CUSTOM_CERTS=true && export AGENT_CERT_FILE=<path-to-agent-certificate-in-Cohesity-format>`
Install the Agent	Run the following command depending on the installer package: • RPM: `rpm -i el-cohesity-agent-7.4.1-1.x86_64.rpm` or `yum localinstall ./el-cohesityagent-7.4.1-1.x86_64.rpm` • Debian: `dpkg -i cohesity-agent_7.4.1-1_amd64.deb` • SUSE RPM: `rpm -i cohesity-agent-7.4.1-1.x86_64.rpm`
Install the agent as root user	Run the executable file with the following command syntax: `/<path_to_installer_file> -- --install -c 0 -S root -G root` Followings are the command options in details: 1. **-S** this is the user to run the agent, specify root. 2. **-G** this is the group permission the Agent will use for files and directories installed by the agent, specify root. 3. **-c** this boolean switch controls if the OS user and group should be created. 0 means do not create the OS user and group, 1 means the agent installation will create the specified OS user and group. (When choose to run with the root user, specify '-c 0' as 'root' already exists)

Action	How to

Install the agent as non-root user	To start the service as a non-root user, create a new user or use an existing user with sudo permission and run the following command:
	• RPM: `export COHESITYUSER=<username> ; rpm -i el-cohesity-agent-7.4.1-1.x86_64`
	• Debian: `COHESITYUSER=<username> dpkg -i cohesity-agent_7.4.1-1_amd64`
	• SUSE RPM: `export COHESITYUSER=<username> rpm -i cohesity-agent-7.4.1-1.x86_64`
	Where, *<username>* is the user name with sudo permission.
	Note:
	By default, the installation uses the root user permission for all the files, and the service is started as root. Therefore, it is required to add non-root users to the sudoers list by making the following changes in the `/etc/sudoers` file:
	`<user_name> ALL=(ALL) NOPASSWD:ALL`
	`Defaults:<user_name> !requiretty`
	An example of the sudoers file is as given below:
	dgoble ALL=(ALL) NOPASSWD:ALL
	Defaults:dgoble !requiretty
	cohbackup ALL=(ALL) NOPASSWD:ALL
	Defaults:cohbackup !requiretty

Action	How to

	Note:
	You can add sudo privileges for non-root users using the syntax given below:
	`<user_name> ALL=(ALL) NOPASSWD:/path/to/command1, /path/to/command2, /path/to/command3`
	Remove `NOPASSWD` from the command above if you do not want to authenticate every sudo commands with password.
	The example below shows the sudo commands for user cohesityagent:
	`cohesityagent ALL=NOPASSWD: /bin/mount,/bin/umount,/bin/chmod,/bin/chown,/bin/hostname,/bin/mkdir,/bin/rm,/usr/bin/tee,/usr/bin/stat,/bin/cp,/bin/sed,/usr/bin/touch,/sbin/blkid,/usr/bin/lsof,/bin/ls,/sbin/losetup,/sbin/dmsetup,/usr/bin/rsync,/usr/bin/timeout,/usr/bin/getfacl,/usr/bin/setfacl, /tmp/installer_ch_cohesity_linux_agent,/home/cohesityagent/installer,/home/cohesityagent/cohesityagent/data/agent/tmp/installer_ch_cohesity_linux_agent,/bin/systemctl,/home/cohesityagent/cohesityagent/software/crux/bin/cohesity_linux_agent.sh`
	The actual paths for each command may vary in each system. Cohesity keeps adding or removing commands across releases.

Location	• Installation directory: `/opt/cohesity` • Log file: `/var/log/cohesity`

Install script installer package

The following table explains the action to be performed for installing the script installer package:

> **Note:** In the following commands, replace the Cohesity agent version with the desired agent version.

Action	How to
Download the Agent installer	From the Download Agents window, based on your Linux distribution, select **Script Installer** and download it to the server you want to protect.
Navigate to the downloaded directory	As the root user with local system privileges on that server, change the directory to the location of the installer package. For SLES 11 SP4, installing the agent as the root user is required.
Make the installer executable	Make the installer executable, for example: `chmod +x cohesity_agent_7.4.1_linux_x64_installer`
Install the Agent	Run the executable: `sudo /cohesity_agent_6.6.0d_u2_linux_x64_installer - --install`
Location	• Installation directory: `/home/<username>/cohesityagent` or `/root/cohesityagent` • Log file: `/home/cohesityagent/cohesityagent/logs` **Note:** You can specify the option `-log-dir <installtion_path>` to specify the directory path where the agent logs will be saved. If this option is not provided, the logs will be saved in the default path

Install Cohesity Agent for Linux PowerPC

> **Note:** To use the Linux PowerPC agent (Java based), Java Runtime Environment (JRE) 8 needs to be installed on the source system.

The following section explains the action to be performed for installing the Linux PowerPC Agent as a ROOT user:

1. Select **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent**.
3. From the **Download Agents** window, select **PowerPC** and download it to the server you want to protect.
4. Change the directory to the location of the installer package and run the following command to install the rpm based installer package for PowerPC

distribution:

```
cohesity-java-agent-7.0.1-1.ppc64le.rpm
```

The agent will be installed in the following directory `/opt/cohesity` and the logs files can be found in `/var/log/cohesity`.

5. To uninstall the agent, run the following command:

```
rpm -e cohesity-java-agent
```

For information about all Cohesity Linux agent installer options, see Linux Script Installer Options.

The agent starts after installation completes.

Agent startup behavior following a machine reboot:

- CentOS and RedHat (distributions with the "systemd" init system) - The agent starts automatically.
- Ubuntu (distributions with the "upstart" init system) - The agent starts automatically.

If a Linux server's `/etc/sudoers` file is managed by a deployment engine such as Chef, Puppet or others, this could impact Cohesity cluster's interaction with servers that have the Linux agent installed. Take the corresponding actions below:

Agent Installation - User type	Action required

As default cohesityagent user	The Cohesity Linux agent is installed using the cohesityagent user by default. For default installations, the cohesityagent is created by the installer. During installation, the installer updates the `/etc/sudoers` file to allow cohesityagent sudo and no-tty sudo access. Ensure the following settings in the /etc/sudoers file for the cohesityagent user are preserved:
	`cohesityagent ALL=(ALL) NOPASSWD:ALL`
	`Defaults:cohesityagent !requiretty`
	An example of the sudoers file is as given below:
	#includedir /etc/sudoers.d
	dgoble ALL=(ALL) NOPASSWD:ALL
	cohbackup ALL=(ALL) NOPASSWD:ALL
	Defaults:cohbackup !requiretty

As non default user, example foo	Ensure the above settings in the `/etc/sudoers` file for the foo user are preserved by replacing the occurrences of cohesityagent with foo.

Agent Installation - User type	Action required
As root user	No changes required

Install Cohesity Agent for Linux on z/Architecture

> **Note:** To use the Linux on z/Architecture agent (Java-based), Java Runtime Environment (JRE) 8+ (JRE 8, 11, 17) needs to be installed on the source system.

The following section explains the action to be performed for installing the Linux on z/Architecture Agent as a ROOT user:

1. Select **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent**.
3. From the **Download Agents** window, select **s390x** and download it to the server you want to protect.
4. Change the directory to the location of the installer package and run the following command to install the rpm based installer package for z/Architecture distribution:

   ```
   rpm -i cohesity-java-agent-0.0.0-1.s390x.rpm
   ```

   The agent will be installed in the following directory `/opt/cohesity` and the logs files can be found in `/var/log/cohesity`.
5. To uninstall the agent, run the following command:

   ```
   rpm -e cohesity-java-agent
   ```

Install Cohesity Agent for Linux on on ARM (aarch64)

> **Note:** To use the Linux on on ARM (aarch64) agent (Java-based), Java Runtime Environment (JRE) 8+ (JRE 8, 11, 17) needs to be installed on the source system.

The following section explains the action to be performed for installing the Linux on on ARM (aarch64) Agent as a ROOT user:

1. Select **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent**.
3. From the **Download Agents** window, select **aarch64** and download it to the server you want to protect.
4. Change the directory to the location of the installer package and run the following command to install the rpm based installer package for z/Architecture distribution:

   ```
   rpm -i cohesity-java-agent-7.3.1-1.aarch64.rpm
   ```

   The agent will be installed in the following directory `/opt/cohesity` and the logs files can be found in `/var/log/cohesity`.
5. To uninstall the agent, run the following command:

   ```
   rpm -e cohesity-java-agent
   ```

cohesity_linux_agent.sh

The table lists how `cohesity_linux_agent.sh` interacts with various `init` systems to issue `start/stop/status` commands.

If Cohesity does not detect any of the `init` systems, the installer uses the **undetected** `init` system commands.

init system	Commands

systemd	`sudo systemctl start cohesity-agent`
	`sudo systemctl stop cohesity-agent`
	`sudo systemctl status cohesity-agent`

upstart	`sudo initctl start cohesity-agent`
	`sudo initctl stop cohesity-agent`
	`sudo initctl status cohesity-agent`

SysV	`sudo /etc/init.d/cohesity-agent start`
	`sudo /etc/init.d/cohesity-agent stop`
	`sudo /etc/init.d/cohesity-agent status`

undetected	`sudo /etc/init.d/cohesity-agent start`
	`sudo /etc/init.d/cohesity-agent stop`
	`sudo /etc/init.d/cohesity-agent status`

AIX Agent

For Cohesity support information, see Supported Software.

1. Download the Cohesity Java Agent package (a **.bff** file) from Cohesity Package Download. Ensure the Java Agent is on the AIX server you want to protect.
2. Log in to the AIX server as a root user and run the following command to install the package:

   ```
   installp -ad <path_to_downloaded_agent_package> all
   ```

   For example:

   ```
   installp -ad cohesity_agent_6.6_aix_java.bff all
   ```

   The Cohesity Java Agent starts automatically. Now you can register the server with the Cohesity cluster. Reboot is not required. The Cohesity Java Agent starts after any future reboots.

Solaris Agent

1. Obtain the Cohesity Solaris agent package. Download the Solaris 10 agent here.
2. Run the following command to install the package:

   ```
   sudo /usr/sbin/pkgadd -d <cohesity_agent_file_path>
   ```

   If sudo is not installed, install the package as the root user.

   The agent starts automatically. Now you can register the server with the Cohesity cluster. Reboot is not required. The agent starts after any future reboots.

HP-UX Agent

You can download the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI. For more information on installing HP-UX agent, see Install and Manage the Agent on HP-UX Servers.

4. Open the following ports on the server:
   - TCP port 50051 for backup operations (incoming)
   - NFS ports 111 and 2049 for restore operations on a Linux server (outgoing)
   - SMB port 445 for restore operations on a Windows server (outgoing)

Register a Physical Server with the Cohesity cluster

Registering a physical server makes its content eligible to be protected or backed up on the Cohesity cluster.

> **Note:** Selecting an Interface Group is not relevant when registering a physical server, despite the option being present on the Physical Server registration page on the Cohesity cluster.

1. Select **Data Protection > Sources**.

   Click **Register > Physical Server.**

   (*Applies only in multi-tenant environment with Data Source Connection enabled.*) From the **Data Source Connection** drop-down, select a Data Source Connection, and then click **Continue**.
2. Specify the hostname (FQDN) or server IP (IPv4 or IPv6) of the server to register. Cohesity recommends using the FQDN.
3. Click **Register**.

   When registration finishes, a new entry appears in the physical section on the Sources page.

> **Note:** You can register a Windows physical server on multiple Cohesity cluster. In a Cohesity cluster failover scenario, when the primary Cohesity cluster goes down, you can use the secondary Cohesity cluster to backup the file/folder on the physical server. However, you cannot register a Windows physical server as an MS SQL server on multiple Cohesity cluster. For more information, Contact Cohesity Support.

Register Veritas Cluster Server (VCS) on Cohesity cluster

To protect VCS, you must first register VCS as a physical server on Cohesity cluster using VCS VIP.

> **Note:**
>
> - You can register a VCS VIP on multiple Cohesity clusters. For more information, Contact Cohesity Support.
>
>   This is an Early Access feature. Contact Cohesity Support to enable the feature.

Red Hat High Availability (Pacemaker) Cluster

To protect the Red Hat High Availability (Pacemaker) Cluster, you must first register it as a physical server on the Cohesity cluster using the VIP address.

Manage Physical Servers

Re-register Cohesity Agent (Windows or Linux)

The re-register functionality on the Cohesity cluster allows you to save the agent settings on the cluster and enables you to perform an agent re-registration with the same settings when required.

With re-registration, the existing file-based protection job would continue to work with the new agent without any intervention. However, the block-based protection job may need to be modified if specific volumes under the job have changed before re-registration.

> **Note:**
> - For block-based protection jobs, the first backup after re-registration will be a full backup.
> - Re-register functionality is currently supported for Windows and Linux servers.
> - During agent re-registration, you must use one of these options, **Apply settings for the agent preserved on Cohesity UI** or **Preserve setting during uninstallation of an agent from the server**.

1. Select **Data Protection** > **Sources**.
2. Navigate in the source hierarchy to the Windows or Linux server you want to re-register, mouse over its actions menu, and click **Re-register**.
3. Click **Apply the Settings for the agent preserved on the <date and time>** checkbox if you want to restore the server settings.

   If this checkbox is not selected, then the settings saved on the cluster are not applied to the newly registered agent.

   > **Note:** During agent uninstallation of a Cohesity Windows or Linux agent, there is an option to preserve settings in the Windows registry or config file for Linux. If you had selected the **Preserve Settings** option during the uninstall process then deselect the above option during the reregistration

> workflow.

4. Click **Re-register**.

Upgrade Agent on a Physical Server

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

When a new version of the Cohesity Agent is available, you can upgrade agents directly from the Source Details page. The upgrade option is enabled only when an update is available for the selected server.

You can choose to:

- Upgrade multiple agents in bulk
- Upgrade an individual agent

If an upgrade is initiated while backup jobs or recovery tasks are in progress, the upgrade starts automatically after the active operations complete.

To upgrade an individual agent

1. Navigate to **Data Protection** > **Sources** and click the physical server source name.
2. In the **Source Details** page, click the **More Options** menu and then select **Upgrade Agent**.

   The Upgrade Agent option is enabled only when a new version of the Agent is available.
3. Click one of the following options:
   a. **Upgrade Now** to upgrade the Agent immediately, then click **Confirm**.
   b. **Schedule for Later**, in the Schedule Agent Upgrade dialog, set the Date & Time for the upgrade and then click **Schedule**.

To perform a bulk agent upgrade

1. Navigate to **Data Protection** > **Sources**.
2. If agent upgrades are available, click the **Upgrade Agents** button in the table header adjacent to Physical Sources.
3. Click **Physical Server**, select the agents you want to upgrade and then click **Continue**.
4. Click **Schedule** to schedule the agent upgrades. Options include:

a. **Upgrade Now** to upgrade the Agent immediately.

b. **Schedule for later**, in the **Schedule** dialog, set the Date & Time for the upgrade.

5. Click **Upgrade**.

Edit Physical Server Registration

1. Navigate to **Data Protection > Sources**.
2. Navigate the source hierarchy to the server, mouse over its actions menu and click **Edit**.
3. Edit the desired information.
4. Click **Update Server**.

Unregister Physical Server from the Cluster

1. Navigate to **Data Protection > Sources**.
2. Navigate the source hierarchy to the server, mouse over its actions menu and click **Unregister**.
3. Click **Unregister**.

Network Bandwidth and CPU Throttling for Physical Server

You can throttle the network bandwidth and CPU on the Windows and Linux physical servers registered with the Cohesity cluster. Network throttling optimizes the data transfer rate from the server to the Cohesity cluster when a Protection Group is run. CPU throttling optimizes the CPU capacity utilized when a Protection Group with source-side deduplication is run on the server.

Supported Servers

The following table lists the physical servers supporting networking bandwidth and CPU throttling.

Operating System	Network Bandwidth Throttling	CPU Throttling
Windows	Yes	Yes*
Linux	Yes	No

* CPU throttling is supported on Windows 2012 R2 and later versions.

> **Note:** The network bandwidth and CPU throttling do not apply to the CBMR.

Enable Network Bandwidth and CPU Throttling

To enable network and CPU throttling, you must first register the physical server as a source and then edit the registered source to enable the options for network bandwidth and CPU throttling:

1. Select **Data Protection > Sources**.
2. Navigate the hierarchy to the registered physical server source for which you want to enable network bandwidth and CPU throttling.

   > **Note:** The edit server page appears.

3. Enable the **Throttle Network Bandwidth** option and in the **Throttle to** field, specify the maximum data transfer rate from the server to the Cohesity cluster.

   To schedule network bandwidth throttling for specific days and time durations:

   1. Click **Add**.
   2. From the **On** field, select the days during which you want to optimize the network bandwidth.
   3. In the **Start Time** and **End Time** duration fields, specify the time duration for throttling.
   4. In the **Throttle to** field, specify the maximum data transfer rate from the server to the Cohesity cluster.
   5. To configure multiple schedules for throttling, click **+** and repeat steps b to d.

      Example: If you want to limit the data transfer rate to **100 MB/S** on **Mondays** and **Wednesdays** from **12.00 to 16.00**, then from the **On** field, select **M** and **W**, set the **Start Time** as **12.00**, **End Time** as **16.00**, and in the **Throttle to** field, enter **100** as the throttling rate.

      > **Note:** For the days and duration for which you have not scheduled the network bandwidth throttling, Cohesity will take the throttling value specified at the server level.

4. Enable the **Throttle CPU** option and in the **Throttle to** field, specify the maximum CPU capacity (in percentage) of the server that Cohesity cluster backup runs with source-side deduplication enabled can utilize.

> **Note:** CPU throttling is only applied when the protection run has source-side deduplication enabled.

5. Click **Save**.

Add or Edit a Protection Group for Physical Servers

In Cohesity Data Cloud, Protection Groups use Protection Policies. Protection Policies reflect the backup and archival frequency, and retention requirements for each Protection Run. A Protection Group defines operational requirements, such as which source objects to protect, the Protection Policy to use, and operational considerations like indexing, exclusions and inclusions, and more.

A Protection Group that protects physical servers captures snapshots which are required for the following tasks:

- Recover files
- Instant volume mount
- Replicate snapshots
- Archive snapshots

> **Tip:**
> You can also watch the following video to know more about Protection Groups:
>
> - Data Protection: Protection Groups

Supported Workflows

For more information about the supported workflows on protection groups for physical servers (either Block-based or File-based), refer to Physical Server Workflows.

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit a Physical Server
- Create or Edit a Standard Policy
- Create or Edit Storage Domains

Compare Block-Based and File-Based Backup Methods

Cohesity offers two distinct approaches for physical server backups: Block-based and File-based. While both methods serve the purpose of safeguarding critical data, understanding their differences is essential for choosing the most suitable backup strategy for your organization's needs.

Cohesity Facet	Cohesity Block-based Physical Server Backups	Cohesity File-based Physical Server Backups
Granularity	Operates at the disk block level, capturing modified or new data blocks.	Operates at the file level, preserving entire files for backup and recovery.
Speed and Performance	Enables faster backup and recovery times by tracking changes at the block level and transferring only modified data blocks.	Offers straightforward backup and recovery operations, relying on the identification and preservation of entire files.
Optimized Storage Efficiency	Significantly reduces storage requirements by capturing only the changed data blocks.	File-based backups rely on deduplication within the storage domain for optimization. It provides file-level accessibility, allowing selective recovery.
Simplicity and Familiarity	Requires familiarity with block-level operations.	Follows traditional file-level backup approaches that administrators are accustomed to.
Flexibility and Accessibility	Focuses on individual data blocks, enabling more granular and incremental backups.	Ensures the accessibility of individual files and directories, simplifying selective recovery.
Metadata and Indexing	Relies on metadata which is available once indexing occurs after the backup is complete.	Includes comprehensive metadata and indexing capabilities for quick search and file location.
Dependencies	Block-based backups depend on the volume CBT driver on Windows and LVM snapshots in Linux.	No dependency in the case of file-based backups.
Supported Configurations	Block-based backups are supported with specific file systems (NTFS for Windows, small group for Linux x64).	File-based backups are less restrictive and can work with a wider range of filesystems and operating systems (AIX etc)
Supported Recoveries	Block-based allows for IVM recovery,	File-based only allows for File-level recovery.
Source Backup	Block-based backup allows only one backup against one Source	File-based backup allows multiple backups to be configured against one Source.

In summary, Cohesity provides two distinct methodologies for physical server backups: Block-based and File-based. While Block-based backups excel in speed, efficiency, and granularity at the disk block level, File-based backups offer simplicity, flexibility, and file-level accessibility. Choosing the appropriate backup method depends on your organization's specific requirements, data protection goals, and recovery objectives.

The decision between Cohesity's Block-based and File-based Physical Server backups should be based on factors such as backup speed, storage efficiency, data granularity, ease of recovery, and the desired level of accessibility.

By understanding the differences between the two backup approaches, organizations can effectively leverage the backup functionality to enhance their data protection strategies to optimize their overall businesses.

To add or edit a Protection Group for Physical Server

Once you have registered your physical server as a source, you're ready to use Cohesity Data Cloud to protect it.

You can either perform Block-based or File-based backups for the physical servers.

- Performing a Block-based Backup
- Performing a File-based Backup

> **Important:** Cohesity recommends not to run backup and restore protection groups in parallel on the same physical server path.

Protect a Physical Server (File-based)

To add or edit a Protection Group for Physical Server:

1. Select **Data Protection > Protection** and do one of the following:
   a. To create a Protection Group— Click **Protect** on the top-right of the page and select **Physical Server** > **File-based**.
   b. To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. In the **New Protection** pop-up, do one of the following:
   a. Click **Add Objects** and select the Servers to protect.

      You can also register a new source with the Cohesity cluster so you can protect it.
   b. Select the physical servers to protect. By default, all Servers are excluded as indicated by the unselected check box. A selected check box (☑) indicates

the Server is selected to be protected.

You can optionally filter the Servers table.

- In the search field, enter part of the Server name.
- Select an option from the **Filter** drop-down:
  - **Unprotected**—Displays Servers that are not protected by any Protection Groups in the Cohesity cluster—not protected by Cohesity.
  - **Protected by Cohesity**—Displays Servers that are protected by any Protection Groups in the Cohesity cluster.
  - **Protected by this Job**—Displays Servers that are currently selected to be backed up by this Protection Group. (This option is available only when editing an existing Job.)

Selecting the check box adjacent to Servers selects all of the filtered servers.

The graphical representations are described below.

- An unchecked box indicates the Object is excluded from any Protection Group.

  A shield indicates the Object is already protected by another Protection Group. A Windows physical Server cannot be protected by a new protection group unless you remove it from the current protection group.

- A checked box indicates the Object will be protected by the current Protection Group.

- Click the icon to see details about the object.
- If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.

c. You can select either Windows, VCS, Solaris, Red Hat High Availability (Pacemaker) Cluster, or AIX physical object. Click the edit icon next to the selected server and use one of the following methods to define the file or a folder path in a file-based protection group for the selected object (Windows, VCS, Solaris, or AIX):

Note:

- When backing up sparse files, consider the following:

  Due to variations across operating systems and file systems, sparse files may not be detected accurately. In such scenarios, they'll be backed up as regular files, with relevant updates in pulse logs.

  When sparse files are detected and processed during backups and recoveries, the amount of data processed during backups and recoveries will reflect the non-zero content. This is likely to differ from previous backups or recoveries where sparse files were processed as normal files.

Options	Descriptions
Follow symlink NAS target (applies only to **Windows file-based** backup)	**Note:** Share user and Cohesity agent login user must be part of the "Domain Admin" group to use option. Enable this option if you want to back up the symbolic link pointing to the NAS target. You can create a symbolic link on the Windows server by running the following command: `mklink /d /<link> /<target>` Where <**link**> specifies the name of the symbolic link that is being created. <**target**> specifies the path (relative or absolute) that the new symbolic link refers to. **Note:** • Cohesity does not support the backup of symbolic links within a symbolic link. • As the symbolic link points to a remote machine, the backup process will be slower than usual. Specify at least one path for each Server you want to protect. By default, all files in the `/ root or c:\` directory are backed up.

Options	Descriptions
All Volumes	Use this option to protect all the local volumes on the specific host. If you select this option, the system dynamically enumerates all the local volumes present during the backup protection group run and protects them. > **Note:** In Windows, local volumes are the volumes listed in the "Disk Management" utility of the Windows host. If you choose to use the Add Exclusions functionality when using this option, you can either specify: • Individual files or folders within a volume by specifying the fully qualified path names • Individual files or folders across all the volumes using the wildcards • Entire volumes by defining the volume path If the exclusion path is not present at the time of a backup run, the system ignores this entry and proceeds to the next. This option overrides any exclusion path specified as global exclude. **Example** If you specify `/mnt/banapps/` as a global exclude path but the local mount points are `/mnt/banapps/ar_techsis_views` and `/mnt/banapps/common`, then all the files and folders inside these local mount points will be included. If you want the Protection Group to exclude all the local mount points inside `/mnt/banapps/`, then you can either add these local mount points as explicit exclude paths or choose a normal backup and then give "/" as the include path, and `/mnt/banapps/` as the exclude path. Add Exclusions You can use this option to add exclusion entries for Individual files and folders and Protect all local volumes. For more information, see Add Exclusions.
Directive File	The directive file defines the locations of the files and folders that you can back up in a protection group. Specify the file path of the directive file (meta-file) on the source while adding it to a protection group. For more information, see Directive File Backup. > **Note:** Applies only to **Windows, AIX, and Linux** servers file-based backup

Options	Descriptions
Specify Directories	**Nested Mount Points to Skip** Based on the Physical server type you want to protect, select the mount points to skip. For Windows Physical Server: Skip Nested Mount Points: Enable this option if you do not want to backup the volumes that are mounted to a subfolder within the selected directory structure. By default, this option is enabled. **For Linux Physical Server**: Nested mount point types to skip: From the drop-down list, select the mount point types that you do not want to back up. By default, the mount point types autofs and nfs on the server are selected. **Inclusions** During File-based Physical Server Backups, all files and directories under the parent directory specified in the Include path of a Protection Group are protected by the Cohesity Cluster. **Add Exclusion** You can use this option to add exclusion entries for Individual files and folders and Protect all local volumes. For more information, see Add Exclusions. **Add Path** Click Add Path to protect an additional path and then click Save. For the Cohesity cluster version below 6.6.0d_u3, after an initial full backup of a cluster, if a parent directory is added, there is a probability of unaltered files going missing in the latest snapshot. This issue is fixed in Cohesity version 6.6.0d_u3 and later. However, Cohesity recommends running at least one full backup after the upgrade. For more information, see Incremental forever backups with parent directory addition.

3. **Protection Group**: To enter a new name for the protection group, select **New Group** and enter the group name in the **Group Name** field. To use an existing protection group name, select **Existing Group** and select an existing group name from the **Group** drop-down list.
4. **Policy**: Select an existing protection policy or create a new one by selecting **New Policy**.
5. **Storage Domain**: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

> **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

6. **Start Time**: Available only if the selected policy is set to Backup Daily. Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.
7. Click **More options** if you want to change or configure any of the additional settings.

Setting	Description
Pause Future Runs	Toggle this option to stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.
End Date	Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying a Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for C6K Platform) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.

Setting	Description
Cancel Runs at Quiet Time Start	Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
Pre and Post Scripts	Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run. For configuration details, see Configure Pre & Post Scripts.
Crash Consistent Backups (applies only to **Windows file-based** backup)	Toggle on this option to read files from the snapshots of volumes on which the files (that need backup) are residing before the Protection Group is executed. Optionally, you can enable **Continue with non-consistent backup in case of failure** option to resume the backup in case the crash consistent backups fail. Crash Consistent backup settings are enabled by default for block-based backup jobs. For file-based backups, this setting is off by default but can be enabled from the Cluster UI. User preferences are taken into account.
Source-Side Deduplication	Toggle on to enable source-side deduplication for all the servers that are part of the Protection Group. Source-side deduplication is not supported on Windows 2008 R2 servers. For more information, see Source-Side Deduplication for Physical Server. Toggle on Source Side Deduplication, click Add Exclusions and select the hostname (IP address) to be excluded from the deduplication protection group. > **Note:** The dedupe parameter must be enabled on the storage domains associated with the Linux and Windows servers.

Setting	Description
Cache Optimization	This feature applies only to AIX, Linux (x86_64), PPC, Solaris (SPARC & x86) and HPUX file-based backups. Cache Optimization is a modified source dedup functionality that offers optimized backup performance on the Cohesity cluster. **Note:** • It is recommended to use this only for new jobs. Cohesity recommends not to switch to Source-side Dedup or normal incremental once you create a job with **Cache Optimization** functionality enabled. You can either select source-side dedup or CacheOptimization for a job. • To use the Cache Optimization functionality, both the Cohesity agent and cluster should have version 7.0 and later.
Indexing	Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recovery task to locate files by name. For information about customizing indexing, see Customize Indexing.

<table>
<tr><th>Setting</th><th>Description</th></tr>
<tr><td>Global Exclude Paths</td><td>
Cohesity version 6.6 and higher versions support global excludes at the protection group level, a standard entry for all servers. This feature applies to all types of Supported operating systems for physical agents.

For a file-based Protection Group, specify the exclusion paths for all the servers that are part of the protection group. By default, all files in the / root or c:\ directory are backed up.

Click +Add to exclude a particular path or a particular file for all the specified sources. An exclusion will be a combination of global exclude and object-level exclude paths specified on each server. For more information on object-level exclusions, see Add Exclusions.

Example

When a user configures a protection group on the Cohesity cluster, they can set up global exclude and include paths to specify which data should be backed up and which should be excluded. For instance,
• If /temp is provided as global exclude for a protection group and /temp/data as include path for a source, then /temp/data will be skipped by the Cohesity cluster.
• If /temp is an include path for a source and /temp/data is provided as a global exclude path for the protection group, then all the data under /temp will be backed up, but the /temp/data folder will be excluded from the backup.

The following table lists the support file systems and default exclusions for the physical server file-based backups.

<table>
<tr><th>OS</th><th>Supported File Systems</th><th>Default Exclusions</th></tr>
<tr><td>Linux</td><td>ext3, ext4, xfs, vxfs, ReiserFS, autofs, cvfs, zfs, fuse.mfs, nfs, nfs3, cifs, gpfs, mmfs, btrfs</td><td rowspan="3">Virtual filesystems are excluded from backups. For example - tmpfs, devtmpfs, CacheFS, procfs, sysfs, etc.</td></tr>
<tr><td>AIX</td><td>jfs2, cdrfs, nfs, nfs3, jfs, vxfs, mmfs, gpfs, autofs</td></tr>
<tr><td>Solaris</td><td>zfs, ufs, nfs, cifs, vxfs</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Setting</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>OS</th><th>Supported File Systems</th><th>Default Exclusions</th></tr>
<tr><td>Windows</td><td>NTFS</td><td>Directories excluded from each physical drive.
For example - C:/, D:/, etc
$Recycle.Bin, Cohesity_Bitmap.cty, System Volume Information, System State</td></tr>
</table>

Note: Windows System State backups are supported as part of Bare Machine Recovery (BMR) backups and not Physical Server File-based backups. For more details, please refer to About Bare Machine Recovery.
</td></tr>
<tr><td>Allow Parallel Runs</td><td>Toggle this option to allow parallel backup runs for the Protection Group.
If the Protection Group is currently executing a protection run and a new run is started (can be manual or scheduled), then the current run continues to back up all the unfinished sources. In parallel, the new runs of the Protection Group will start backups for all the objects or entities that are processed completely in the current protection run.</td></tr>
<tr><td>Alerts</td><td>Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:
Success: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.
For more information, see Alerts.</td></tr>
<tr><td>Priority</td><td>Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with High priority are given the highest priority to execute, Medium the second-highest priority, and Low the lowest priority.</td></tr>
</table>

Setting	Description
SLA	The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster. • Incremental: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run. • Full: Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
Description	Optional. Click the plus sign and enter a description for the Protection Group.
Ignorable Errors	Select this option to ignore the following errors during a file-based protection group run for physical servers. You can select one or both options as per your requirements. **End of File** - If this option is selected, the Cohesity cluster ignores the End of File errors during a file-based protection group run and the file will be backed up successfully. **Non-Existent** - If this option is selected, the Cohesity cluster ignores the Non-Existent errors during a file-based protection group run.
Continue on Error	Enable this option to continue physical backups even if errors are encountered during the backup of any included file in the backup runs. This option is enabled by default. Disabling this option would cause the physical backup job runs to fail if any errors are encountered.

8. Click **Protect**.

   The Protection page is displayed and includes the new Protection Group.

Note:

- If a server is being backed up and files will also be recovered to that same server, the file recover task starts after the backup completes. The converse is also true; if files are being recovered to a server and that same server will be backed up, the backup starts after file recovery completes.
- For a File-Based backup, the "Protected" and "Protected Size" fields on the source page for physical sources are updated after the source refresh.

Protect a Physical Server (Block-based)

To add or edit a Protection Group for Physical Server:

1. Select **Data Protection > Protection** and do one of the following:
   a. To create a Protection Group— Click **Protect** on the top-right of the page and select **Physical Server** > **Block-based**.
   b. To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.
2. In the **New Protection** pop-up, do one of the following:
   a. Click **Add Objects** and select the Servers to protect.

      You can also register a new source with the Cohesity cluster so you can protect it.

      A block-based protection group can protect Windows and non-Windows Servers but not both at once. You have to create separate protection jobs to protect Windows and non-Windows Servers.

      > **Note:** A physical server (block-based) can be protected by only one Protection Group at a time. If you want to protect a Server that is already protected by a protection group, you must remove it from its current Group. If the Server being removed is the protection group's only Server, you can delete the Group and its Snapshots or the protection group only. If you delete a snapshot, it may impact the recovery workflows.

   b. Select the physical servers to protect. By default, all Servers are excluded as indicated by the unselected check box. A selected check box (☑) indicates the Server is selected to be protected.

      You can optionally filter the Servers table.

      - In the search field, enter part of the Server name.
      - Select an option from the **Filter** drop-down:
        - **Unprotected**—Displays Servers that are not protected by any Protection Groups in the Cohesity cluster—not protected by Cohesity.
        - **Protected by Cohesity**—Displays Servers that are protected by any Protection Groups in the Cohesity cluster.
        - **Protected by this Job**—Displays Servers that are currently selected to be backed up by this Protection Group. (This option is available only when editing an existing Job.)

Selecting the check box adjacent to Servers selects all of the filtered servers.

The graphical representations are described below.

- An unchecked box indicates the Object is excluded from any Protection Group.

  A shield indicates the Object is already protected by another Protection Group. A Windows physical Server cannot be protected by a new protection group unless you remove it from the current protection group.

- A checked box indicates the Object will be protected by the current Protection Group.

- Click the icon to see details about the object.
- If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.

c. You can select either Windows or a non-Windows physical object and click the edit icon next to the selected server and perform the following actions:

   i. **Take BMR backup**: Toggle on for bare machine recovery (CBMR) backups of the Server. For more information, see About Bare Machine Recovery.

   ii. **Automatically protect all volumes on this Server**: Toggle off and deselect the volumes you do not want to protect. Even if all volumes are selected for protection, new volumes will not be protected until they are added to the protected list.

   iii. Click **Continue**.

3. **Protection Group**: To enter a new name for the protection group, select **New Group** and enter the group name in the **Group Name** field. To use an existing protection group name, select **Existing Group** and select an existing group name from the **Group** drop-down list.
4. **Policy**: Select an existing protection policy or create a new one by selecting **New Policy**.
5. **Storage Domain**: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

> **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

6. Click More options if you want to change or configure any of the additional settings.
   a. **Start Time**: Available only if the selected policy is set to Backup Daily. Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.
   b. **SLA**: The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
   c. **Incremental**: Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
   d. **Full**: Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
   e. Click **Additional Settings** if you want to change or configure any of the additional settings.

Setting	Description
Pause Future Runs	Toggle this option to stop the protection runs of the Protection Group from executing. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.
End Date	Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.

Setting	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying a Backup HDD, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Pre and Post Scripts	Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run. For configuration details, see Configure Pre & Post Scripts.

Setting	Description
Crash Consistent Backups (applies only to **Windows file-based** backup)	Toggle on this option to read files from the snapshots of volumes on which the files (that need backup) are residing before the Protection Group is executed. Optionally, you can enable **Continue with non-consistent backup in case of failure** option to resume the backup in case the crash consistent backups fail. **Exclude Writers**: All the VSS writers associated with the source are included in the backup job by default. However, you can exclude a VSS Writer from the backup job by toggling on the Exclude Writers option. To exclude a VSS Writer at the job level, You can either enter the name of the Writer you want to exclude or select it from the Exclude drop-down list. Additionally, if you have multiple Windows servers in a backup job, you can set the Exclude Writers option at the server level. To do this, click on the edit icon next to the selected server and toggle on the Exclude Writers option to specify the writers to be excluded from the backup job at the server level. Note that excludes defined at the server level will take precedence over the excludes defined at the job level. > **Note:** This is applicable only for Windows block-based backup jobs.
Incremental After Restart	For Windows physical servers, toggle on if you prefer incremental backup for the first scheduled backup job, which runs after an intentional server restart. If an incremental backup is not possible, a full backup is performed. For example, if the server restarts after a power failure, a full backup is performed. > **Note:** Incremental backup consistency across reboots is not guaranteed if I/O operations are performed on the volume outside of the fully functional state of the server. This includes I/O operations that are performed by booting into the recovery mode. For example, if a server is shut down and its volume is attached to another server that does not have the Cohesity agent, ensure that the next backup is a full backup operation. This option is not available for file-based backups.

Setting	Description
Source-Side Deduplication	Toggle on to enable source-side deduplication for all the servers that are part of the Protection Group. Source-side deduplication is not supported on Windows 2008 R2 servers. For more information, see Source-Side Deduplication for Physical Server. Toggle on Source Side Deduplication, click Add Exclusions and select the hostname (IP address) to be excluded from the deduplication protection group. **Note:** The dedupe parameter must be enabled on the storage domains associated with the Linux and Windows servers.
Indexing	Indexing is enabled by default and required for file recovery. The Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recovery task to locate files by name. For information about customizing indexing, see Customize Indexing.
Cancel Runs at Quiet Time Start	Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with High priority are given the highest priority to execute, Medium the second-highest priority, and Low the lowest priority.

Setting	Description
Description	Optional. Click the plus sign and enter a description for the Protection Group.

7. Click **Protect**.

   The Protection page is displayed and includes the new Protection Group.

Download Physical Server File List for a Protection Group

System-generated file lists are used to provide more details about the backup process. You can now download CSV files for inclusion-exclusion reports and files backed up successfully or with warnings on the Cohesity cluster.

To download the file lists for a protection run:

1. Select **Data Protection** > **Protection**.
2. On the **Protection** page, select the protection group for which you want to download the file lists.
3. On the **Group Details** page, select the backup run and then select the server instance for which you want to download the file list.
4. Click one of the following options:
   - **Download Inclusion-Exclusion Reports**
   - **Download Success Logs**

Download Error Report for a Physical Protection Group

You can now download the list of all the errors encountered during the protection run for a physical server backup. Once the protection run is complete, you can download the details of the files that are skipped or have encountered errors in a CSV format.

Follow the steps below to download the error report for a Protection Group:

1. On the Cohesity Dashboard, select **Data Protection** > **Protection**.
2. On the **Protection** page, select a protection group that has the warning icon (ICON) in the **Status** column.
3. On the **Group Details** page, click the **Warning** icon .
4. On the **Run Details** page, click **View More** in the **Message** column to view the number of files that were skipped due to errors.

This will download all the error logs in a CSV file format.

Manage the Physical Server Backup Configuration

In Cohesity Data Cloud, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the assigned Storage Domain in the Protection Group settings.

Add a Physical Server to the Protection Group

You can add a new Physical Server to the existing Protection Group for backup. The newly added Physical Server is backed up from the next protection run and it will be a full backup.

Perform the following steps to add Physical Server to the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for Physical Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the edit icon.
5. Select the Physical Server you want to add to the Protection Group.
6. Click **Save Selection** > **Save**.

The newly added Physical Servers are backed in the next scheduled protection run.

Remove the Physical Server from the Protection Group

From the Edit Protection page, you can remove the Physical Server from the Protection Group that you do not want to back up in the next protection run.

Perform the following steps to remove the Physical Server from the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for Physical Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** field, click the edit icon.
5. Unselect the Physical Server you want to remove from the Protection Group.
6. Click **Save Selection** > **Save**.

Edit Existing Protection Policy

You can edit the existing Protection Policy that is attached to the Protection Group based on your requirements.

Perform the following steps to edit the policy settings:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for Physical Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Policy** field, click the **Edit Policy** icon .

   The **Edit Policy** page appears
5. Edit the settings and then click **Save**.
6. Click **Save**.

Attach New Protection Policy to the Protection Group

You can attach a new Protection Policy to the Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

Perform the following steps to attach a new protection policy to the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for Physical Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, from the **Policy** drop-down, select a new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can instantly start a protection run of the Protection Group that you have created to back up Physical Servers. It provides you options to back up all the objects or only selected objects in the Protection Group. If the policy associated with the Protection Group:

- Supports multiple backup types, you can select an incremental or full backup.
- Is configured with replication targets, you can select those targets.
- Is configured with archival targets, you can select those targets.

Perform the following steps to start an on-demand protection run:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for Physical Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
   a. From the **Backup** drop-down list, select any one of the following:
      - **All objects in the Protection Group**.
      - **Selected objects in the protection run**.
   b. In the **Backup Type** drop-down list, select the backup type.
   c. Optional. In the **Replicate** section, select the replication target and the retention duration, or click **Add** to add the replication target and the retention duration.
   d. Optional. In the **Archive** section, click **Add** to add the external target to which you want to archive the data and the retention duration. Also, select the **Archive only fully successful runs** checkbox to archive the data only from the completed protection runs.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

Perform the following steps to pause the future protection runs:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for Physical Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered

Perform the following steps to resume paused protection runs:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you paused.

3. Click the actions menu (⋮) next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

You do not want to retain the backups and the local and archived snapshots generated by the Protection Group.

If you want to reclaim the storage space utilized by the Protection Group.

Cohesity Data Cloud provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that this Protection Group captured are not deleted.
- **Delete Protection Group and Snapshots**: Delete the Protection Group and all snapshots, any local snapshots captured by this Protection Group and stored on this cluster are deleted. In addition, any archived snapshots stored on external targets that this Protection Group captured are also deleted.

After deleting the Protection Group, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

Perform the following steps to delete the Protection Group:

1. Select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you created for Physical Server.
3. Click the actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - Delete Group - done by typing 'yes' and clicking on 'Delete Group'
   - Delete Group and Snapshots - done by typing 'yes', selecting the toggle to 'Delete Snapshots', and clicking on 'Delete Group'

   Suppose the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target. In that case, you should cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Run.

   If the Protection Group replicates to a remote cluster, the failover job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job have not been deleted. You can log in to the remote cluster to delete the snapshots and the failover job.

Start On-demand Multiple Protection Runs for Physical Server

After creating a directive file-based Protection Group for a physical server, Cohesity supports to simultaneously trigger more than one protection run of the Protection Group to backup only a specific set of files and folders in each run. This functionality can be achieved using the Cohesity REST API, `/public/protectionJobs/run/{id}`.

To trigger a protection run, add the `metadata file path` parameter in the JSON body of the API and provide the directive file path having the file or folder you plan to backup as the value.

Example to trigger a protection run from using the REST API:

URI

```
https://<cluster_name>/iris-
services/api/v1/public/protectionJobs/run/<protectionGroupID>
```

Method

POST

JSON body

```
{
 "copyRunTargets": [
 ],
 "runNowParameters": [
   {
     "physicalParams": {
       "metadataFilePath": "/home/cohesity/dirfile2"
     },
     "sourceId": 5
   }
 ],
 "runType": "kFull"
}
```

Where **/home/cohesity/dirfile2** is the directive file path that contains a list of files and folders folder that will be backed up using this API.

Similarly, you can simultaneously trigger multiple protection runs of a Protection Group.

Limitation

You cannot simultaneously trigger multiple protection runs to back up the same file or folder.

For example, you cannot start a protection run with the directive file path as **/home/cohesity/dirfile2** if there is already another active protection run of the same Protection Group, which is using the same directive file path.

Source-Side Deduplication and Cache Optimization for Physical Server

Source-side deduplication is a Cohesity feature designed to optimize backup efficiency by eliminating redundant data at the source (physical server) before transferring it to the Cohesity cluster. This process reduces network usage and speeds up backup operations, ensuring efficient data protection for your physical servers.

With source-side deduplication enabled for a physical server during Protection Group creation, unique data blocks that have not been previously transferred are sent to the Cohesity cluster. This significantly minimizes network bandwidth utilization and reduces data transfer times.

Additionally, source-side deduplication integrates seamlessly with Cohesity's built-in deduplication, further optimizing storage by processing only new data blocks.

Advantages of Source-Side Deduplication

- **Reduced Network Traffic**: This feature minimizes the amount of data transferred between physical servers and the Cohesity cluster, conserving network bandwidth.
- **Faster Backups**: Speeds up backup processes, particularly over slower networks such as WANs, enabling quicker recovery point objectives (RPOs).
- **Efficient Resource Utilization**: Offloads deduplication processing from the Cohesity cluster, allowing the system to handle a higher volume of backup workload sessions.
- **Lower Storage Costs**: Reduces overall storage consumption by transmitting only new or changed data blocks.

Cache Optimization

Cache Optimization is a specialized feature that enhances the performance of backups for large files (over 64 GiB) by further refining source-side deduplication functionality. When enabled, only the changed blocks of target files are backed up during subsequent incremental backup runs, rather than the entire file. This ensures efficient storage utilization and reduced backup times.

Requirements

To use Cache Optimization, ensure that both the Cohesity agent and the Cohesity cluster are running version 7.0 or later.

Advantages of Cache Optimization

- **Reduced Network Traffic**: Optimizes data transfers between physical servers and the Cohesity cluster.
- **Lower CPU Usage**: Reduces CPU load on the source server compared to standard source-side deduplication.
- **Improved Performance**: Provides enhanced backup efficiency, particularly for large files, by transferring only changed blocks.

Protect Windows Failover Cluster

With the Failover Clustering feature, users using the Cohesity clusters experience minimum disruptions in service. If one or more nodes of the cluster fail, the backup nodes begin to provide service by a process known as failover.

Currently, Cohesity does not have out-of-box support for backing up Clustered disk resources managed by Windows Failover Cluster (WFC). However, it is still possible to protect Cluster data leveraging existing Physical Protection Jobs using one of the following methods:

Protecting Shared and Local Drives using Block-based Physical Backup

Protecting Shared and local Drives using a File-based Backup (Volume backup)

Protecting Shared and Local Drives using Block-based Physical Backup

To create a block-based protection group for WFC, follow the steps listed below:

1. Install Cohesity Agent on each node of the Windows Failover Cluster and register each Node of WFC as a separate Physical Server. For more information, see Register or Edit a Physical Server.
2. Create a Protection Group (physical server block-based). For more information, see Add or Edit a Protection Group for Physical Servers.
3. Enable **Automatically protect all volumes on this Server** option on every node when creating a protection group.

   If you enable this option, disk changes are automatically detected when a resource failover occurs and required volumes are backed up without requiring any changes to the Protection jobs created.

   > **Note:** If there is no failover, all the backups are incremental (Leveraging CBT) for the Shared disks. If there is a failover/failback on a shared disk over to the other node, the next backup will be a full backup.

Protecting Shared and local Drives using a File-based Backup (Volume backup)

To create a file-based protection group for WFC, follow the steps listed below:

1. Install Cohesity Agent on each node of the Windows Failover Cluster and register each Node of WFC as a separate Physical Server. For more information, see Register or Edit a Physical Server.
2. Create a Protection Group (physical server file-based). For more information, see Add or Edit a Protection Group for Physical Servers.
3. Enable **Protect all local volumes** option on every node when creating a protection group.
4. Provide the volume path in **Exclusions** to exclude any volumes (local drives).

Recovery

The shared drive could be on any of the WFC nodes at the time of backup, You have to search the volume on the backed-up physical servers. Once Volume is found it can be restored to either the original node or any other Physical Server registered with Cohesity. For more information, see Recover Physical Servers.

Protect File Server Role deployed on Windows Server Fail-over Cluster

A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). Windows Server Failover Cluster (WSFC) supports multiple roles, this feature supports the protection of a WSFC File Server role.

Prerequisites

- Make sure the following prerequisites are met before registering a File Server role as a source in the Cohesity cluster:
- You should create shares for the volumes to be discovered as WSFC File Server role volume on the Cohesity cluster.
- The minimum supported version for WSFC Role is the Cohesity version 6.7 and onwards (both Cohesity Windows Agent and Cohesity cluster).
- WSFC is not supported in multi-tenancy and DMaaS installations.
- Only file servers for general use are supported, scale-out file servers also known as active-active are not supported.

Considerations

Review the following considerations before protecting a WSFC File server role using Cohesity:

- Cohesity supports WSFC running on Windows Server 2012, Windows Server 2012 R2, 2016, and 2019 versions.
- WSFC File Server role protection is supported only on on-prem deployments. Data Management as a Service (Backup as a Service) and Multi-tenancy are not supported.
- This feature supports backup and recovery of data associated with the File server role in general use.
- For a crash-consistent backup, the first backup after failover will be full.
- For a crash-consistent backup, the backup will fail if there is a failover during backup.
- If there is no failover, all the backups are incremental (Leveraging CBT) for shared disks. If there is a failover/failback on a shared disk over to the other node, the next backup will be a full backup.
- A standalone physical server protection Group can also protect all volumes associated with the File Server role. However, Cohesity does not recommend protecting volumes associated with the File Server roles through a standalone physical source.
- Bare Metal Restore (BMR) is not supported for WSFC roles. However, BMR can be done on the Windows Servers which are part of the Failover Cluster.
- Upgrading the Cohesity agent on the File Server role registered as a source will also upgrade the agents on all the nodes within the cluster.
- The backups will run through the owner node of the File Server role cluster as the cluster disks are offline and in a reserve state on other nodes.

Protecting File Server role volumes using Block-based Physical Backup

To create a block-based protection group for File Server Roles, follow the steps listed below:

1. Install Cohesity Agent on each node of the Windows Failover Cluster.
2. To protect a share or a volume attached to a File Server role, register **File Server Role VIP/FQDN**. For more information, see Register or Edit a Physical Server.
3. Create a Protection Group (physical server block-based). For more information, see Add or Edit a Protection Group for Physical Servers.

> **Note:**
> - If there is no failover, all the backups are incremental (Leveraging CBT) for the Shared disks. If there is a failover/failback on a shared disk over to the other node, the next backup will be a full backup.
> - You should create shares for the volumes to be discovered as WSFC File Server role volume on the Cohesity cluster.

Protecting FileServer role volumes using a File-based Backup

To create a file-based protection group for File Server Role, follow the steps listed below:

1. Install Cohesity Agent on each node of the Windows Server Failover Cluster and register File Server VIP as a Physical Server. For more information, see Register or Edit a Physical Server.
2. Create a Protection Group (physical server file-based). For more information, see Add or Edit a Protection Group for Physical Servers.
3. Enable **Protect all local volumes** option while creating a protection group if you want to protect all volumes associated with the File Server role.

   > **Note:**
   > - You should create shares for the volumes to be discovered as WSFC File Server role volume on the Cohesity cluster.
   > - For protection with **Crash Consistency** enabled. If there is no failover, all the backups are incremental for the Shared disks. If there is a failover/failback on a shared disk over to the other node, the next backup will be a full backup.
   > - For protection with **Crash Consistency** disabled. All backups are incremental even if there is a failover/failback of shared disks

4. Provide the file path in Exclusions to exclude any files/folders/drives from backup.

Recovery

Once we have backed up the WSFC resources as file-based or block-based, recovery options like “Files or Folders”, and “Instance Volume Mount” can be used to restore data.

For more information, see Recover Physical Servers and Recover Physical Server Files or Folders.

Physical Server Recovery

The physical servers that are backed up on to the Cohesity cluster can be recovered as physical servers, and files and folders. You can also perform bare machine recovery and instant volume mount of the physical server.

Instant Volume Mount for Physical Server

Instant volume mounting is supported for virtual and physical environments. This feature allows the selected backup volumes to be made available at the target location where you can then complete the desired operations. Instant Volume Mount is available for backed-up volumes, whether they are stored locally on the Cohesity cluster or archived to a registered external target.Use cases include granular recovery of Microsoft Exchange, SQL and SharePoint data.

You can mount the volume to its original location or to a new location. For more information, see:

- Instant Volume Mount to the Original Location for Physical Server
- Instant Volume Mount to a New Location for Physical Server

The following table lists the sources and the guest operating systems on which instant volume mounting is supported.

Source	Support for IVM on guest OS	
	Windows	Linux
AHV	No	No
VMware	Yes	Yes
HyperV	Yes	No
Physical*	Yes	Yes

*Supported only for block-based backups.

Prerequisite:

The default path of the recovered file or instantly mounted volume is *<Cluster_IP_address>/<Cloned view name>\<PATH>*.

To access this location, both the target and the Cohesity cluster must be in the same AD domain.

Considerations:

- When recovering a file or instantly mounting a volume from a Windows VM or Physical Server Backup Source that has Windows deduplication installed and enabled for one or more volumes, you must choose a target machine that also has Windows deduplication installed (it does not have to be enabled for any volume). (However, this rule does not apply to Nutanix AHV VMs. If AHV VMs are enabled with Windows deduplication, the only supported recovery option is full VM recovery.)

  If the target does not have Windows deduplication installed:

  - File level recovery will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."
  - Instant volume mounting will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."

  To determine if Windows deduplication is installed on the Source or target machine, follow the steps given below:

  1. Open **Server Manager**.
  2. Select **Roles and Features > File and Storage Services > File and iSCSI Services**.
  3. Select the **Data Deduplication** check box, if necessary.
  4. Click **Next** until the **Install** button is enabled and then click **Install**.

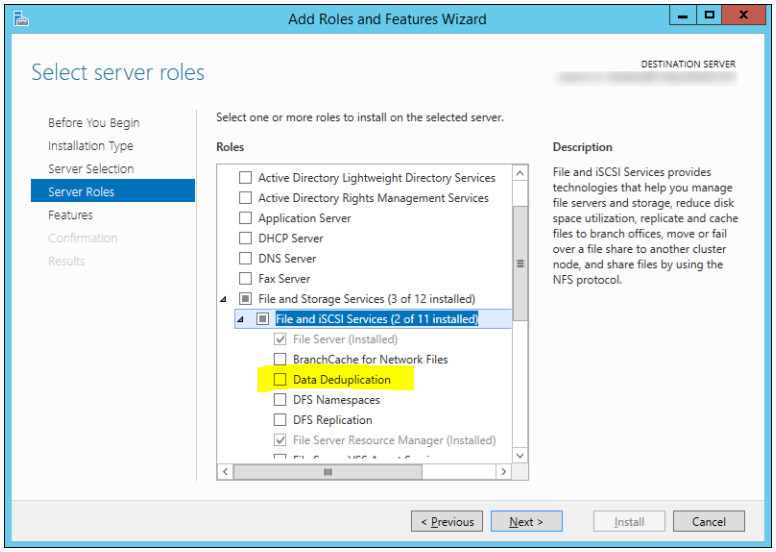

- Instant Volume Mount (IVM) restore of ReFS volumes backed up using Windows physical block-based jobs cannot be restored to an alternate Windows server running a lower version of ReFS.
- When mounting volumes on a Linux physical Server, the loop devices present on the Server are used for mounting. Therefore, the number of volumes that can be

mounted depends on free loop device availability. By default, the number of available loop devices is 8, but this number can be customized. If the number of configured loop devices is the default of 8, up to eight volumes can be mounted. In this example, if an attempt to mount more than 8 volumes occurs, the mounting of all the volumes after the 8th volume fails and errors are reported.

- Tearing down a cloned database or instant volume mount deletes the mounted volumes. Any new or modified data on these volumes will be deleted along with the volumes, so ensure you back up any important data before teardown.
- You cannot instantly mount a volume from a VM to a physical server, and vice-versa.
- Replicating snapshots is supported on the receiving Cohesity cluster when block-based physical backups were replicated. Register the mount-target physical server on that cluster and ensure network and agent prerequisites are met.

Before you can instantly mount a volume, a Snapshot of the Object you want to recover must exist.

Troubleshooting

This section provides troubleshooting information for issues that might occur during IVM recovery.

Authentication Failures

Issue

The IVM recovery fails due to an authentication failure and displays the following error:

"The specified network password is not correct".

Possible Root Cause

The cluster joined the AD but the Windows host has not joined the AD.

Troubleshooting steps

- Create a local user on the Cohesity cluster and use the local user to access the SMB share. To create a local user on the Cohesity cluster, contact Cohesity Support.

or

- Disable SMB Authentication on the Cohesity Data Cloud. Perform the following steps to disable SMB Authentication:
  1. Log in to the Cohesity Data Cloud.
  2. Select **File Services** > **Views** and click the **Global Settings** tab.
  3. Toggle off the **SMB Authentication** setting.

Instant Volume Mount to the Original Location for Physical Server

To perform instant volume mounting to the original server:

1. Select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Physical Server> Instant Volume Mount**.

   The Instant Volume Mount page appears.
3. Search for the server that contains the data you want to mount by entering characters of the server name or by specifying the wildcard character *.

   A list of servers displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, server type, or date.
4. From the list of search results, select the server that contains the data you want to mount.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <server_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

     b. Select the snapshot to be recovered by clicking the corresponding blue dot.

     c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     d. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the **List** mode:

a. Select the **List** tab from the top-right corner of the Edit recovery point for <server_name> page.

   By default, the page displays the snapshots available in the last one month.

b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

c. Select the snapshots to be recovered.

d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Select Volume(s)** option, perform one of the following actions:
   - Enable **Select all Volumes** to recover all the volumes of the selected server
   - From the **Select one or more options** drop-down list, select the individual volumes you want to recover.

   After volumes are mounted, they may be displayed as different volumes than their source volume. For example, D: volume may be mounted as X: volume.
8. From the **Recover To** options, select **Original Location** to recover the volume to the original location.
9. (Applies only for VMs) With the **Ensure Disks are Online** option enabled, select one of the following options to deploy Cohesity Agent on the target VM.

   > **Note:** For Linux target VMs, you cannot disable this option. For Window target VMs, you can disable this option.

   - **Auto Deploy Cohesity Agent:** Select this option and provide the username and password of the target VM to automatically deploy the Cohesity Agent on the target VM.
   - **Use Existing Cohesity Agent**: Select this option to use the existing Cohesity Agent deployed on the target VM

     The following table provides information on how the volume is mounted when the **Ensure Disks are Online** option is enabled and disabled.

Environment	Ensure Disks are Online option is Turned On	Ensure Disks are Online option is Turned Off
VMware	After the VMDKs are attached to the VM, Cohesity cluster explicitly mounts the volume. To mount the volume to the VM, Cohesity cluster requires the built-in administrator credentials, such as the local system administrator of the VM. It uses the credentials provided in the **Mount Target Credentials** fields for mounting the volume. If the UAC property "Admin Approval mode" is disabled, then you can specify any user who is part of the system's Local Administrator group. • **Note:** VMware Tools are required to bring the disks online, but mounting is still supported without VMware Tools.	The VMDKs are just attached to the VM and the VM may bring the disks online. Cohesity cluster does not explicitly bring the volumes online within the VM. In this case, the Target Volume displays as "n/a" after mounting the volume succeeds.
Standalone Hyper-V host or cluster	You must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.	Not Applicable

10. Optional. Edit the default settings of the following **Recovery Options**:
    - **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
    - **Task Name**: Change the default name of the recover task.
11. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

> **Note:** If a drive letter was provided as the file path for a file-based backup of NTFS volumes, the default attributes for drive letters were preserved, so hidden drives remain hidden. They are still recovered and can be accessed by providing the full path in the search bar.

Instant Volume Mount to a New Location for Physical Server

To perform instant volume mounting to the original server:

1. Select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Physical Server> Instant Volume Mount**.

   The Instant Volume Mount page appears.
3. Search for the server that contains the data you want to mount by entering characters of the server name or by specifying the wildcard character *.

   A list of servers displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, server type, or date.
4. From the list of search results, select the server that contains the data you want to mount.

   By default, the latest snapshot is selected for recovery.

   > **Note:** To instant mount volumes from replicated snapshots, sign in to the cluster that stores the replicated snapshots and start the workflow from that cluster. Replicated snapshots appear in recovery searches the same way as local snapshots. You cannot select a remote replication target as a Location on the source cluster. Use Cloud Archive locations only for archived copies.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <server_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

b. Select the snapshot to be recovered by clicking the corresponding blue dot.

c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

d. Click **Select Recovery Point** and then go to step 6.

- To select a snapshot from the **List** mode:

a. Select the **List** tab from the top-right corner of the Edit recovery point for <server_name> page.

By default, the page displays the snapshots available in the last one month.

b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

c. Select the snapshots to be recovered.

d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived in the cloud by clicking the cloud ( ) icon if available in the **Location** option.

e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Select Volume(s)** option, perform one of the following actions:
   - Enable **Select all Volumes** to recover all the volumes of the selected server
   - From the **Select one or more options** drop-down list, select the individual volumes you want to recover.

   After volumes are mounted, they may be displayed as different volumes than their source volume. For example, D: volume may be mounted as X: volume.
8. From the **Recover To** options, select **New Location** to recover the volume to the original location.
9. From the **Source** drop-down menu, select a registered source having the server to which you want to recover the volume. You can also register a new source by clicking **Register Source** from the **Registered Source** drop-down menu.
10. From the **Target** drop-down menu, select a physical server to which the volume will be mounted.

> **Note:** To mount to a physical server, the server must already have the Cohesity agent installed and registered with the Cohesity cluster. If not, the server will not be available as a mount target.

11. (Applies only for VMs) With the **Ensure Disks are Online** option enabled, select one of the following options to deploy Cohesity Agent on the target VM.

> **Note:** For Linux target VMs, you cannot disable this option. For Window target VMs, you can disable this option.

- **Auto Deploy Cohesity Agent:** Select this option and provide the username and password of the target VM to automatically deploy the Cohesity Agent on the target VM.
- **Use Existing Cohesity Agent**: Select this option to use the existing Cohesity Agent deployed on the target VM

  The following table provides information on how the volume is mounted when the **Ensure Disks are Online** option is enabled and disabled.

Environment	Ensure Disks are Online option is Turned On	Ensure Disks are Online option is Turned Off
VMware	After the VMDKs are attached to the VM, Cohesity cluster explicitly mounts the volume. To mount the volume to the VM, Cohesity cluster requires the built-in administrator credentials, such as the local system administrator of the VM. It uses the credentials provided in the **Mount Target Credentials** fields for mounting the volume. If the UAC property "Admin Approval mode" is disabled, then you can specify any user who is part of the system's Local Administrator group. • **Note:** VMware Tools are required to bring the disks online, but mounting is still supported without VMware Tools.	The VMDKs are just attached to the VM and the VM may bring the disks online. Cohesity cluster does not explicitly bring the volumes online within the VM. In this case, the Target Volume displays as "n/a" after mounting the volume succeeds.
Standalone Hyper-V host or cluster	You must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.	Not Applicable

12. Optional. Edit the default settings of the following **Recovery Options**:
    - **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
    - **Task Name**: Change the default name of the recover task.
13. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

> **Note:** If a drive letter was provided as the file path for a file-based backup of NTFS volumes, the default attributes for drive letters were preserved, so hidden drives remain hidden. They are still recovered and can be accessed by providing the full path in the search bar.

Recover Physical Server Files or Folders

Cohesity provides the ability to recover files and folders from a Snapshot created earlier by a Protection Group. Files and folders can be recovered to their original location (limitations apply to physical Servers and specific operating systems) or to a newly specified location, which can be within the original Source or a different one. You can choose to retain the recovered files' and folders' original (at the time of the backup) permissions and attributes. You can also download files and folders from selected Snapshots that were created by a CohesityProtection Group. However, only items that were indexed when the Snapshot was created can be downloaded. Recovering and searching for files and folders in Snapshots created by other backup software is not supported. For information on managing these Protection Groups, Add or Edit a Protection Group for Physical Servers.

To learn more, see About File Recovery.

For more information about the recovery options:

- Recover Files or Folders to the Original Location
- Recover Physical Server Files or Folders to a New Location
- The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

  On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

  Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

- When recovering a file or instantly mounting a volume from a Windows VM or Physical Server Backup Source that has Windows deduplication installed and enabled for one or more volumes, you must choose a target machine that also has Windows deduplication installed (it does not have to be enabled for any volume). (However, this rule does not apply to Nutanix AHV VMs. If AHV VMs are enabled with Windows deduplication, the only supported recovery option is full VM recovery.)

  If the target does not have Windows deduplication installed:

  - File level recovery will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."
  - Instant volume mounting will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."

  To determine if Windows deduplication is installed on the Source or target machine, follow the steps given below:

  1. Open **Server Manager**.
  2. Select **Roles and Features > File and Storage Services > File and iSCSI Services**.
  3. Select the **Data Deduplication** check box, if necessary.
  4. Click **Next** until the **Install** button is enabled and then click **Install**.

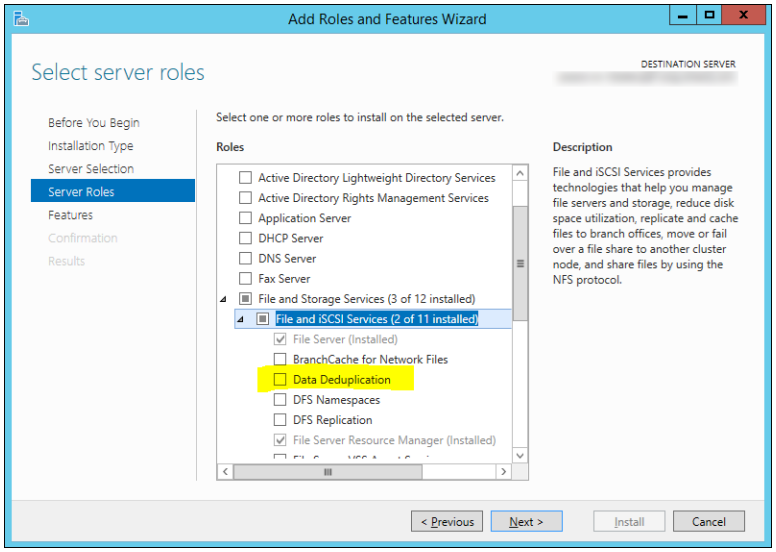

- When unzipping a zip file that was created by downloading files and folders from an archived Snapshot, if the file or folder name has encoded characters, unzip the zip file using the corresponding encoding. For example if a file name in the zip file has a UTF-8 character, unzip the file using the following command:

```
unzip -O UTF-8 Download-Files_Sep_20_2018_3-17pm_3090.zip
```

- When recovering a Linux file, the Cohesity Linux Agent runs the following commands in sudo:
  - mount
  - umount
  - findmnt
  - timeout
  - blkid
  - lsof
  - ls
  - rsync
  - losetup
  - dmsetup
  - lvs
  - vgs
  - lvcreate
  - lvremove
  - lvchange
- Windows behavior (giving each user its own set of mapped networked drives) prevents recovering to an SMB share that was mounted by a different user.

Considerations:

- For Linux Logical Volume Manager (LVM), if all the disks for a volume group are not found by the Cohesity cluster, the Cohesity cluster will not process that volume group. As a result of that, no volumes of this volume group will be recognized or indexed by the Cohesity cluster.
- File recovery is not supported for ReFS volumes in these environments: physical, VMware, Hyper-V and AHV.
- Encrypted folders that have been renamed or deleted cannot be recovered.
- Recovering files/folders with names longer than 200 characters may return an error. This limitation is due to Windows behavior when handling files/folders with long names.
- Downloading files and folders from tape archive locations is not supported.
- Recovering files and folders from VMs to physical servers and from physical servers to VMs is not supported.
- The downloadable zip file can contain regular files and folders only; symlinks are not supported. When unzipping the downloaded files/folders, use a zip utility that supports the ZIP64 format.

- For RHEL5, if LVM is not configured, running a protection job fails with the "No LVM volume found in idbmantasdev" error. For more information, see the Registering or running a Linux protection job fails when no LVM volumes are detected KB article.

Before you can recover a file from a Snapshot, the Snapshot must exist. To recover files from a Snapshot, the files must be indexed during the backup process. To learn more, Add or Edit a Protection Group for Physical Servers.

Recover Physical Server Files or Folders to the Original Location

To recover files and folders of a physical server to the original location:

1. Select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Physical Server > Files or Folders**.
3. Select any one of the following options to search for the file or folder:
    - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
    - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select..	Then..
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select..	Then..
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:

- **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.
- **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

  Character replacement explanation for : and \ in file and directory names in zip files

  If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
  Example:
  - The original name is "a:b" then the archived name is "a_COLON_b"
  - The original name is "a:b:c" then the archived name is "a_COLON_b_ COLON_c"

  If the original file or directory name contains a '\', it is replaced with '_BACK_ SLASH_' in the archive.
  Example:
  - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
  - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_ y_BACK_SLASH_z"

5. From the **Recover To**, select **Original Server** to recover the files or folders to a Original server.
6. To recover to a different location in the original server, disable the **Recover to Original Path** option, and then provide the location to which the files or folders are to be recovered in the **Recover To** field. By default, the files and folders will be recovered to the original location.
7. Optional. Edit the default settings of the following **Recovery options**:
   - **Authorized User** – Select the user account that will be used to authenticate the SMB mount during file and folder recovery to a Windows machine. Option include:
     - **Automatic** - Choose this option to let Cohesity automatically select the appropriate user for SMB mount authentication. Use this setting when the Cohesity cluster and the Windows host have a trust relationship at the Active Directory domain or forest level.
     - **Local User** - Select this option to authenticate the SMB mount using a local user account on the Windows machine. Use a local user who has the required permissions to authenticate SMB mounts created by the Cohesity cluster. Such a user should have been granted 'Allow SMB

Access Tokens' as a privilege. For more details, see User Settings.

- **Specific AD Identity** - Select this option when the agent is not joined to Active Directory but the Cohesity cluster is joined to AD. Specify the Active Directory domain and provide the AD user credentials that should be used to authenticate the SMB mount.

- **Overwrite Existing File/Folder**: By default, this option is disabled to create the files and folders in the specified location. Enable this option to overwrite the existing files and folders. Any duplicate files are skipped.

  Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

  Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

  - C:\restore\1.txt
  - C:\restore\cohesity<taskid>_0\1.txt
  - C:\restore\bar
  - C:\restore\cohesity<taskid>_1\bar

  If **Overwrite Existing File/Folder** is enabled, this will overwrite the existing files and folders. Whether overwriting occurs depends on the application using the file.

- **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only, files, then files will receive the new timestamps.
- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.
- **Save Success Files**: Enable this option to download the success logs for a recovery run.

8. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Download List Information for a Recovery Run

System-generated lists are used to verify success or errors during recovery. You can now download the list information as a CSV file for both successes and failed recovery job runs on the Cohesity cluster.

To download the logs for a recovery run:

1. Select **Data Protection** > **Recoveries** > **Recover**.
2. On the **Recoveries** page, select the recovery run for which you want to download the debug logs.
    1. If you select a successful recovery run, on the Recovery Details page, select the **Show Subtask** button and then click **Download Success** Logs*.
    2. If you select a failed recovery run, on the Recovery Details page, select the **Show Subtask** button and then click **Download Errors**.
    3. If the recovery run is partially successful, select the **Show Subtask** button and then click either the **Download Errors or Download Success Logs***.

> **Note:** * You must enable the **Save Success Files** option under **Recovery** Options while creating a recovery job to be able to download the success logs for a recovery run.

Recover Physical Server Files or Folders to a New Location

To recover files and folders to a new server:

1. Select **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Physical Server > Files or Folders**.
3. Select any one of the following options to search for the file or folder:
    - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
    - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select..	Then..
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select..	Then..
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up the browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select..	Then..
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:

- **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.
- **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

  Character replacement explanation for : and \ in file and directory names in zip files

  If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
  Example:
  - The original name is "a:b" then the archived name is "a_COLON_b"
  - The original name is "a:b:c" then the archived name is "a_COLON_b_ COLON_c"

  If the original file or directory name contains a '\', it is replaced with '_BACK_ SLASH_' in the archive.
  Example:
  - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
  - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_ y_BACK_SLASH_z"

5. From the **Recover To**, select **New Server** to recover the files or folders to a new server.
6. From the **Registered Source** drop-down menu, select a registered server to which you want to recover the files and folders. You can also register a new source by clicking **Register Source** from the **Source** drop-down menu.
7. From the **Target** drop-down menu, select a server to which the files and files are to be recovered.
8. Change the default location to which the files or folders are recovered. By default, the files and folders are recovered to the **/tmp/** directory.
9. Optional. Edit the default settings of the following **Recovery options**:
   - **Authorized User** – Select the user account that will be used to authenticate the SMB mount during file and folder recovery to a Windows machine. Option include:
     - **Automatic** - Choose this option to let Cohesity automatically select the appropriate user for SMB mount authentication. Use this setting when the Cohesity cluster and the Windows host have a trust relationship at the Active Directory domain or forest level.

    - **Local User** - Select this option to authenticate the SMB mount using a local user account on the Windows machine. Use a local user who has the required permissions to authenticate SMB mounts created by the Cohesity cluster. Such a user should have been granted 'Allow SMB Access Tokens' as a privilege. For more details, see User Settings.
    - **Specific AD Identity** - Select this option when the agent is not joined to Active Directory but the Cohesity cluster is joined to AD. Specify the Active Directory domain and provide the AD user credentials that should be used to authenticate the SMB mount.
- **Overwrite Existing File/Folder**: By default, this option is disabled to create the files and folders in the specified location. Enable this option to overwrite the existing files and folders. Any duplicate files are skipped.

  Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

  Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

    - C:\restore\1.txt
    - C:\restore\cohesity<taskid>_0\1.txt
    - C:\restore\bar
    - C:\restore\cohesity<taskid>_1\bar

  If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in Fuse may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
- **Preserve File/Folder Attributes**: By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only, files, then files will receive the new timestamps.
- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

- **Save Success Files**: Enable this option to download the success logs for a recovery run.

10. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Download List Information for a Recovery Run

System-generated lists are used to verify success or errors during recovery. You can now download the list information as a CSV file for both successes and failed recovery job runs on the Cohesity cluster.

To download the logs for a recovery run:

1. Select **Data Protection** > **Recoveries** > **Recover**.
2. On the **Recoveries** page, select the recovery run for which you want to download the debug logs.
    1. If you select a successful recovery run, on the Recovery Details page, select the **Show Subtask** button and then click **Download Success** Logs*.
    2. If you select a failed recovery run, on the Recovery Details page, select the **Show Subtask** button and then click **Download Errors**.
    3. If the recovery run is partially successful, select the **Show Subtask** button and then click either the **Download Errors or Download Success Logs***.

> **Note:** * You must enable the **Save Success Files** option under **Recovery** Options while creating a recovery job to be able to download the success logs for a recovery run.

Recover Physical Servers

Cohesity enables you to perform:

- Bare machine recovery (BMR): To learn more about Cohesity's Cristie BMR integration, see About Bare Machine Recovery. For Cristie product documentation, see CBMR Documentation.
- Physical volumes recovery: For the Windows and Linux servers, you can perform block-level recovery of physical volumes to another volume in the same or different target.

Before you can recover a volume or perform BMR, a snapshot or BMR data must exist. To include CBMR data in physical server protection, the policy must include a BMR backup schedule (Create or Edit a Standard Policy) and the Protection Group must have BMR backup enabled (Add or Edit a Protection Group for Physical Servers).

The options to recover a physical server are as follows:

- Recover Physical Server to the Original Location
- Recover Physical Server to a New Location

Recover Physical Server to the Original Location

Before you can recover volume or bare machine recovery (BMR) data, a snapshot must exist. To create Snapshots, create and run a Protection Group.

You can recover the volumes of a physical server or perform bare machine recovery (CBMR) data, to its original location.

> **Note:** Cohesity's Cristie bare machine recovery (CBMR) is supported for Windows and Linux physical servers. For more information, see About Bare Machine Recovery.

To recover a volume or BMR:

1. Select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Physical Server > Physical Server**.

   The Physical Server page appears.
3. Search for the physical server that contains the snapshots to recover from by entering characters of the server name, or by specifying the wildcard character *.

   A list of servers displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, and date.
4. From the list of search results, select the server whose volume you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

- To select a snapshot from the **Timeline** mode:
  a. From the **Choose a date** drop-down, select a date.

     The snapshots available on the selected date are displayed in blue dots on the timeline bar.

  b. For recovering CBMR snapshots, select the **Show BMR snapshots** check box.

     The BMR snapshots are also displayed.

  c. Select the snapshot to be recovered by clicking the corresponding blue dot.

  d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

  e. Click **Select Recovery Point** and then go to step 6.
- To select a snapshot from the **List** mode:
  a. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page.

     By default, the page displays the snapshots available in the last one month.

  b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

  c. For recovering CBMR snapshots, select the **Show BMR snapshots** checkbox.

     The BMR snapshots are also displayed.

  d. Select the snapshots to be recovered.

  e. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

  f. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the volume to the original location.
8. From the **Select Volume(s)** option perform the following actions:

   a. Select the **Force unmount volume** check box to forcefully unmount the target volume if it is already being used by other processes.
   b. From the list of volumes, select the volumes you want to recover, and from the corresponding **Select** drop-down list, select the volumes onto which the volume should be recovered.
9. **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
10. **Task Name**: Change the default name of the recovery task.
11. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Cohesity also provides the ability to tear down (delete) BMR jobs for Windows and Linux servers. If you tear down a recovery job, all the views created as part of the recovery task will be inaccessible from the Cohesity cluster.

To tear down a BMR job:

1. Select **Data Protection** > **Recoveries**.
2. Click the BMR job you want to tear down.
3. Click **Tear Down** located at the top-right of the page.
4. Click **Yes, Teardown** to confirm.

Recover Physical Server to a New Location

Before you can recover volume or bare machine recovery (BMR) data, a Snapshot must exist. To create Snapshots, create and run a Protection Group.

You can recover the volumes of a physical server or perform bare machine recovery (CBMR) data, to a new location.

> **Note:** Cohesity's Cristie bare machine recovery (CBMR) is supported for Windows and Linux physical servers. For more information, see About Bare Machine Recovery.

To recover a volume or BMR:

1. Select **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Physical Server > Physical Server**.

   The Physical Server page appears.

3. Search for the physical server that contains the snapshots to recover from by entering characters of the server name, or by specifying the wildcard character *.

   A list of servers displays.

   You can optionally filter the search results by source, Protection Group, Storage Domain, and date.

4. From the list of search results, select the server whose volume you want to recover.

   By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **edit** icon.

   The Edit recovery point for <VM_name> page appears.

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

   By default, the option to select the snapshot is displayed in the **Timeline** mode. You can also choose to view the option in List mode by clicking the **List** tab in the top-right corner of the page.

   - To select a snapshot from the **Timeline** mode:
     a. From the **Choose a date** drop-down, select a date.

        The snapshots available on the selected date are displayed in blue dots on the timeline bar.

     b. For recovering CBMR snapshots, select the **Show BMR snapshots** check box.

        The BMR snapshots are also displayed.

     c. Select the snapshot to be recovered by clicking the corresponding blue dot.

     d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option.

     e. Click **Select Recovery Point** and then go to step 6.
   - To select a snapshot from the **List** mode:
     a. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page.

By default, the page displays the snapshots available in the last one month.

b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range.

c. For recovering CBMR snapshots, select the **Show BMR snapshots** check box.

The BMR snapshots are also displayed.

d. Select the snapshots to be recovered.

e. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived in the cloud by clicking the cloud ( ) icon if available in the **Location** option.

f. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** to recover the volume to a new location.
8. From the **Registered Source** drop-down menu, select a registered source having the server to which you want to recover the volume. You can also register a new source by clicking Register Source from the **Registered Source** drop-down menu.
9. From the **Physical Server** drop-down menu, select a server to which the volume will be mounted.

> **Note:** The target server must already have the Cohesity agent installed and registered with the Cohesity cluster. If not, the server will not be available as a mount target.

10. From the **Select Volume(s)** option perform the following actions:
    a. Select the **Force unmount volume** check box to forcefully unmount the target volume if it is already being used by other processes.
    b. From the list of volumes, select the volumes you want to recover, and from the corresponding **Select** drop-down list, select the volumes onto which the volume should be recovered.
11. **Cluster Interface**: By default, the **Auto Select** option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the **Interface Group** drop-down.
12. **Task Name**: Change the default name of the recovery task.

13. Click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Cohesity also provides the ability to tear down (delete) BMR jobs for Windows and Linux servers. If you tear down a recovery job, all the views created as part of the recovery task will be inaccessible from the Cohesity cluster.

To tear down a BMR job:

1. Select **Data Protection** > **Recoveries**.
2. Click the BMR job you want to tear down.
3. Click **Tear Down** located at the top-right of the page.
4. Click **Yes, Teardown** to confirm.

Chapter 17: MS Applications

Cohesity supports the backup and recovery of the following applications:

Active Directory 3927
MS Exchange Server 3944

Active Directory

This chapter includes the following topics:

- Active Directory Key Concepts
- Active Directory Protection Requirements
- Active Directory Best Practices
- Active Directory Non-Cohesity-Specific Information
- Set Up Active Directory for Data Protection
- Protect Active Directory
- Recover Active Directory Objects
- Recover Active Directory Domain Controller
- Active Directory Data Protection FAQs

Active Directory Key Concepts

The following topics introduce you to key concepts about Active Directory protection and object recovery.

About Active Directory Data Protection

A Cohesity Active Directory (AD) backup captures an app-consistent VSS snapshot of all volumes on the Active Directory domain controller (DC). For instructions on how to use the Cohesity UI to back up AD, review Protect Active Directory.

Cohesity's AD backup implementation is intended to protect sufficient AD data to enable granular object recovery and object attribute compare and recovery.

If the objective is to recover an entire AD DC virtual server, protect the VM as described in Add or Edit a Protection Group for Virtual Servers. For AD DC bare metal server recovery, protect the process described in Add or Edit a Protection Group for Physical Servers. App-consistent VM/bare metal snapshots and crash-consistent snapshots are available.

All configurations are for illustrative purposes only. Cohesity recommends following Active Directory best practices.

The following diagram displays an example Microsoft application environment that contains Active Directory.

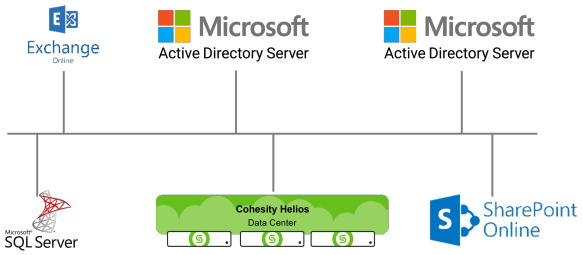

The AD backup task consists of the following actions:

1. When a Protection Group starts, the Cohesity cluster requests an AD snapshot be taken.

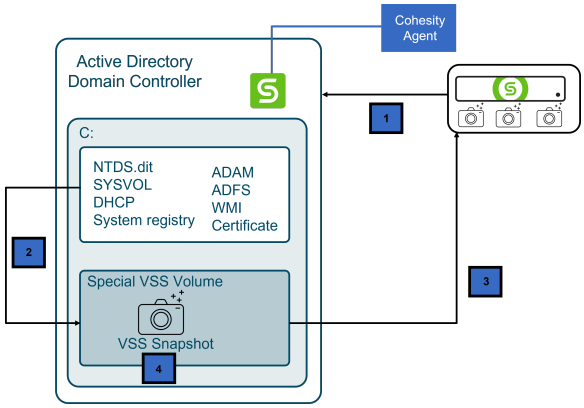

2. A Volume Shadow Copy Service (VSS) snapshot of AD is captured using VSS.

   The snapshot is a full VSS backup and uses Cohesity change block tracking (volume CBT).

3. The snapshot data is copied to the Cohesity cluster.
4. The VSS snapshot is deleted.

About Active Directory Object Recovery

A Cohesity Active Directory (AD) recovery mounts the volume that contains NTDS.dit. For instructions on how to use the Cohesity UI to recover AD, review Recover Active Directory Objects.

You can use Cohesity for these Active Directory recovery types:

- Granular object recovery

  For example, Users, Computers and Groups
- Granular object attribute compare and recovery

  For example, description, email address, group membership, custom attributes
- Group Policy Object (GPO)

The AD recover task consists of the following actions:

1. When an AD recovery starts, the Cohesity cluster creates an instant volume mount (IVM) on the cluster that contains the requested snapshot of AD data (NTDS.dit).

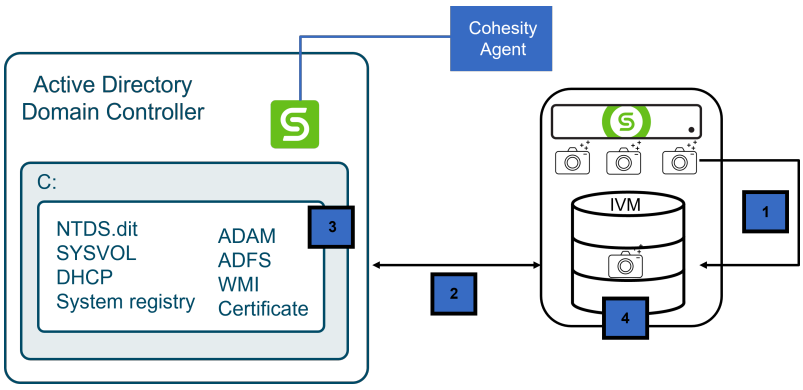

2. The Cohesity agent mounts the IVM and accesses it through SMB. If the Cohesity cluster has joined an AD domain, then the cluster must join that AD domain (Join an AD Domain) to access the SMB volume.
3. The agent accesses the snapshot on a custom LDAP port, such as 53389, and the three subsequent ports.

   LiveAD is accessed to execute LDAP queries based on user input to compare objects, accept object or attribute and update LiveAD.
4. Optionally unmount the IVM.

   You can tear down the IVM when you finish using it.

Group Policy Object Recovery

You can recover the GPO if you have accidentally deleted it or if the GPO was updated and you do not want the updates.

GPO recovery involves recovering the following two components of the GPO and also the GPO links:

- **Group Policy Container** - It represents the attributes relevant to the GPO content. Each GPO contains two containers Machine and User and the respective containers contain specific information related to the machine and user nodes of the GPO. For

more information, see Group-Policy-Container class.

- **Group Policy Template** - Contains the content of the GPO. It resides in a share known as SYSVOL. The share is replicated to every DC in the domain and the share contains separate directories named with GPO GUID for each GPO.
- **GPO links** - A GPO can be linked to different AD objects such as Domain, Site, and OrganizationalUnit. The links are maintained as the attribute (gPLink) of the corresponding AD objects.

Cohesity enables you to recover the following using the Cohesity UI:

- the Group Policy Container in AD.
- the SYSVOL folder for the GPO
- GPO links for all the objects that are linked to the GPO in cases where the GPO was deleted.

Active Directory Protection Requirements

Supported Versions

For more information on the supported versions, see Microsoft Applications.

Prerequisites

- Verify anti-virus exclusions are added for Cohesity Windows agent binaries under `C:\Program Files\Cohesity` and the log folder, `C:\ProgramData\Cohesity`.
- For the protection and granular recovery, the minimum required Cohesity version and agent is 6.4.
- Group Policy Object (GPO) on PowerShell execution. On all AD DCs where the Cohesity agent is running, Windows PowerShell script execution must be enabled.

Protection

- On all DCs where the agent is installed, open Cohesity port 50051 (For more information, see Manage Firewall Ports.)

Recovery

- On the recovery target DC, open the custom LDAP ports specified in the recover task. See Recover Active Directory Objects for more information.
- The AD Adapter has the same port requirements as filesystem backup. For recovery, it is the same as filesystem via an SMB (445), plus the addition of any 4 free consecutive ports of your choice.
- If the Cohesity cluster has joined an AD domain, then the cluster must join the target recovery AD domain to access the SMB volume. If the cluster has joined a forest (root/parent domain) and Trusted Domain Discovery is enabled from the cluster, the cluster does not need to join the child domains to perform recovery.

  The SMB volume is accessed using the Cohesity cluster computer object credentials registered in the AD. Ensure that there is a DNS server entry for the cluster computer object FQDN and that clocks are synchronized between AD and the Cohesity cluster (Kerberos allows no more than 5 minute skew).
- To perform object recovery to a Windows 2019 AD DC, you must join the Cohesity cluster to the Windows 2019 domain due to SMB security requirements.
- The recovering administrator user must have View privileges (the minimum permissions) on the production AD object. Otherwise, object recovery is prevented and the object is displayed as 'missing' and its properties are not shown.
- When recovering a deleted object from an earlier forest/domain functional backup (such as Windows 2000 functional level) to a later AD DC (such as Windows 2012 R2),

you may see differences in attributes (such as ntSecurityDescriptor) on the recovered object due to AD schema upgrades and prebuilt group memberships.

- Get-SystemAttributesInternal attributes (such as objectGuid, objectSid, usnChanged) are considered system properties, which are not compared and cannot be recovered due to AD restrictions. They are, however, displayed when browsing a snapshot. The Cohesity REST API has the ability to include system properties in attribute comparison.
- Manually teardown and rerun the domain recovery search and attempt object recovery if the connection to the mounted AD snapshot is lost. This happens if your Windows machine reboots in the middle of a granular restore or if the single-node cluster gets upgraded/rebooted.

Multitenancy

- The Cohesity cluster is part of both domains.
- A DNS server is set with conditional forwarding that forwards requests to the right Domain DCs.
- AD replicas are in the same DC as the Cohesity cluster.

Supported Workflows

The AD adapter supports the following workflows:

- Backup of Domain Controllers.
- Granular Recovery of AD objects.
- Granular Recovery of GPOs (Group Policy Objects).
- Granular Recovery of AD object attributes.

Minimum Permissions

- Agent installation - Install the agent on all AD DCs with FSMO roles that you want to protect.

  Use an AD Domain Admin account to install the Cohesity agent with the Volume CBT driver.

  The service can run as the LocalSystem account.
- Recovery - The AD account that performs recovery must be one of the following:
  - Domain Admin or Enterprise Admin
  - Domain user with delegated permissions to the OU where objects will be recovered, and reanimate permission on the AD Recycle Bin, if enabled (Delegation of Active Directory Recycle Bin).

Limitation

- Mounting a snapshot that was captured as a lower version domain/forest functional level/AD schema than the target DC FFL/Schema version, requires setting this gflag to true: magneto_active_directory_mount_ntds_allow_upgrade. Otherwise, the mount may fail after a 1 hour timeout.

  Obtain the AD schema version from:

  ```
  Get-ADObject (Get-ADRootDSE).schemaNamingContext -Property
  objectVersion
  ```

- If a Cluster or Windows machine is not part of Active Directory, workflows involving mounting Cohesity View on the Windows machine over SMB, including IVM restore, file folder restore from block-based backup, CBMR backup, and other Hyper-V workflows, will fail on the Cohesity cluster from Cohesity version 6.8 onwards.

Active Directory Best Practices

Cohesity recommends the following best practices when protecting Active Directory:

- Protect all domain controllers that have FSMO roles.

  For example, in the case of role failover this practice helps to ensure a backup (snapshot) is present even if a backup DC becomes primary at some point.

Active Directory Non-Cohesity-Specific Information

This section discusses Active Directory requirements, limitations and behavior that are not specific to Cohesity.

Target AD Server Reboot During Recovery

If a target AD server is rebooted or a single-node Cohesity cluster is upgraded/rebooted during a recovery, the recovery will fail. The Cohesity admin must perform a tear down operation of the recovery task and then re-attempt object recovery.

Objects Recovered with New SIDs when AD Recycle Bin is not Enabled

Any recovered user, group and computer objects will have a new security descriptor (SID) if the Recycle Bin is not enabled on the AD forest, or if you recover after the tombstone lifetime has expired.

This is due to an AD limitation that prevents APIs setting object SIDs. The SID is permanent within a domain during an object's lifetime. This is a general limitation for all backup vendors and Microsoft API users. If these SIDs were inherent to any groups or ACLs, they would need to be adjusted manually. Cohesity always attempts to recover a missing object from the Recycle Bin and creates a new object (thus, a different SID) only when it is not found in the Recycle Bin. An object comparison reveals the SID difference.

NTDS VSS Limitation

Microsoft NTDS VSS writer does not support Ntds.dit files on a mount point (such as C:\mount).

ntSecurityDescriptor ACL Attribute Recovery

The ntSecurityDescriptor security ACL attribute of an object within an OU may be delegated and inherited. In this case, the ACL cannot be recovered for the object because AD prevents it, showing the attribute as different. Recover the ntSecurityDescriptor of the OU.

Object Recovery by a Non-Domain Admin/Enterprise Admin

If a non-domain administrator/enterprise administrator account (such as HelpDesk) is delegated to perform object recovery and the Recycle Bin feature is enabled on the AD forest, the account must have Reanimate Tombstone object permissions to look up and recover objects from the AD Recycle Bin. For more information, see MSDN.

Delegated Permissions for a Non-Admin for Recovery

When giving delegated permissions to a non-admin user, choosing everything under "Delegate the following common tasks" gives delegate permissions for create/delete/manage user accounts, groups and inetOrgPerson accounts. You should specify permissions for other object types (such as computers and organizational units) separately. For object recovery to succeed, you must grant permissions to read and write all attributes in addition to create and delete object permissions.

Set Up Active Directory for Data Protection

This topic describes how to set up an Active Directory (AD) domain controller (DC) to be protected by a Cohesity cluster.

If the AD DC is already registered as a physical server with the Cohesity cluster, skip to the Register Active Directory steps. If the server is not registered with the Cohesity cluster at all, start with Set Up the Cohesity Agent on an AD DC.

Set Up the Cohesity Agent on an AD DC

Install the Cohesity agent on each AD DC that you want Cohesity to protect.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an AD Domain Admin, run the executable and complete the installation wizard.
4. **Volume CBT (Changed Block Tracker):** Install this component for the best incremental backup performance. Installing this component requires a one-time reboot to load the Cohesity Volume CBT driver.

   File System CBT (Changed Block Tracker): This component is not required.

5. **Service Account Credentials:** The service can run as the LocalSystem account.
6. Wait for the installation to finish. The agent starts automatically.
7. Reboot the machine.
8. Repeat the agent installation process on each AD domain controller you want to protect.

Register Active Directory

Register each AD DC that you want Cohesity to protect.

1. Select **Data Protection > Sources**.
2. Click **Register** on the top-right of the page and then select **Applications > Active Directory**.
3. Specify the hostname (FQDN) or IP address (IPv4 or IPv6) of the server to register.
4. Click **Register**.

   When registration finishes, a new entry appears in two sections on the Sources page, Active Directory and Physical. If the DC was already registered as a physical server, a new entry appears in the Active Directory section.

   The Sources page displays the following information for AD sources:

   - Total Size shows the Windows physical volume size.
   - Protected Size shows the AD database size.

Related Topic

Active Directory Protection Requirements.

Protect Active Directory

This topic describes how to protect Active Directory (AD).

Protect Active Directory Overview

1. Register an AD source with the Cohesity cluster if one is not already registered. See Set Up Active Directory for Data Protection
2. Select **Data Protection > Sources**.
3. Locate the AD source you want to protect and click **Protect Active Directory** from the actions menu.

Create an AD Protection Group

> **Note:** This page displays registered AD servers only, so only AD servers are selected regardless of the other servers within the source object hierarchy.

1. In the Cohesity Dashboard, select **Data Protection > Protection**
2. Click **Protect** on the top-right of the page and then select **Applications > Active Directory**. You can select multiple objects to protect.
3. In the **New Protection** pop-up, do one of the following:
   - Protect the Active Directory by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered Active Directory source or click **Register Source** to register an Active Directory. From the **Objects** list, select the objects to protect.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

     > **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

     v. Click **Protect**.
   - To protect an Active Directory by specifying all the options, click **More Options**.
4. **Source:** Select the servers to protect.
5. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

   The selected storage domain stores the snapshots by this Protection Group.
6. Click **Protect**.

   The Protection page is displayed and includes the new Protection Group.

Name: Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.

Advanced Settings

- **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

  The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
- **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
- **QoS Policy:** Select an appropriate quality of service (QoS) policy.

  > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

  - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
  - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
  - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
  - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
- **Pre & Post Scripts:** Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group run.

  For configuration details, see Configure Pre & Post Scripts.
- **Cancel Runs at Quiet Time Start:** Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a protection run starts, it continues to execute even when a quiet time period specified for this protection run starts. However, a new protection run will not start during a quiet time period.
- **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:

  - **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
  - **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
  - **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

  For more information, see Alerts.
- **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
- **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

  > **Note:**
  > The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

- **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
  - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
  - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).

Recover Active Directory Objects

This topic describes how to recover Active Directory (AD) objects. Before you can recover an AD object, a successful snapshot must exist.

Cohesity cluster must join the AD domain for AD source backup and to perform AD granular object recovery. If the cluster has joined a forest (root/parent domain) and Trusted Domain Discovery is enabled from the cluster, the cluster does not need to join the child domains. For information about joining a domain, see Join Active Directory. Kerberos authentication handles recovery credentials. AD users can be defined as Cohesity cluster users and

assigned different roles for cluster operations (for more information, see Users and Groups).

Recover AD Object Overview

1. Select **Data Protection > Recoveries**.
2. Click **Recover > Applications > Active Directory**.
3. Select the AD DC and snapshot you want to recover.

   For Sources that are protected as Virtual machines or as Physical Hosts, choose the Active Directory Snapshot instead of the VM Snapshot or Physical host Snapshot.
4. Set recovery options.

Set Recovery Options

1. **Domain Controller:** Select a domain controller.
2. **AD Administrative Username:** Enter the username to use for the recovery. It must be one of the following:
   - Domain Admin or Enterprise Admin
   - Domain user with delegated permissions to the OU where objects will be recovered and reanimate permission on the AD Recycle Bin, if enabled.
3. **Password:** Enter the account password.
4. **Port:** Enter the port number to use for the recovery.

   This port and the 3 consecutive port numbers after it must be currently unused and opened for inbound traffic in the Windows Firewall.

   Windows Server 2008 and newer versions of Windows Server have increased the dynamic client port range for outgoing connections. The new default start port is 49152, and the default end port is 65535.

   For more information on Dynamic Port Range, see Active Directory and Active Directory Domain Services Port Requirements.

   > **Note:** These LDAP ports do not need to be opened for hosts outside the AD DC. Only the Cohesity agent must use these ports.

   Example:

   ```
   C:\Users\Administrator>netsh advfirewall firewall add rule
   name="Cohesity_Agent_In" dir=in action=allow protocol=TCP
   localport=1033-1036
   ```
5. **Recover from:** Select a snapshot to recover from.
6. Click **Continue**.

Browse Snapshots

After successfully mounting a snapshot, you can browse the snapshot to review the differences between the backup and live data.

1. Select **Data Protection > Recoveries**.
2. Click the task name.
3. Click **Browse Snapshot**.

Entity Hierarchy Navigation

The protected source and where the snapshot is mounted are displayed at the top left. The AD domain's entire entity hierarchy is displayed on the left of the page. Navigate the hierarchy to display desired objects.

Search

You can use the search bar for text searches or run your own customized AD LDAP queries. Search queries are executed against the currently selected entity hierarchy level. To search the entire hierarchy, ensure the top level is selected.

An example LDAP query to search for a user named AutomationUser02:

```
(&(objectcategory=person)(sAMAccountName=AutomationUser02))
```

Differences

If a difference exists between the backup and live data, an object's difference column displays one of the following:

- Yes - The backup and live data differ
- Missing - The object is missing from the source

Recover

You can recover only objects that have differences.

Select an object's check box and click **Recover**. Drilling down to an object's lowest level allows you to view its individual properties and whether to view only recoverable changes.

When recovering multi-value objects, select one of the following:

- Merge Values - Merge the values from the backup and live data
- Overwrite Values - Overwrite the live values with the backup values

When recovering a missing user, you must provide a password that meets your enterprise password requirements or the recovery will fail. The following options are also available:

- User must change the password at the next login. This is on by default.
- Recover in a disabled state. This is off by default.

After Recovering Objects or Attributes

After a successful recovery, Recovered is displayed next to any recovered objects.

When you finish recovery, click **Done** in the top right corner and select an option:

- Stay - Continue browsing the snapshot
- Keep mount and Leave - Keep the snapshot mounted but leave the Browse Snapshot page
- Unmount and Leave - Unmount the snapshot and leave the Browse Snapshot page. This destroys the snapshot. To use the snapshot later, you must start another Recover task.

Audit Recover Tasks

Use the following methods to audit AD recoveries:

- View recoveries on the Cohesity cluster.
  1. Select **Data Protection > Recoveries**.
  2. Click the task name.
  3. The Details tab lists objects recovered from that AD snapshot. Mouse over the object name for additional details.
- View the Windows application event log in the event viewer. The source for Cohesity recoveries will be CohesityAgent.

Tear Down

If you left a snapshot mounted, tear it down using these steps:

1. Select **Data Protection > Recoveries**.
2. Click the snapshot's task name.

   Its status should be Success.
3. Click **Tear Down** and confirm.

   After teardown finishes, the status should be Destroyed.

Recover Active Directory Domain Controller

Cohesity provides the ability to recover the Active Directory (AD) Domain Controllers that are backed up on the Cohesity cluster. The procedural steps below align with Microsoft best practices.

Prerequisites and Considerations

For recovering AD:

- Snapshots of the AD domain controller must be available on the Cohesity cluster.

- You must have the local administrator username and password.
- Windows Firewall should permit TCP Port 50051 inbound traffic from the Cluster for the instant mount
- Permission should be given to the instant volume mount to access the temporary SMB mount used during the recovery.
- The steps listed here are specific to the recovery of the AD Domain Controller.

Recover AD Domain Controller

Perform the following steps to recover AD Domain Controller:

1. On the Domain Controller, configure the Domain Controller host to boot in Directory Services Repair Mode (DSRM):
    1. Press **WIN+R** on the keyboard.
    2. Type **msconfig** and click **OK**.
    3. Select the **Boot** tab.
    4. Under **Boot Options**, select the **Safe Boot** check box.

        > **Note:** Once the **Safe Boot** is enabled, the AD services are disabled. DNS and all other AD services are not available.

    5. Select **Active Directory repair** and click **OK**.
    6. Restart the Domain Controller host.
2. Log in to the Domain Controller host as the local administrator.
3. In the Cohesity UI, select the snapshot for restore, then download the folders containing the AD data (NTDS and SYSVOL). For more information, see Instant Volume Mount.
4. On the Domain Controller host:
    1. Rename the old NTDS and SYSVOL folders.
    2. Copy over the same folders (NTDS and SYSVOL) from the Instant Volume Mount to the Domain Controller host.
5. From the Cohesity UI, tear down the Instant Volume Mount. For more information, see Teardown a Mounted Volume.
6. On the Domain Controller, configure the Domain Controller host to boot in normal mode:
    1. Press **WIN+R** on the keyboard.
    2. Type **msconfig** and click **OK**.
    3. Select the **Boot** tab.

4. Under **Boot Options**, clear the **Safe boot** check box.
5. Restart the Domain Controller host.

7. Log in to the Domain Controller host as the domain administrator.

Active Directory Data Protection FAQs

This topic answers frequently asked questions about protecting and recovering Active Directory (AD) objects.

Sources

Where do I need to install the Cohesity agent?

Install the agent on all DCs with FSMO roles that you want Cohesity to protect.

How do I upgrade the agent?

You can upgrade agents to the latest version from the Cohesity cluster.

Permissions

What permissions are required?

- Agent installation - Install the agent on all AD DCs with FSMO roles that you want to protect.

  Use an AD Domain Admin account to install the Cohesity agent with the Volume CBT driver.

  The service can run as the LocalSystem account.
- Recovery - The AD account that performs recovery must be one of the following:
  - Domain Admin or Enterprise Admin
  - Domain user with delegated permissions to the OU where objects will be recovered and reanimate permission on the AD Recycle Bin, if enabled (Delegation of Active Directory Recycle Bin).

Protection

If I protect the DCs, will the servers they run on be protected?

AD protection is intended to protect sufficient AD data to enable granular object and attribute recovery. The implementation does not protect the entire VM or physical server. To accomplish that, define separate VM/physical server protection.

If I protect a forest DC, are all DCs in the forest protected?

No. Protecting forest a DC does not protect all DCs in the forest. You must protect individual DCs.

Recovery

Do I need to open firewall ports when recovering?

Yes. Recovery requires opening firewall ports to mount a snapshot on AD. For example, port 33389 and the 3 subsequent ports must be opened.

A snapshot is missing a user that was added to production AD. What is wrong?

If a new user was added to production AD after the snapshot was captured, the user will not be present in the snapshot on the cluster.

How do I recover an entire OU?

First recover the OU container and then recover all objects inside it.

Can I recover multi-value attributes?

Yes. Merge/overwrite options are provided for multi-value attributes such as url, telephone, etc.

Can I recover a deleted OrganizationalUnit object?

Yes. If the OrganizationalUnit (OU) is accidentally deleted, you can recover the OU from the snapshot. When you recover an OU object, all the attributes of the OU object including the `gPLink` attribute are recovered.

Can I recover GPO links for Domain, Site or OrganizationalUnit?

Yes. You can recover GPO links by initiating the recovery at an attribute level for an object. For example, recovery of attribute `gPLink` for OrganizationalUnit.

Can I recover Active Directory Domain Controller?

Yes. For more information, see Recover Active Directory Domain Controller.

Can I recover all AD objects?

Yes. You can recover all AD objects including GPOs except system attributes.

Can I recover system attributes?

No. You cannot recover system attributes. For example, system IDs or object IDs. Cohesity cannot restore system attributes if the object is deleted or not available in the AD recycle bin.

Can I recover attributes from the AD recycle bin if the AD recycle bin is enabled?

If the AD recycle bin is enabled, Cohesity will attempt to restore from the AD recycle bin and attempt to preserve everything of that object from the AD recycle bin restore.

If the AD recycle bin is not enabled or the object is past TTL, the system attributes of that object are not restored.

MS Exchange Server

This chapter includes the following topics:

- MS Exchange Server Requirements
- Setup MS Exchange Server

- Protect MS Exchange Server
- Recover MS Exchange Server
- Unregister MS Exchange Server
- MS Exchange Server Limitation
- Truncate Logs in Microsoft Exchange Servers on Physical Windows Servers
- MS Exchange Server Troubleshooting

MS Exchange Server Requirements

This topic provides information about the supported platform, permission, requirements, and limitation for protecting the Exchange Server by Cohesity cluster.

Supported Versions

For more information, see Microsoft Applications.

Minimum Permissions

- The Cohesity agent service logon should be running as a specific AD account (not LocalSystem or local computer account such as local administrator) that has privileges to run Exchange Management PowerShell and query AD for Exchange objects. Exchange Management PowerShell is required for executing the following sets of cmdlets to get the Exchange server, DAG, and database topologies:
  - Get-ExchangeServer
  - Get-Mailboxdatabase
  - Get-Mailbox
  - Get-MailboxDatabaseCopyStatus
  - Get-DatabaseAvailabilityGroup

  The following permissions are required to register and backup Exchange Servers:
  - **Backup Operators group** on AD Domain.
  - **Exchange Servers and Organization Management groups** under the **Microsoft Exchange Security Groups** Organizational Unit.
  - **Local Administrators group** on the Exchange Server.
- For recovering mailbox or databases:
  - You need special privileges for the Ontrack user. Refer the Ontrack documentation for the permission requirements.
  - The Exchange server must have joined the same AD domain as the Cohesity cluster for SMB authentication.

Prerequisites

- Verify anti-virus exclusions are added for Cohesity Windows Agent binaries under `C:\Program Files\Cohesity including RunE2E.exe` and the log folder, `C:\ProgramData\Cohesity`.
- Install the Cohesity agent on the Microsoft Exchange Servers.
- Install Kroll Ontrack® PowerControls™ 9.2 or later on a Windows server that has joined the same active directory as the Microsoft Exchange Servers. Kroll Ontrack is required for:
  - Exchange mailbox recovery.
  - A Windows server with 32-bit Microsoft Outlook installed.
- Verify MS Exchange Server services are running.
- On the Server's Windows system, set the Power Plan to High performance.
- On physical servers with an installed agent, open the following ports:
  - TCP port 50051 for backup operations (incoming)
  - Standard AD and SMB port (445) for restore operations (outgoing)

Setup MS Exchange Server

To back up Exchange Servers (standalone and DAG) on the Cohesity cluster, you must first register them as a source on the Cohesity cluster. This topic describes how to set up standalone Exchange Servers and Exchange servers on DAG to be protected by a Cohesity cluster. This involves the following steps:

1. Download and Install the Agent.
2. Register Exchange Server

Download and Install the Agent

Install the Cohesity agent for Windows on each Exchange server that you want Cohesity to protect.

1. Select **Data Protection > Sources**.
2. Click **Download Cohesity Agent** and download the agent for Windows. Ensure the agent has been downloaded to the appropriate server.
3. Run the agent executable as an AD account admin having privileges to run the Exchange Management PowerShell and query AD for Exchange objects. For more information see, minimum permissions.

   The welcome page of the Cohesity agent setup wizard appears.
4. Click **Next**.

5. **File System CBT (Changed Block Tracker)**: Install this component to perform MS Exchange Server file-based protection.
6. **Service Account Credentials**: Enable **The following user account** and enter the specific AD account that has privileges to run the Exchange Management PowerShell and query AD for Exchange objects. For more information see, minimum permissions.
7. Wait for the installation to finish.

   The agent starts automatically.
8. Repeat the agent installation process on each Exchange server you want to protect. This includes any standalone MS Exchange servers and MS Exchange Server nodes on DAG.

Register Exchange Server

Once you install the Cohesity agent on the server, you must register the Exchange Server with the Cohesity cluster.

- Register Standalone Exchange Server
- Register Exchange Server on DAG

Ensure one of the following:

- Cohesity cluster has joined the DNS server of the Exchange Server.
- Exchange servers are specified in the host mapping on the Cohesity cluster.

Prerequisite

You must install the Cohesity agent on the server before you register the standalone Exchange server or Exchange server on DAG. Click the **Download Cohesity Protection Service** link and then download and install the Windows agent on the server. For more information, see Download and install agent.

Register Standalone Exchange Server

1. Navigate to **Data Protection > Sources**.
2. Click **Register** > **Applications** > **Exchange Server**.
3. Specify the hostname or FQDN of the Exchange server.

   > **Note:** Ensure you have installed the Cohesity agent on the server.
4. Click **Register**.

The standalone Exchange Server is registered as a source.

Register Exchange Servers on DAG

1. Navigate to **Data Protection > Sources**.
2. Click **Register** > **Applications** > **Exchange Server**.

3. Specify the hostname, FQDN of the DAG endpoint, or the FQDN of a node in the DAG.

   > **Note:** Ensure you have installed the Cohesity agent on the server.

4. Click **Register**.

   The Register MS Exchange Server screen appears. Provide the additional information to complete the registration.

5. **DAG Preference**: Select one of the following options from the drop-down list to set the backup preference:
   - **Passive and Active Copy**: Takes backup from healthy passive copy. However, if all the passive database copies on DAG are unhealthy, then Cohesity will take the backup from the active database copy.
   - **Passive Copy Only**: Takes backup only from the passive database copy.
6. **Backup Preference Order**: Select one of the following options to set the backup preference order of the Exchange Servers (nodes) of a DAG :
   - **Reverse of the databases' Activation Preference Order**: Backup passive copies with the highest activation preference order and least lag time specified in Exchange Server Management.
   - **Specify backup preference order**: Allows you to set your preferred backup order of the servers of a DAG. This option is recommended if the network distance between the Cohesity cluster and the DAG node is large.
7. Select the nodes you want to register. In the above step, if you have selected to specify the backup preference order, then the order of backup will be as per the order you select the discovered nodes/server.

   > **Note:** All the DAG nodes must be registered with one cluster. The backup will fail if you select nodes registered with different clusters.

8. If you have not installed the Cohesity agent on the server, click the **Download Cohesity Protection Service** link and then download and install the Windows agent on the server.
9. Click **Complete Registration**.

The Exchange Servers on DAG are registered as a source. However, the nodes and databases of the DAG that are in maintenance mode will not be registered. Cohesity recommends you to register all the nodes of a DAG if you have selected **Passive and Active Copy** as the backup preference option. Otherwise, if a passive copy is not found in the registered nodes, it may back up from the active copy of the database, which is undesirable.

Protect MS Exchange Server

Cohesity provides support for backup and recovery of Microsoft Exchange Servers (Standalone and DAG) running on a physical server or VM. This topic describes how to protect MS Exchange Server.

Protect Exchange Server Overview

1. Register an Exchange Server source with the Cohesity cluster if one is not already registered. See Setup MS Exchange Server.
2. Select **Data Protection > Sources**.
3. Locate the Exchange Server source you want to protect and click the action menu (⋮) and select **Protect**.

You can also protect MS Exchange Server by performing the following steps:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** > **Application** > **Exchange Server**.
3. Select the source to protect.

Create an Exchange Server Protection Group

> **Note:** This page displays registered MS Exchange Server only, so only MS Exchange servers are selected regardless of the other servers within the source object hierarchy.

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** > **Application** > **Exchange Server**.
3. In the **New Protection** pop-up, do one of the following:
   - Protect the Exchange Server by specifying a subset of options such as objects, protection group and policy by clicking the edit icon corresponding to the following options:
     i. **Add Objects**: Select the registered Exchange Server source or click **Register Source** to register an Exchange Server. From the **Objects** list, select the mailbox databases in the Exchange Server to protect.
     ii. **Protection Group**: Specify a name for the Protection Group.
     iii. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy.
     iv. **Storage Domain**: Select the appropriate storage domain.

> **Note:** This option is visible on the **New Protection** pop-up only if there is more than one storage domain.

v. Click **Protect**.

- To protect an Exchange Server by specifying all the options, click **More Options**.

4. **Source**: Select the source to protect.
5. **Objects**: Select the mailbox databases in the Exchange Server to protect and then click **Save Selection**. By default, all objects in the source are excluded as indicated by the unselected check box. Expand the source objects and select the objects you want to include in the Protection Group. A selected check box ( ) indicates the object is selected to be backed up.

   If appropriate for your organization, enable Auto Protect.

   Click to turn on Auto Protect. By enabling this option, any new mailbox database added to the Exchange Server will automatically be protected by this Protection Group and become protected on the next run.

   For example, you create a Protection Group that backs up an Exchange Server. The Protection Group runs for two weeks and then a new database is added to the Exchange Server. The next time the Protection Group runs, if the object hierarchy has been refreshed on the Cohesity cluster, the new database is also backed up even though the new database has not been explicitly selected to be included in the Protection Group. The object hierarchy is automatically refreshed every four hours. To manually refresh the object hierarchy, select **Data Protection > Sources**, find the source in the list and click . To exclude an object from Auto Protect, click .

   - Indicates an Auto Protected object.
   - Excluded objects are not protected. The icon indicates that an ancestor of the object is Auto Protected but this object is explicitly excluded.

   Protection of dismounted databases, recovery database types, and deleted databases are automatically skipped by the Protection Group.

   You can optionally use the filter icon to filter users or groups based on.

   - **Show All**—Displays all the Objects in the Cohesity cluster.
   - **Selected Objects**—Displays the selected Objects in the Cohesity cluster.
   - **Unprotected**—Displays Objects that are not protected by any Protection Groups in the Cohesity cluster—not protected by Cohesity.

- **Protected by Cohesity**—Displays Objects that are protected by any Protection Groups in the Cohesity cluster.

The graphical representations are described below.

- An unchecked box indicates the Object is excluded from any Protection Group.

  A shield indicates the Object is already protected by another Protection Group. A Windows physical Server cannot be protected by a new Job unless you remove it from the current job.

- A checked box indicates the Object will be protected by the current Protection Group.
- Click the icon to see details about the object.
- An outlined shield indicates that not all Objects are protected.
- If the object name is in red, it is inaccessible, invalid, or orphaned. If the object is in one of the specified states when the Protection Group runs, no snapshot is created. If the object becomes accessible, a snapshot is created the next time the Protection Group runs.

6. **Protection Group**: To enter a new name for the Protection Group, select **New Group** and enter the Group Name in the **Group Name** field. To use an existing Protection Group name, select the **Existing Group** and select an existing group name from the **Group** drop-down list.
7. **Policy**: Set the policy for your Protection Group. Under **Policy**, select the appropriate Protection Policy, choose one of the defaults (Bronze, Silver, or Gold), or create a New Policy if the available policies do not meet your schedule and retention needs.

   Cohesity recommends you to create a policy with the

   - **Log Backup (Databases)** option if your incremental backup window exceeds the desired time window. The backup will automatically skip the circular logging enabled databases during the log backup window.
   - **Retry Options** set to 3 and the time between retries as five minutes.

     For more information see, Create or Edit a Standard Policy.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

   > **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

A Storage Domain defines the level of storage efficiency, security, and resilience for the data being protected in Cohesity. In a Storage Domain, you can enable or disable deduplication, compression, and/or encryption for the domain, and select the necessary level of resilience (that is, the number of faults tolerated).

If a Storage Domain with the right settings exists, select that one. Otherwise, create a new one with Add Storage Domain. See Create or Edit Storage Domains.

If you need to create a new Storage Domain:

- You can enable or disable deduplication and compression at any time.
- You can enable or disable encryption only when you create the Storage Domain for the first time.
- There are no benefits if you perform deduplication and compression on the encrypted data.

8. **Start Time**: Available only if the selected policy is set to Backup Daily. Indicates when the job should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting
9. After selecting a Storage Domain, if you need to change any of the Advance settings, scroll down and click **Edit** on the right.

Advance Settings	Description
End Time	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. but it doesn't change or take effect on the currently running task.
Indexing	This option is enabled by default and is required for file recovery. By enabling this option, Cohesity cluster will scan all the files in the Protection Group and create an internal index that can be used later by a recover task to locate files by name.
Exchange Log Truncation	This option is enabled by default and truncates the Exchange server transaction logs. If this option is disabled, then "copy only" backup is taken and transaction logs are not truncated.
Cancel Runs at Quiet Time Start	Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing job runs should abort if a quiet time period specified for the Protection Group starts. By default this toggle is off, which means after a job run starts, it continues to execute even when a quiet time period specified for this job starts. However, a new job run will not start during a quiet time period.

Advance Settings	Description
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: 1. High-priority jobs 2. Medium-priority jobs 3. Low-priority jobs.
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the job.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Description	Optional. Click the plus sign and enter a description for the Protection Group.

10. Click **Protect**.

Your Protection Group runs appear in the list on the Protection Groups page.

Once your Protection Group completes its first run, the Exchange Server will be ready for recovery, replication, and archival.

Recover MS Exchange Server

This topic describes how to recover the Exchange Server objects (Exchange Databases) that you have successfully backed up on the Cohesity cluster. Cohesity supports two types of recovery options for Exchange Server objects.

View-Based Recovery

Cohesity mounts the view over SMB on the target server. The mounted views contain the database files and transaction logs. You can use Ontrack to perform recovery of mailbox databases including granular recoveries such as individual mailboxes, emails, files, and folders.

> **Note:** Microsoft Exchange (on-premises) requires Ontrack® software and a license for the granular recovery of mailbox databases such as individual mailbox, emails, files and folders.
> Cohesity recommends you to install Ontrack® on a separate Windows machine (than Exchange server) with an AD account that has appropriate mailbox permissions. The Windows machine must be registered as a physical source on Cohesity. On recovery, Cohesity will expose an SMB share with Exchange database files that can be opened using Ontrack®. You can also use Exchange tools such as **ESEUTIL.EXE** to perform manual inspection and recovery.

The Exchange database files will be recovered to the following SMB path of the Cohesity View:

```
\\<cluster_Ip>\<storage _domain>.<view_name>.fs\exchange
```

For example:

```
\\10.10.10.10\DefaultStorageDomain.ExchView_1585178933.fs\exchange
```

Database recovery

MS Exchange database recovery allows you to restore Exchange Server databases to a healthy state during a failure or data loss event. These databases store email messages, calendar items, contacts, and other vital information, making their recovery crucial for business continuity. When recovering an MS Exchange database, you can either overwrite the original database or restore to an alternate database.

Important:
If you want to **Overwrite the Original Database**, you must enable the **This database can be overwritten by a restore** parameter from the **Exchange Admin Center**.

To enable the parameter:

1. Log in to the **Exchange Admin Center**.
2. Select **Servers** from the left-nav and then navigate to **databases**.

   This will list all the databases on the Exchange Server.
3. Select the database you want to overwrite and then click the **Edit** icon at the top-right corner to open the database properties window.
4. In the database properties window, navigate to the **Maintenance Tab** and select the **This database can be overwritten by a restore** check-box.
5. Click **Save**.

To recover Exchange Server Objects:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** > **Applications** > **Exchange Server**.
3. On the Recover page, On the Recover page, select Database or View.
4. Search for the mailbox database you want to recover by entering any of the following options:
   a. Mailbox Database name
   b. User name of the mailbox
   c. E-mail address of the mailbox
   d. IP address or hostname of Exchange Server.
5. You can optionally specify the wildcard character *. You can also optionally narrow the search results by specifying filter criteria, for example, you can filter the search results by a specific Protection Group, source or date.

   Note: If you choose multiple databases, recovery to the original location is the only option available.
6. Select the object(s) you want to recover.

7. From the **Selected** object list, click the edit icon and select the exact snapshot you want to recover, and then click **Select Recovery Point**. By default, the latest Snapshot is recovered.
8. When performing a **View-Based** recovery, provide the following details:
    a. In the **Cohesity View** field, change the default name of the View where you want to recover the data to. This cannot be an existing view name. Cohesity will perform a mount to the target server and you can use this SMB path to access the view from the target server.
    b. In the **Server Running Recovery Tool** drop-down list, select the target server to which the SMB view should be mounted.

    > **Note:** For granular recovery, access the view from the target server running the recovery tool using the SMB path and then perform the desired Granular Level Recovery (GLR) operations to recover data to the server.

9. When recovering a **Database**, provide the following details:
    a. **Targets**: Select **Overwrite Original Database** or **Alternate Database**.

    If you want to **Overwrite the Original Database**, the original DB files on the Exchange Server will be deleted. To perform this operation, you must enable the **This database can be overwritten by a restore** parameter from the Exchange Admin Center. For more information, see Database Recovery.

    In addition, the **Overwriting Original Database** has an additional **Roll-Forward recovery** option. If this is enabled, the latest transaction logs available on the Exchange Server will be used when replying logs on the EDB.

    If you select an alternate database, provide the following details:

    a. **Exchange Server**: Select the target server from the drop-down where you want to recover.

    > **Note:** The target server should be registered as an Exchange Server on the Cohesity cluster.

    b. **Database Name**: Provide a database name for the restored database.
    c. **Database File Path**: Specify the database file path for the restored database.
    d. **Log Directory**: Specify the database transaction logs directory path for the restored database.

10. Click **Recover Options**.
    a. **Restore as Recovery Database**: Enable this option to restore the database as a recovery database. A recovery database (RDB) is a special kind of mailbox database that allows you to mount and extract data from a restored mailbox database as part of a recovery operation. RDBs allow you to recover data from a backup or copy of a database without disrupting user access to current data.
    b. **Mount Database**: Enable this option to restart the Information Store service and mount the new Database.

    > **Note:** When a new Database is created Microsoft recommends restarting the Information Store service. However, restarting the Information Store service will bring down all the other databases hosted on that Exchange server. This may take a long time and cause client connection failures.

    c. **Task Name**: Optionally, change the default name for the recover task in the Task Name field.
11. Click **Recover**.

The restore process is initiated. You can monitor the recovery task from the Recoveries page.

> **Note:** After a successful recovery of the Exchange Database, the user must run a full backup of the recovered database.

Unregister MS Exchange Server

You can unregister an Exchange Server from the Cohesity cluster. Before you unregister it, you must delete the Protection Groups associated with the Exchange Server.

1. Delete a Protection Group
2. Unregister Exchange Server
3. Uninstall the Cohesity Agent

Delete a Protection Group

Perform the following steps to delete a Protection Group:
If the Protection Group is deleted, then overwrite restores cannot be performed on the job.

1. Select **Data Protection > Protection**.
2. Find the Job in the list and in the **Actions** column select **Delete**.

Alternatively, drill-down on a Job name in the **Job** column of the list. In the Job Details page, click the (**Delete Job**) icon.

3. Choose one of the following:
   - **Delete Protection Group Only:** Select this option to delete only the Protection Group. Any local snapshots and or archived snapshots stored on the external targets that were captured by this Job are not deleted.
   - **Delete Protection Group and Snapshots:** If you select the option to delete the Job and all Snapshots, any local Snapshots captured by this Job and stored on this Cluster are deleted. In addition, any archived Snapshots stored on external targets that were captured by this Job are also deleted. (To selectively delete the Snapshots for specific protection runs without deleting the Job, see Deleting Snapshots or Objects.)

   If the current job is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should probably cancel these replication or archive tasks before deleting the job. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Group Run.

   If the Job is replicating to a remote cluster, the inactive job created on the remote cluster is not deleted. In addition, the snapshots that were replicated (copied) to the remote cluster by this job are not deleted. You can log in to the remote cluster to delete the snapshots and the inactive job.

Unregister Exchange Server

Follow the steps listed below to unregister the Exchange Server:

1. Navigate to **Data Protection > Sources**.
2. Navigate the source hierarchy to the Exchange Server, mouse over its actions menu, and click **Unregister**.
3. Click **Unregister**.

The Exchange Server will be unregistered from the Cohesity cluster. However, the physical server on which the Exchange Server is running will not be unregistered from the Cohesity cluster.

Uninstall the Cohesity Agent

For information about uninstalling the Cohesity agent from an Exchange server, see Install and Manage the Agent on Windows Servers.

MS Exchange Server Limitation

- **Microsoft Exchange limitation**: Exchange Management Powershell does not work with **Group Managed Service Account (gMSA)**. Therefore, the Cohesity agent

service account cannot be a gMSA. If gMSA is specified (even if the account is part of Backup Operators, Exchange Servers, Organization Management), the Powershell remote session fails.

- Only full/incremental backup runs are archived. LOG backups are not archived.

Truncate Logs in Microsoft Exchange Servers on Physical Windows Servers

Microsoft Exchange Email Server truncates transaction logs dynamically after a successful backup of a physical Windows Server that hosts the Exchange Email Server.

Cohesity has qualified truncating logs for the following Microsoft Exchange Servers hosted on physical Servers:

- Microsoft Exchange 2019 running on supported Windows Servers (Standalone and DAG)
- Microsoft Exchange 2016 running on supported Windows Servers (Standalone and DAG)
- Microsoft Exchange 2013 running on supported Windows Servers (Standalone and DAG)
- Microsoft Exchange 2010 SP3 running on supported Windows Servers (Standalone and DAG)

When a protection run protecting physical servers occurs, an app-consistent backup is always triggered so if the physical server is a Windows server with Microsoft Exchange, the Exchange logs are dynamically truncated by Microsoft Exchange. To disable log truncation for Microsoft Exchange with VMware VADP backups, see Enable or Disable Log Truncation.

If the Windows Server has an Exchange server but the log truncation fails, an error is reported. However, the backup can successfully complete.

> **Note:** The physical Server adapter includes Exchange Log Truncation functionality. You can also backup the Exchange Server using the Exchange Server adapter.

For more information on how to leverage Cohesity for Exchange Server backups using the physical adapter, see the How to protect and truncate logs for Exchange VMs (Standalone or DAG) on Hyper-V KB article.

Cohesity provides support for backup and recovery of Microsoft Exchange Servers (Standalone and DAG) running on a physical server or VM. For more information, see Protect MS Exchange Server.

Enable or Disable Exchange Log Truncation

To enable or disable Exchange Log Truncation, perform the steps below:

1. Navigate to **Data Protection** > **Protection**.
2. Click the protection job (VM that is running Exchange).
3. Click the **Settings** tab.
4. Click **Edit**.
5. Click **Edit** on the exchange object.
6. Toggle on or off **Exchange Log Truncation**.
7. Click **Save**.
8. Click **Protect**.

MS Exchange Server Troubleshooting

Category	Issue	Workaround
Source Registration	IP address-based DAG Exchange registration may fail due to Cohesity Cluster not able to resolve FQDN of the DAG nodes.	Add the Exchange node DNS server to Cohesity cluster DNS server or to Cohesity host mapping.
Recovery	Repeatedly mounting SMB share of the Exchange database recovery with the same view name and tearing down may cause issues in accessing the view in Windows Explorer.	Provide a unique view name for each database recovery.
Backup	Incremental backup schedule after a Windows reboot or DAG node switchover pulls the entire Exchange database EDB file.	None. File- based change block tracking (CBT) information is not preserved across Windows reboots. Log backup is not affected by this and only new log files are pulled for a DAG database copy.
Snapshot delete	Deletion of snapshots from protection runs in the middle or expiration of the log backup may subsequent snapshots until incremental/full to be not recoverable .	None. Delete only the oldest snapshots of backup runs.

Chapter 18: SAN Volume

Cohesity supports the backup and recovery of the following SAN devices:

- IBM Storage FlashSystem
- Pure Storage
- HPE Alletra Storage MP

IBM Storage FlashSystem

The Cohesity cluster supports backing up IBM Storage FlashSystem volumes and volume groups in a Protection Group. Protection Groups can run repetitively on a schedule.

To learn more about the relationship between Protection Groups and Policies, see About Policies and Protection Groups. For instructions on how to use the Cohesity UI to back up IBM Storage FlashSystem volumes, see Protect an IBM Storage FlashSystem Volume or a Volume Group.

Plan and Prepare for IBM Storage FlashSystem

This topic provides information on minimum permissions, prerequisites and considerations to protect IBM Storage FlashSystem with Cohesity cluster.

Minimum Permissions

Ensure that the IBM Storage FlashSystem user account you use to register, manage, and protect the FlashSystem resources has an **Administrator** role. For security reasons, Cohesity recommends not to use the **IBM superuser** account or any account configured within the **SecurityAdmin** group.

Support Matrix

Before you register your IBM Storage FlashSystem with the Cohesity cluster, ensure that Cohesity supports the IBM Storage FlashSystem versions. For more information, see Supported Software.

Ports Used for Communication

Ensure all necessary ports are open in the firewall between the Cohesity cluster and IBM Storage FlashSystem. For more information, see Manage Firewall Ports.

Prerequisite

- IBM Storage FlashSystem API and data traffic ports should be accessible from all the Cohesity nodes. For more information, see Ports Used for Communication.
- It is important to note that the Cohesity cluster's node(s) should not be added manually as a host or host cluster on the IBM Storage FlashSystem as this will be done automatically when registering the array as a source on the Cohesity cluster.

Considerations

- Due to the IBM Storage FlashSystem volume name character limit, Cohesity does not support backing up volumes with names greater than 45 characters.
- IBM Storage FlashSystem supports block-based backups of volumes and volume groups.
- With IBM Storage FlashSystem, you have the flexibility to create volumes while utilizing various capacity-saving methods such as thin-provisioning, compression, deduplication, or a combination of these methods. However, it's important to note that if there is a predefined provisioning policy set for the storage pool, the volume-level settings will not be applicable. In case of a restore, the system will restore the capacity-saving options based on the target storage pool's properties and the existence of any predefined provisioning policies.
- IBM Storage FlashSystem supports Safeguarded Snapshot creation for basic volumes, but it is not supported for mirrored volumes.

Register or Edit an IBM Storage FlashSystem

You can register IBM Storage FlashSystem on the Cohesity cluster to protect the IBM Flash volumes.

Before you Begin

Meet Prerequisites

Ensure that the requirements for the IBM FlashSystem, such as software version, user account role privileges, and firewall ports, are met.

Check iSCSI connectivity

By registering an IBM FlashSystem as a storage array, you can ensure that its Logical Unit Number (LUN) becomes eligible for protection on the Cohesity cluster. The LUN of both the Internet Small Computer Systems Interfaces (iSCSI) or Fiber Channels (FC) is needed to access the storage device.

After completing the iSCSI configuration and enabling the iSCSI discovery on the IBM FlashSystem, verify the iSCSI connectivity between the Cohesity cluster and the IBM FlashSystem. To check if the connection is working, log into any of the Cohesity cluster nodes and run the following command:

```
iscsiadm -m discovery -t sendtargets -p <ipaddress>
```

For additional iSCSI commands, see Additional iSCSI Commands.

Register IBM Storage FlashSystem

1. Navigate to **Data Protection** > **Sources**.
2. Click **Register** > **SAN**.
3. Select the source type as **IBM Storage FlashSystem**.
4. Enter the information described below:

Field	Description
Management Host or IP Address	Specify the Management Host or IP Address of the IBM Storage FlashSystem to register.
Username	Specify a user with a Storage Admin role.
Password	Enter the password for the specified user.

5. Click **Register**.

Protect an IBM Storage FlashSystem Volume or a Volume Group

A Protection Group that protects IBM Storage FlashSystem volumes captures Snapshots and these Snapshots are required for the following tasks:

- Recover volumes
- Replicate Snapshots
- Archive Snapshots

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit an IBM Storage FlashSystem
- Create or Edit a Standard Policy

- Create or Edit Storage Domains

To Protect an IBM FlashSytem Volume or a Volume Group:

You can now protect IBM Storage FlashSystem volumes or volume groups using a Protection Group on the Cohesity cluster.

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **SAN**.
3. In the **New Protection** pop-up, you can select either a single volume and volume group, or multiple volumes or volume groups, but not a combination of both.

   To Protect a Volume:

   Select the **Volume List** View ( ) on the top-right to list the volumes. From the **Volume** list, select the volumes you want to back up.

   To Protect a Volume Group:

   Select the **Volume Group** View ( ) on the top-right to list the volume groups. Select the volume groups and then click the **Auto-Protect Volume Group** option to protect the selected volume groups. If you select this option, any new volumes added to the volume group will be automatically protected from the next protection group run.

   > **Note:** Selecting an item in Volume View and changing the view to Volume Group causes the selection to be lost.

   After selecting one of the above options, click the edit icon and provide the following information:

   1. **Add Objects**: From the **Objects** list, select the **Volume Groups** you want to backup or click **Register Source** to register a new Volume Group.
   2. **Protection Group**: Specify a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
   3. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy. For **Storage Array Snapshot creation**, you have to define the storage array snapshot parameters during the policy creation. For more information, see Create or Edit a Standard Policy.

      > **Note:** If you define a storage array snapshot in the protection policy, the Protection Group will trigger a Storage Array Snapshot as per the frequency defined in the protection policy.

4. **Storage Domain**: Select the appropriate storage domain. For more information, see Manage Storage Domains.
5. To protect Volume Group by specifying all the options, click **More Options**.

Advance Settings	Description
Start Time	Available only if the selected policy is set to Backup Daily. Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.
End Time	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts before this date runs until completion even if it completes after this date.

Advance Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: a. High b. Medium c. Low
QoS Policy	Select an appropriate quality of service (QoS) policy. > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Safe Guarded Copy Snapshot	Select this option to enable safeguarded snapshots for the protection group run.

Advance Settings	Description
Pre & Post Scripts	Toggle Enable to run scripts on the Unix server. When creating a Protection Group, you can configure: • **Pre Scripts**: Scripts that will be executed before the protection run. • **Post Backup Scripts**: Scripts that will be executed after the protection run. • **Post Snapshot Scripts**: Scripts that will be executed after the snapshot has been created on the source. If enabled, the scripts are executed during the protection run as configured. You can distinguish between the scripts that are run on the different objects (entities) using the COHESITY_BACKUP_ENTITY environment variable. The Cohesity cluster populates this environment variable with the name of the backed object (or entity). For more information, see Configure Pre & Post Scripts. For more information, see Configure Pre & Post Scripts.
Quiet Times	Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. If the option is toggled off, then after a protection run starts, it continues to execute even when a quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period.
Alerts	Select one or more of the following **Alerts On** settings if you want alerts to be created for the following triggers: • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • Success: Create an informational alert when the Protection Group completes successfully. Emails are not sent when informational alerts are created. To send emails when alerts are triggered for the Protection Group, in the **Email Recipients** field, add the email addresses of the recipients.
Description	Enter a brief description about the Protection Group.

6. Click **Protect**.

Recover an IBM Storage FlashSystem Volume or Volume Group

You can recover the volume and volume groups backups from the Cohesity cluster to an IBM Storage FlashSystem.

To recover an IBM Storage FlashSystem Volumes or a Volume Group follow the steps listed below:

1. Select **Data Protection** > **Recoveries**.
2. Click on **Recover** located at the top-right of the page and then select **SAN**.
3. Search for the Volume or Volume Group in the backup run by entering the Volume Group name, a volume name, a name sub-string, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, Storage Domain, and date.
4. From the search results list, select the volume or volume group you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **Edit** icon.

   By default, the latest snapshot is selected for recovery. You can also view and select the older snapshots by clicking the **Protection Run** drop-down. When you select a snapshot, the volumes associated with that snapshot are displayed.
6. From the **Recover To** options, select **Original Location** to recover to the same storage pool of the IBM Storage FlashSystem or select **New Location**. If a new location is selected, then select the **Storage Pool** for the new location.
7. In the **Rename** field, you can specify the name of the recovered volume group by adding prefixes and/or suffix strings to the name of the original volume group.
8. **Task Name**: Change the default name of the recovery task.
9. Click Start **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the **Recoveries** page.

Recover a Storage Array Snapshot for IBM Storage FlashSystem Volume or Volume Group

You can recover a Storage Array Snapshot for IBM Storage FlashSystem Volumes or a Volume Group by following the steps listed below:

To recover a Storage Array Snapshot for IBM Storage FlashSystem Volumes or a Volume Group follow the steps listed below:

1. Select **Data Protection** > **Recoveries**.
2. Click on **Recover** located at the top-right of the page and then select **SAN**.
3. Search for the Volume or Volume Group in the backup run by entering the Volume Group name, a volume name, a name sub-string, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, Storage Domain, and date.
4. From the search results list, select the volume or volume group you want to recover.

   By default, the latest snapshot is selected for recovery.
5. To recover a Storage Array Snapshot, mouse over the snapshot from the selected list, and then click the **Edit** icon.
6. In the **Edit Recovery Point** window, select **Storage Array Snapshot** as the backup type.
7. From the **Recover To** option, **Original Location** is selected by default. You can recover the storage array snapshot to the original location only.
8. Select the **Storage Pool** to recover the storage array snapshot.
9. In the **Rename** field, you can modify the name of the recovered Volume Group by adding prefixes and/or suffix strings to the name of the original Volume Group.
10. **Use Thin Clone**: Enable this option to create a thin clone of the selected snapshot instead of a complete copy.
11. **Task Name**: Change the default name of the recovery task.
12. Click Start **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the **Recoveries** page.

Pure Storage

The Cohesity cluster supports backing up multiple Pure Storage volumes in a single Protection Group. Protection Groups can run repetitively on a schedule.

To learn more about the relationship between Protection Groups and Policies, see Manage Protection Policy. For instructions on how to use the Cohesity UI to back up Pure Storage volumes, see Manage Protection Groups.

For example, you could create a Protection Group with a schedule that backs up specific Pure Storage volumes every hour as illustrated in the following figure.

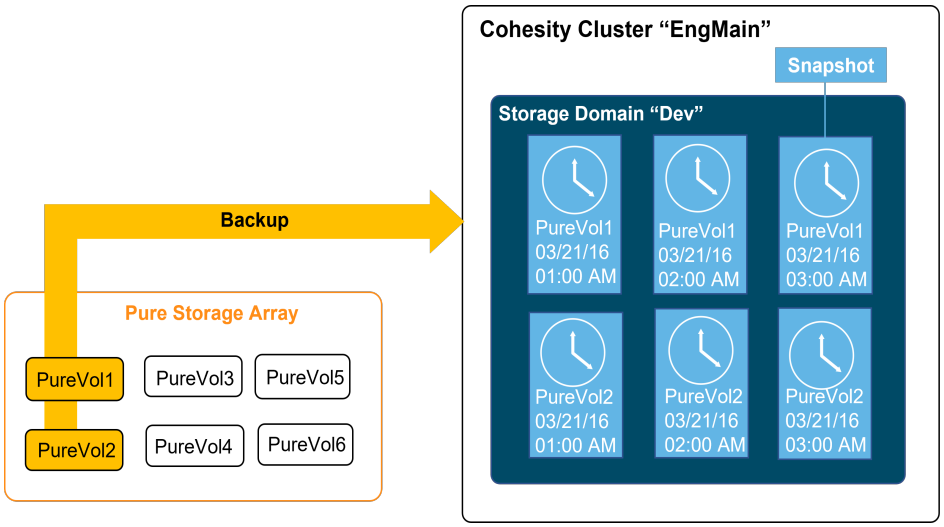

During a Protection Group Run, Cohesity Snapshots of the Pure Storage volumes are saved on the Cohesity cluster with a date and time stamp. For incremental backups, Cohesity uses Pure Storage's ability to track changed blocks so that only those blocks are backed up to the Cluster.

For instructions on how to use the Cohesity UI to back up Pure Storage volumes, see Protect a Pure Volume or a Protection Group.

Plan and Prepare for Pure FlashArray

This topic provides information on minimum permissions, prerequisites and considerations to protect Pure FlashArray with Cohesity cluster.

Minimum Permissions

Ensure you have the required minimum permissions to register Pure FlashArray as a source with the Cohesity. For more information, see Ensure Adequate Privileges for Cohesity on the Source.

Support Matrix

Before you register your Pure FlashArray with the Cohesity cluster, ensure that Cohesity supports the Pure FlashArray versions. For more information, see Supported Software.

Ports Used for Communication

Ensure all necessary ports are open in the firewall between the Cohesity cluster and Pure FlashArray. For more information, see Manage Firewall Ports.

Prerequisite

The Pure Flash array needs to have a Network Interface Card (NIC) and iSCSI discovery IP configured on the array.

Considerations

- For storage-backed snapshots to work, the SAN transport must be configured over iSCSI.
- iSCSI needs to be configured on the Pure FlashArray for volume backup.
- Recovering a Pure Storage volume with the same name as a volume that is pending eradication is not supported.

  Workaround - Recover the volume with a different name or log in to the Pure FlashArray and eradicate the deleted volume.
- If a Pure FlashArray reaches its maximum number of volumes or snapshots, Pure Storage backup jobs and recovery tasks for that array will fail. A Pure Storage volume also has a maximum number of allowed snapshots. If it reaches this limit, backing up the volume fails.

  Workaround - Delete and eradicate some volumes or snapshots and then trigger the backup or recovery.
- Due to a Pure Storage volume name character limit, Cohesity does not support backing up volumes with names greater than 45 characters.
- Cohesity does not currently support backing up Pure FlashArrays with Challenge-Handshake Authentication Protocol (CHAP) authentication enabled.
- If a Cohesity cluster has been destroyed and then recreated, its Cohesity host entry from the original Pure FlashArray registration is retained on the array. This causes IQN collision when registering the same Pure FlashArray with the recreated Cluster. IQN collision causes Cohesity backups of the array to fail.

  Workaround - Follow the below steps:

  - Run the following command on any Node in the Cluster to obtain the Cluster's IQN: `allssh cat /etc/iscsi/initiatorname.iscsi`
  - On the Pure FlashArray, remove the host entry that matches the obtained IQN.
  - Unregister and re-register the Pure FlashArray with the Cohesity cluster.
- If a Cluster upgrade is running during a Pure Storage backup job, the job may fail. The job will succeed on the next retry attempt.
- Cohesity does not currently support Pure Storage PODS or Active Cluster but you can register the individual arrays within the Active Cluster separately as standalone arrays.
- There is a limitation on the number of hosts that can be accommodated in the Pure FlashArray. When this limit is reached, it may lead to iSCSI communication failures during the backup processes.

  Error Message:

```
Encountered HTTP error kBadRequest with GET request for /host/cohesity-
1732787028580538/volume [Host does not exist.]
```

Workaround

Follow the below steps:

- Create a free slot for a new host entry in the Pure FlashArray.
- Choose one of the following options:
  - Restart the Backup Manager service using the following CLI commands:

    ```
    iris_cli cluster stop service-names=magneto
    ```

    ```
    iris_cli cluster start service-names=magneto
    ```

  - Alternatively, you can wait for an hour for the host creation process to complete.

- The following lists the considerations for a Pure Protection Group during the backup and recovery process:
  - Cohesity does not support the protection of empty Pure Protection Groups.
  - The name of the Pure Protection Group and Pure volumes created during recovery will be based on the specified suffix and prefix.
  - If the newly created PPG or a volume has the same name as that of the PPGs or volumes belonging to an existing PPG, then the recovery jobs will fail.
  - Volume Group will not be created. It should pre-exist. The recovery volumes will be created in the same volume group in which the original volume was present. If the volume group is not present at the time of recovery then the recovery job will fail.

Register or Edit a Pure FlashArray

You can register Pure FlashArray on the Cohesity cluster to protect the Pure Storage volumes. You can also register Pure FlashArray to leverage storage array-based snapshots for VMDKs.

Prerequisite

The Pure FlashArray needs to have a Network Interface Card (NIC) and iSCSI discovery IP configured on the array.

Register Pure Storage

Before registering or editing a Pure FlashArray, ensure it meets prerequisites, minimum permissions, and firewall port requirements in the Plan and Prepare for Pure FlashArray.

Registering a Pure FlashArray as a Storage array makes its LUN (both iSCSI and FC) content eligible to be protected on the Cohesity cluster.

After you complete the iSCSI configuration and have an iSCSI discovery set/enabled on the Pure Flash array, verify the iSCSI connectivity between the Cohesity cluster and the Pure Flash array. To verify the connectivity, you can manually log in to one of the configured iSCSI targets on the Pure Flash array from any one of the Cohesity cluster nodes.

Run the following command to check the iSCSI connectivity:

```
iscsiadm -m discovery -t sendtargets -p <ipaddress>
```

For additional iSCSI commands, see Additional iSCSI commands.

1. Navigate to **Data Protection > Sources**.
2. Click **Register > SAN**.
3. Select the source type as **Pure**.
4. Enter the information described below:

> **Note:** To leverage storage array-based snapshot for VMDKs, you must also register the corresponding VMware vCenter where FlashArray Datastore is mounted with the Cohesity cluster.

Field	Description
Management Host or IP Address	Specify the Management Host or IP Address of the Pure FlashArray to register.
Username	Specify a user with a Storage Admin role.
Password	Enter the password for the specified user.

5. Click **Register**

Additional iSCSI Commands

The table below lists the additional iSCSI commands:

Command	Description
iscsiadm -m node -T <targetname> -p <ipaddress> -l	Use this command to log in to a specific target.
iscsiadm -m node -T <targetname> -p <ipaddress> -u	Use this command to log out from a specific target.
iscsiadm -m node -T< targetname> -p <ipaddress>	Use this command to display the target information.

Command	Description
iscsiadm -m node -s -T <targetname> -p <ipaddress>	Use this command to display the statistics of a target.
iscsiadm -m session	Use this command to display all the current logged-in sessions.
iscsiadm -m discovery -o show	Use this command to display the iSCSI database regarding discovery.
iscsiadm -m node -o show	Use this command to display the targets to log into.
iscsiadm -m session -o show	Use this command to display the sessions logged into.

Related Topic

Ensure Adequate Privileges for Cohesity on the Source

Protect a Pure Volume or a Protection Group

A Protection Group that protects Pure Storage volumes captures Snapshots and these Snapshots are required for the following tasks:

- Recover volumes
- Replicate Snapshots
- Archive Snapshots

Before you begin

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Register or Edit a Pure FlashArray
- Create or Edit a Standard Policy
- Create or Edit Storage Domains

To Protect a Pure Volume or a Protection Group :

You can now protect Pure volumes or multiple volumes using a Pure Protection Group on the Cohesity cluster.

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **SAN**.

3. In the **New Protection** pop-up, you can either select a Volume or a Pure Protection Group:

   To Protect a Volume:

   Select the **Volume List** View ( ) on the top right to list the volumes. From the **Volume** list, select the volumes you want to back up.

   To Protect a Pure Protection Group:

   Select the **Pure Protection Group** View ( ) on the top-right to list the **Pure Protection Groups**. From the Pure Protection Group list, select the Pure Protection Group you want to back up.

   After selecting one of the above options, click the edit icon and provide the following information:

   1. **Add Objects**: From the **Objects** list, select the **Pure Protection Groups** you want to backup or click **Register Source** to register a new Pure Protection Group.
   2. **Protection Group**: Specify a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
   3. **Policy**: Select an existing policy or click **Create Policy** to create a new protection policy. For more information, see Create or Edit a Standard Policy.
   4. **Storage Domain**: Select the appropriate storage domain. The default time zone in the Start Time field is the browser's time zone. You can change the time zone of the job by selecting a different time zone. For more information, see Manage Storage Domains.
   5. To protect Pure Protection Group by specifying all the options, click **More Options**.

Advance Settings	Description
End Time	Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of jobs exceeds the ability to process jobs, this will be the priority of implementation of the jobs: a. High b. Medium c. Low
QoS Policy	Select an appropriate quality of service (QoS) policy. > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
Retain on Pure Storage Array	Enter the number of the most recent snapshots to retain on the Pure FlashArray. Retaining the last 5 snapshots is the default. You can change this number or choose to retain All Snapshots. To take incremental snapshots, this value must be at least 1.

Advance Settings	Description
Pre & Post Scripts	Toggle Enable to run scripts on the Unix server. When creating a Protection Group, you can configure: • **Pre Scripts**: Scripts that will be executed before the protection run. • **Post Backup Scripts**: Scripts that will be executed after the protection run. • **Post Snapshot Scripts**: Scripts that will be executed after the snapshot has been created on the source. If enabled, the scripts are executed during the protection run as configured. You can distinguish between the scripts that are run on the different objects (entities) using the COHESITY_BACKUP_ENTITY environment variable. The Cohesity cluster populates this environment variable with the name of the backed object (or entity). For more information, see Configure Pre & Post Scripts. For more information, see Configure Pre & Post Scripts.
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Quiet Times	Available only if the selected policy has at least one quiet time period. Toggle it on to specify that all currently executing protection runs should abort if a quiet time period specified for the Protection Group starts. If the option is toggled off, then after a protection run starts, it continues to execute even when a quiet time period specified for this Protection Group starts. However, a new protection run will not start during a quiet time period.

Advance Settings	Description
Alerts	Select one or more of the following **Alerts On** settings if you want alerts to be created for the following triggers: • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • Success: Create an informational alert when the Protection Group completes successfully. Emails are not sent when informational alerts are created. To send emails when alerts are triggered for the Protection Group, in the **Email Recipients** field, add the email addresses of the recipients.
Pause Future Runs	Toggle on this option to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
Description	Enter a brief description about the Protection Group.

6. Click **Protect**.

Recover a Pure Volume or a Protection Group

You can recover the Pure Protection Groups backed up onto the Cohesity cluster.

To recover Pure Volumes or a Protection Group follow the steps listed below:

1. Select **Data Protection** > **Recoveries**.
2. Click on **Recover** located at the top-right of the page and then select **SAN**.
3. Search for the Pure Protection Group that contains the snapshots to recover from by entering characters of the Pure Protection Group name, a volume name, or any constituent volume (which is part of a PPG) by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, Storage Domain, and date.
4. From the search results list, select the Pure Protection Group you want to recover.

   By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, mouse over the snapshot from the selected list, and then click the **Edit** icon.

   By default, the latest snapshot is selected for recovery. You can also view and select the older snapshots by clicking the **Protection Run** drop-down. When you select a snapshot, the volumes associated with that snapshot are displayed.
6. After you select the snapshot to be recovered then click **Select Recovery Point**.
7. From the **Recover To** options, select **Original Location** to recover the volume to the same Pure FlashArray or select **New Location** to recover the volume to a new Pure FlashArray.
8. In the **Rename** field, you can specify the name of the recovered Pure Protection Group by adding prefixes and/or suffix strings to the name of the original Pure Protection Group.
9. **Task Name**: Change the default name of the recovery task.
10. Click Start **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

HPE Alletra Storage MP

Cohesity now supports using storage array snapshots for VMware backups on the HPE Alletra MP storage array. This integration allows Cohesity to leverage the native snapshot capabilities of the storage array, resulting in more efficient and reliable backups of VMware VMs.

Register or Edit HPE Alletra Storage MP

You can register the HPE Alletra Storage MP array to leverage storage array snapshot- for VMware VM based backups.

Registering a Nimble Storage as a Storage array makes its LUN (both iSCSI and FC) content eligible to be protected on the Cohesity cluster.

After you complete the iSCSI configuration and have an iSCSI discovery set/enabled on the Nimble array, verify the iSCSI connectivity between the Cohesity cluster and the PE Alletra Storage MP array. To verify the connectivity, log into any of the Cohesity cluster nodes and run the following command:

```
iscsiadm -m discovery -t sendtargets -p <ipaddress>
```

For additional iSCSI commands, see Additional iSCSI commands.

Perform the following steps to register the HPE Alletra Storage MP Array:

1. Navigate to **Data Protection > Sources**.
2. Click **Register > SAN**.

3. Select the source type as **HPE Alletra Storage MP**.

> **Note:** You can register the HPE Alletra Storage MP by selecting the source type as HPE Alletra Storage MP.

4. Enter the following information:

> **Note:** You must also register the corresponding VMware vCenter where Storage Array Datastore is mounted with the Cohesity cluster.

Field	Description
Management Host or IP Address	Specify the Management Host or IP Address of the HPE Alletra Storage MP to register.
Username	Specify the user name used for accessing the HPE Alletra Storage MP. Ensure that the user being used has the Edit role.
Password	Enter the password for the specified user.

5. Click **Register**.

   When registration finishes, a new entry appears in the Sources table.

Additional iSCSI Commands

The table below lists the additional iSCSI commands:

Command	Description
iscsiadm -m node -T <targetname> -p <ipaddress> -l	Use this command to log in to a specific target.
iscsiadm -m node -T <targetname> -p <ipaddress> -u	Use this command to log out from a specific target.
iscsiadm -m node -T< targetname> -p <ipaddress>	Use this command to display the target information.
iscsiadm -m node -s -T <targetname> -p <ipaddress>	Use this command to display the statistics of a target.
iscsiadm -m session	Use this command to display all the current logged-in sessions.

Command	Description
iscsiadm -m discovery -o show	Use this command to display the iSCSI database regarding discovery.
iscsiadm -m node -o show	Use this command to display the targets to log into.
iscsiadm -m session -o show	Use this command to display the sessions logged into.

After registering the snapshot provider, to leverage the VMware snapshots from HPE Alletra Storage MP, you must enable the **Leverage Storage Snapshots for Data Protection** field and select **Storage Array/NAS** when creating the protection group for VMware VMs. For more information, see create a Protection Group for VMware.

Using the **Leverage Storage Snapshots for Data Protection** option with the **Leverage SAN Transport for Data Protection** option is not recommended.

Chapter 19: Hadoop

The Cohesity Hadoop DataProtect solution natively integrates with Cloudera distributions to provide consistent backup and recovery for Hadoop deployments. The Cohesity solution supports Hive, HBase, and HDFS data sources.

Based on the policy settings, the Cohesity NoSQL and Hadoop Service initiates a backup of the HDFS, Hive, or Hbase. The first backup is a full backup and everything after that is an incremental backup only. The service scales horizontally based on the available compute and memory resources on Cohesity. Multiple backup and recovery streams run in parallel to reduce the backup and recovery window. The necessary Cohesity Views are automatically generated where the backups are going to be written. Communication between Cohesity and Hadoop is managed by the NoSQL and Hadoop Service as described in the section below. See the following topics for details:

Core components 3985
Hadoop Features 3986
Plan and Prepare for Hadoop Protection 3986
Register a Hadoop Source 3992
Backup Hadoop 3994
Recover Hadoop 4003

Core components

The following lists the core components of the Cohesity Hadoop solution:

- Physical or virtual Hadoop clusters.
- Cohesity Data Cloud.
- Cohesity Compute Nodes.
- Cohesity NoSQL and Hadoop Service running on Cohesity Compute Nodes.
- Optional. Archive or tiering target at your data center or cloud provider of choice.

As illustrated in the below figure, you can integrate Cohesity with an existing Hadoop infrastructure by installing the Cohesity NoSQL and Hadoop Service on Compute Nodes. Once the service is installed and running, you can create protection policies, Protection Groups, and replication schedules for the Hadoop file system and databases. Once the data is protected, it can be archived, sent to tape, and or replicated to another instance running in the cloud or on-premises as shown in the below image.

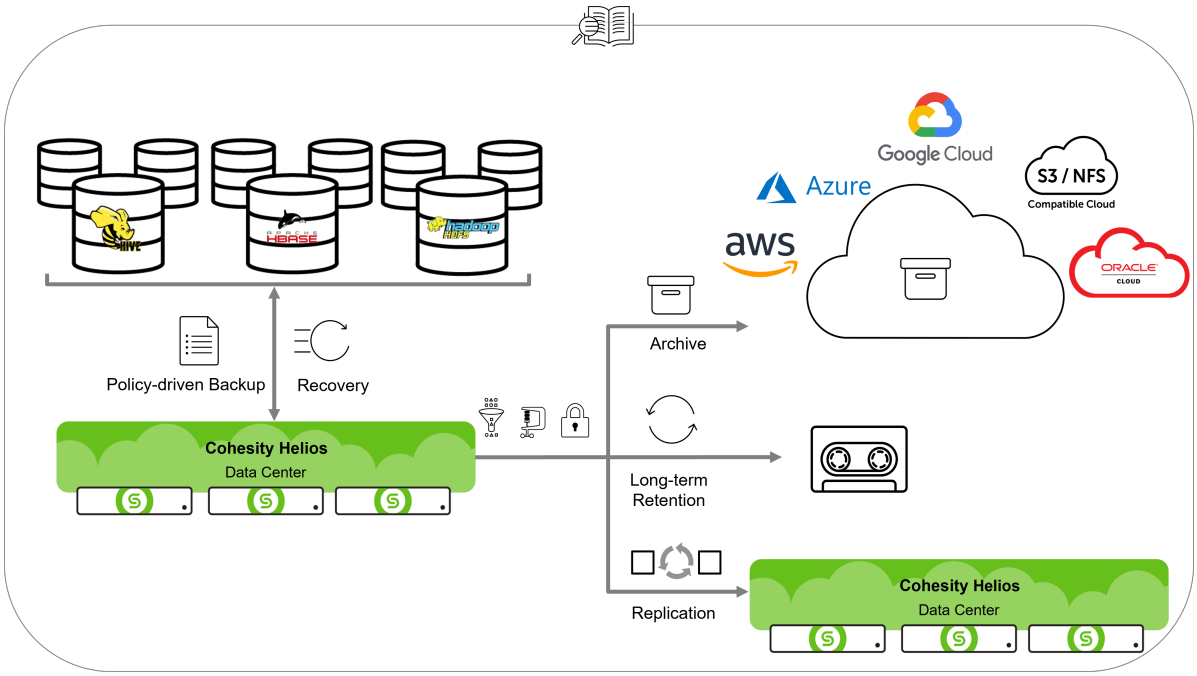

Hadoop Features

Following are the key Hadoop-specific features of the solution:

- Support Cloudera distributions of Hadoop.
- Support Hive, Hbase, and HDFS.
- Incremental-forever backups - after the first full backup, all subsequent backups are incremental only forever.
- Granular backup and recovery - backup and recover individual tables, databases, or files.
- Automatically handle topology changes seamlessly (node addition, removal, IP changes, etc.)
- Support Kerberos enabled clusters.
- Support SSL-enabled clusters.

Plan and Prepare for Hadoop Protection

The following topics provide guidance that helps you get ready to implement the Hadoop DataProtect Solution:

Supported Versions

For more information, see NoSQL and Hadoop Platforms.

> **Note:** Users have to create new workflows when upgrading their Hadoop cluster to Cloudera Data Platform 7.x versions. For more information, see Cloudera Data Platform Upgrade and Migration.

Prerequisites

Make sure the following prerequisites are met before using Hadoop DataProtect Solution.

- The NoSQL and Hadoop Service is in Running status. See NoSQL and Hadoop Service for details.
- The hostname to IP address and reverse resolution of all the hadoop primary nodes and Kerberos server (KDC) is working from all the NoSQL and Hadoop Service PODs.
- Refer to the NoSQL and Hadoop Service - Post Installation Procedures to get a general idea about the following operations:
  - Manage SSL enabled NoSQL primary clusters
  - Work with Kerberos Security enabled primaries for data protection
- Kerberos-Enabled Hadoop
- SSL-Enabled Hadoop
- Kerberos + SSL-Enabled Hadoop

Kerberos-Enabled Hadoop

The following section describes the steps you need to perform for adding a Kerberized Hadoop data repository.

Adding Kerberized Hadoop data repository:

This section describes the steps that must be performed on the primary source cluster prior to successfully add new Kerberized HDFS, HBase, or Hive data repositories. These steps are also needed when you are adding a new data repository with a new set of Kerberos principals.

> **Note:** At the end of these steps, service restart is required on both Primary and NoSQL Hadoop clusters. Protection runs will be killed. Please plan this activity when no protection runs are in progress.

1. Add the new realm to NoSQL and Hadoop Services as mentioned in the section Work with Kerberos Security enabled primaries for data protection.
2. Make changes to the Primary cluster's HDFS, Hive and HBase configuration.

   Let us assume that you plan to use the following as backup/recovery user principals on your primary Hadoop cluster:

   ```
   imanis_hdfs@PRIMARY.REALM - Kerberos principal used for HDFS workflows
   imanis_hbase@PRIMARY.REALM - Kerberos principal used for HBase workflows
   imanis_hive@PRIMARY.REALM - Kerberos principal used for Hive workflows
   ```

   These Kerberos principals map to corresponding Unix users. For example, imanis_hdfs@PRIMARY.REALM maps to Unix user imanis_hdfs on your primary cluster nodes.

   Similarly, imanis_hbase@PRIMARY.REALM maps to imanis_hbase and imanis_hive@PRIMARY.REALM maps to imanis_hive Unix user. All these three users should be added to a group, for example, imanis_supergroup in all cluster Nodes of your primary hadoop cluster.

   The steps for performing the changes on the primary cluster differ from vendor to vendor. The below example is applicable for Cloudera 6.0. For other vendors and versions please refer to the vendor-specific documentation.

   - Grant superuser privileges in the primary Hadoop cluster to the Kerberos principals used for backup/recovery.
     a. In **Cloudera Manager Admin Console**, click **HDFS** service > **Configuration** tab.
     b. Search for the dfs.permissions.superusergroup property and add imanis_supergroup to this property.
     c. Save the changes and restart HDFS service.
   - Set hbase.superuser to imanis_hbase
     a. In **Cloudera Manager Web UI**, click **HBase** > **Configuration** tab.
     b. Search for hbase.superuser and set it to imanis_hbase.
     c. Deploy the changed configuration and restart HBase service.
   - On the primary Hadoop cluster, add hbase.superuser property to the hbase-site.xml if not already present.

```
<property>
<name>hbase.superuser</name>
<value>imanis_hbase</value>
</property>
```

- On the primary Hadoop cluster, add dfs.webhdfs.enabled property in the hdfs-site.xml file:

```
<property>
<name>dfs.webhdfs.enabled</name>
<value>true</value>
</property>
```

SSL-Enabled Hadoop

The following section describes the tasks you need to perform before adding SSL-enabled Hadoop as a source data repository in the Cohesity cluster.

Before adding SSL-enabled Hadoop, do the following:

1. Consult your Hadoop system admin to get the SSL certificates of all nodes in your production Hadoop system.
2. Import SSL certificates from your primary cluster into Cohesity by following the steps documented in Manage SSL enabled NoSQL primary clusters .
3. After importing the SSL certificates(step #2), run the following command from any one node of the Cohesity cluster to verify if your Hadoop primary is accessible from the NoSQL and Hadoop Service node.

   Log in to one of the Cohesity cluster nodes and run the following:

   ```
   $ kubectl.sh get pods --all-namespaces   | grep  imanis-nodemanager-0
   athena-104    imanis-nodemanager-0                   1/1        Running
           3       41h
   ```

   Log in to imanis container using the following command:

   ```
   $ kubectl.sh exec -it imanis-nodemanager-0   -n athena-104 /bin/bash
   ```

   Check if you can reach the primary namenode host from Imanis container using the following command:

   ```
   $ ping <primary_namenode_hostname>
   hdfs dfs -ls swebhdfs://<primary_namenode_hostname>:50470/
   ```

4. Go ahead and add this SSL-enabled Hadoop as a data repository in Cohesity.

Kerberos + SSL-Enabled Hadoop

The following section describes the tasks you need to perform before adding Kerberos + SSL-enabled Hadoop as a source data repository in Cohesity.

Before adding Kerberos + SSL-enabled Hadoop, do the following:

1. First, follow the steps mentioned in the Kerberos-Enabled Hadoop section for Kerberos-related configuration required for adding Kerberized Hadoop primary.
2. Then, consult your Hadoop system admin to get the SSL certificates of all nodes in your production Hadoop system.
3. Refer to SSL-Enabled Hadoop.
4. Log in to one of the Cohesity cluster nodes and run the following:

```
$ kubectl.sh get pods --all-namespaces   | grep  imanis-nodemanager-0

athena-104    imanis-nodemanager-0                   1/1        Running
        3       41h

$ kubectl.sh exec -it imanis-nodemanager-0   -n athena-104 /bin/bash
```

5. Run the following command to log in with user hdfs on the Cohesity container:

```

$ kinit -kVt $INSTALL_DIR/conf/talena.keytab imanis_hdfs@PRIMARY.REALM
```

> **Note:** If the Kerberos configuration steps are correctly performed, this command should succeed. Use klist command to verify that imanis_ hdfs@PRIMARY_REALM is the active kerberos principal.

6. Run the List command (ls) as the following:

```
$ hdfs dfs -ls swebhdfs://<namenode>:50470/
```

7. Execute the kdestroy command for imanis_hdfs@PRIMARY.REALM principal.

```
$ kdestroy
```

8. Go ahead and add this Kerberos + SSL-enabled Hadoop as a data repository in Cohesity.

Considerations

Hadoop HDFS

- Unicode characters in hdfs file and directory names are not supported.
- Special characters % ; + (percent semicolon plus) in hdfs file and directory names are not supported.
- Access Control Lists (ACLs) are not supported.

Hive

- Backup and Recovery across versions/vendors is not supported. For example, performing a data backup from CDH 5.8.2 and then recovering it to CDH 5.12.2 might cause data loss.
- Special characters, except underscore(_), are not supported for database and table names.
- Access Control Lists (ACLs) is not supported.
- For Cloudera Data Platform 7.x, backup and Recovery of hive databases: sys and information_schema is not supported.

HBase

- Backup and Recovery across versions/vendors is not supported. For example, performing a data backup from CDH 5.8.2 and then recovering it to CDH 5.12.2 might cause data loss.
- Access Control Lists (ACLs) is not supported.

Firewall Ports

For more information on the required firewall ports that are to be open to enable communication between Cohesity and Hadoop DataProtect Solution, see Open-Source Software Platform.

Minimum Permissions

This section describes the minimum set of permissions needed for the Hadoop DataProtect Solution to communicate with the source.

For kerberized setup:

User principal provided in UI must have super user privileges of production hadoop cluster and must have full access to the production HDFS/HIVE/Hbase system.

For non-kerberized setups:

NoSQL and Hadoop Service user (hdfs) must have full access to the production HDFS/HIVE/Hbase setups.

Register a Hadoop Source

Register Hadoop Applications as Sources

To configure Hadoop applications for data protection in Cohesity, you need to register them as sources. All three applications - HDFS, Hbase and Hive - can be registered together.

Before you register Hadoop as a source, ensure that you have reviewed and completed the instructions provided in the Plan and Prepare for Hadoop Protection topic.

To register Hadoop as a source:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** on the top-right of the page and then select **Hadoop**. The **Discover Hadoop** dialog is displayed.
3. Under the **Discovery** tab, enter the required details. Refer the table below for the field descriptions.
    - Choose either **Password** or **Private Key** as the authentication type. If you choose **Password**, enter the SSH Username and Password in the respective fields. Else if you choose **Private Key**, enter the Login(SSH username), SSH private key and passphrase in the respective fields.

Field	Description
Host	IP or hostname of any host from which the configuration file can be read.
Hadoop Configuration Directories	Directory path for the Hadoop configuration file.
Distribution	Select the Hadoop distribution for this cluster.
Hadoop Version	Enter the Hadoop version for this cluster.
HDFS Connection Protocol	Select the protocol for connecting to HDFS. The available options are: • Web HDFS • DFS **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

4. Click **Register**. The **Register Sources** tab is displayed.
5. If the primary source is kerberized, enter the Kerberos principal and edit the application details that are identified by the system. Below table contains the fields and their descriptions.

Field Descriptions

Field	Description
HDFS	
namenode Address	HDFS Namenode IP or hostname.
webhdfsPort	HDFS WebHDFS port.
HBase	
Zookeeper Quorum	Zookeeper Quorum for this HBase.
HBase Data Root Directory	Data Root Directory for this HBase
Kerberos Principal	The kerberos principal to be used to connect to this Hbase source.
Hive	
Metastore Address	MetastoreAddress for this Hive.
Metastore Port	MetastorePort for this Hive.
Kerberos Principal	The kerberos principal to be used to connect to this Hive source

6. Click **Register**.

If the registration is successful, the registered cluster appears in the list on the **Sources** page. After registration, each application - HDFS, Hive and HBase - functions as an independent source.

Reconfigure HDFS, Hive and HBase Sources

You can individually reconfigure the registered HDFS, Hive and Hbase sources and update the configuration details that you had provided during the registration process with the latest configuration.

To reconfigure HDFS/ Hive/ HBase:

1. Navigate to **Data Protection > Sources**.
2. Browse to the required HDFS/ Hive/ HBase source you want to edit and click ⋮ > **Reconfigure**. The respective **Reconfigure** dialog is displayed.
3. Edit the authentication details and other configurations as required. Refer Field Descriptions for more details.

   > **Note:** Performing step 4 leads to rediscovering the source. The source ID, protection runs and Protection Groups that are already created are retained.

4. To update the details, click **Update**.

Refresh HDFS, Hive and HBase Source

You can manually refresh the HDFS, Hive and HBase source configuration and fetch the latest changes on the source.

To refresh the Source Configuration:

1. Navigate to **Data Protection > Sources**.
2. Click ⋮ next to the required HDFS/ Hive/ HBase source and select **Refresh**.

Unregister HDFS, Hive and HBase Source

If you want to stop taking backups of the HDFS/ Hive/ HBase source, you can unregister the source from Cohesity.

> **Note:** Make sure that you delete all the Protection Groups and backup snapshots associated with the source.

To unregister the HDFS/ Hive/ HBase Source:

1. Navigate to **Data Protection > Sources**.
2. Click ⋮ next to the HDFS/ Hive/ HBase source and select **Unregister**.

Backup Hadoop

Cohesity provides a simple sequence of tasks that constitute the procedure for backing up your HDFS/ Hive/ HBase source. Once HDFS/ Hive/ HBase is registered as a source on Cohesity, you can start configuring your backups.

HDFS/ Hive/ HBase objects are protected using a Protection Group and a Protection Policy. A Protection Group enables you to select the objects that you want to back up. It uses the

schedules and settings defined in the Protection Policy to determine when and how backups are captured and replicated.

A Protection Policy is a reusable list of settings that define the type of backup, backup schedule and backup snapshot retention. Also, for Disaster Recovery, you can add replication for the Protection Group using the Protection Policy. For detailed instructions on how to create a Protection Policy, refer Create or Edit a Standard Policy.

Protect HDFS Objects

To protect a HDFS object:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Hadoop > HDFS**. The **New Protection** page is displayed.
3. Select the required **Source** from the dropdown.
4. To protect all the objects in the selected data source, select **Protect Entire Data source**.
5. To protect specific paths in the selected data source, select **Protect Specific Paths**.

   Enter the paths to include and exclude and click + Add to add the paths.

   > **Note:** Specific path fields accept absolute paths and regular expressions.
6. Click **Save Selection**.
7. Under **Protection Group**, enter a protection group name. Existing protection groups cannot be used while creating a new HDFS protection.
8. Under **Protection Policies**, select an existing protection policy or create a New Policy if the available policies do not meet your schedule and retention needs.
9. Select the appropriate **Storage Domain**. For more information, refer Manage Storage Domains.
10. After selecting a Storage Domain, if you need to change any of the additional settings on the HDFS Protection page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
11. Click **Protect**.

Your new Protection Group is now active and running.

Protect Hive Objects

To protect a Hive object:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Hadoop > Hive**. The **New Protection** page is displayed.
3. Select the required **Source** from the dropdown.
4. Select the required objects to protect such as hive databases, tables, or views.

   > **Note:** If the number of objects exceeds the threshold (10000 objects), the UI will display only the hive database objects and it will not show the tables or views.

5. To auto protect an object, click adjacent to it. For more details, see Enable or Disable Auto Protect.
6. Click **Save Selection**.
7. Under **Protection Group**, enter a protection group name. Existing protection groups cannot be used while creating a new Hive protection.
8. Under **Protection Policies**, select an existing protection policy or create a New Policy if the available policies do not meet your schedule and retention needs.
9. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
10. After selecting a Storage Domain, if you need to change any of the additional settings on the HDFS Protection page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
11. Click **Protect**.

Your new Protection Group is now active and running.

Protect HBase Objects

To protect a HBase object:

1. Navigate to **Data Protection > Protection**.
2. Click **Protect** on the top-right of the page and then select **Hadoop > HBase**. The **New Protection** page is displayed.
3. Select the required **Source** from the dropdown.
4. Select the required objects to protect.
5. To auto protect an object, click adjacent to it. For more details, see Enable or Disable Auto Protect.

6. Click **Save Selection**.
7. Under **Protection Group**, enter a protection group name. Existing protection groups cannot be used while creating a new HBase protection.
8. Under **Protection Policies**, select an existing protection policy or create a New Policy if the available policies do not meet your schedule and retention needs.
9. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
10. After selecting a Storage Domain, if you need to change any of the additional settings on the HBase Protection page, scroll down and click **Edit** on the right. For more information, see Additional Settings.
11. Click **Protect**.

Your new Protection Group is now active and running.

Additional Settings

Field	Description
End Date	Optional. Click and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Concurrency	Maximum number of backup streams that will be created to exchange data with the cluster. By default, maximum scaling is capped at 16.

Field	Description
Bandwidth Throttling	Maximum network bandwidth that each backup stream can use for exchanging data with the cluster.
Indexing	• HDFS - Files/directories are indexed. • Hive, HBase - Databases/tables are indexed. **Note:** The indexing option for HDFS is not mandatory since the default indexing policy is sufficient. For Hive and Hbase, the indexing options are not displayed since the default indexing policy is selected automatically.
Alerts	Optional. Select one or more of the following settings if you want alerts to be created for the following triggers: • Success: Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • Failure: Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • SLA Violation: Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts. You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection group.
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups. However, if the number of protection groups exceeds the ability to process protection groups, this will be the priority of implementation of the protection groups: • High-priority protection groups • Medium-priority protection groups • Low-priority protection groups.

Field	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • Incremental: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. • Full: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks).
Pause Future Runs	Once enabled, no protection runs will be scheduled.
Include & Exclude	You have the option to include or exclude objects for protection in addition to the currently selected objects. To enable the Include and Exclude options: • Click icon. • Toggle on Enable Partial Protection. • Select the include or exclude option. • Click Select File. **Note:** You can upload a file with a maximum size of 30 MB. This file should contain fully qualified names of hive databases, tables, or views, with each name separated by a new line. For example: • mydb.table1 • empdb.table1 • empdb.table2 • empdb.view1
Description	Optional. Enter a brief description about the Protection Policy.

Manage Backup Configuration

A backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

Add Objects to an existing Protection Group

You can add new objects to the Protection Group created for HDFS/ Hive/ HBase protection. The newly added objects are backed up from the next protection run.

To add objects to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for HDFS/ Hive/ HBase.
3. Click ⋮ next to the Protection Group and select **Edit**. The **Edit Protection** page is displayed.
4. For HDFS, click + and enter the paths.
5. For Hive and HBase, click ✎ next to **Objects**.
6. Make the required changes. For more information, see Protect HDFS Objects, Protect Hive Objects and Protect HBase Objects.
7. Click **Protect**.

Enable or Disable Auto Protect

The Auto Protect option enables you to automatically protect the objects that are added to the Hive/ HBase source. You can auto protect either the entire cluster or a Keyspace.

When the Auto Protect option is enabled:

- All the objects that are added to the Hive/ HBase source in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the objects that are removed from the Hive/ HBase source in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backups are preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly added objects, in the Hive/ HBase source, to the Protection Group.
- Remove existing objects, that are removed from the Hive/ HBase source, from the Protection Group. The existing backups are preserved until the snapshot expires.

To enable or disable Auto Protect:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Hive/ HBase. For more information, see Protect Hive Objects and Protect HBase Objects.
3. Click ⋮ next to the Protection Group and select **Edit**.
4. On the **Edit Protection** page, in the **Object** area, click **Edit**.

5. In the pop-up window , toggle the to enable or disable Auto Protect.
6. Click **Protect**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group in the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Hive/HBase. For more information, see Protect Hive Objects and Protect HBase Objects.
3. Click ⋮ next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Protect**.

Start an On-demand Protection run

You can manually start a run for the Protection Group that you have created to backup Hive/HBase objects. It provides you with options to:

- Back up all the objects in the Protection Group.
- Select an incremental or full backup.
- Add replication and archive.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Hive/HBase.
3. Click ⋮ next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    1. The **Backup** drop-down list is defaulted to **Backup all Objects in the Protection Group**
    2. In the **Backup Type** drop-down list, select the backup type.
    3. If required add details for Replication and Archival.
5. Click **Backup Now**.

Pause Future Protection Runs

You can pause future protection runs of the Protection Group. However, if a protection run is in progress, it continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Hive/ HBase.
3. Click ⋮ next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered.

To resume future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had paused.
3. Click ⋮ next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the snapshots(local, archived, and replicated) that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete the Protection Group

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived snapshots generated by the Protection Group and reclaim the space utilized by the Protection Group.

> **Note:** The storage space reclaimed by deleting the Protection Group takes some time to reflect.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for Hive/ HBase.
3. Click the ⋮ next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
    - Delete Protection Group Only
    - Delete Protection Group and Snapshots

Recover Hadoop

You can recover the protected HDFS/ Hive/ HBase Objects using Cohesity. You can recover to the same location or an alternate location.

> **Important:**
> - Recovering at the protection group level is supported for HDFS/ Hive/ HBase.
> - Recovering multiple protection groups in a single recovery task is not supported.
> - Recovering of items in the same protection group but from different recovery points is not supported in the same recovery task. You need to create different recovery tasks for different recovery points.
> - Recovering objects from more than one protection group in a single recovery task is not supported.

Recover HDFS Objects

To recover HDFS objects you can search through the backups using the protection group, keyspaces or table names for the objects you plan to recover, and then start the recovery.

To recover HDFS objects:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Hadoop > HDFS**. The **New Recovery** dialog is displayed.
3. Search for the HDFS protection groups or objects that need to be recovered.

    > **Note:** Wildcards are supported in the search.

4. Accept the latest Recovery Point or click ✎ and select a different recovery point.
    - Snapshots: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
5. Click **Next: Recover Options**. The recovery details are displayed.
6. Under **Recover As** select **Replace Original** or **New Object**.
    - **Replace Original** - Replace the object(s) back to the source from where it was backed up.
    - **New Object** - Specify the target for the recovered object. Use the toggle to **Overwrite existing objects with the same name**. A warning is displayed, enter **YES** to proceed.
7. Under **Recovery Options**, select the required options for the recovery. For more information, see Recovery Options.
8. Click **Start Recovery** to start the recovery operation.

Recover Hive Objects

To recover Hive objects you can search through the backups using the protection group, keyspaces or table names for the objects you plan to recover, and then start the recovery.

To recover Hive objects:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Hadoop> Hive**. The **New Recovery** dialog is displayed.
3. Search for the Hive protection groups or objects such as hive databases, tables, or views. that need to be recovered.

    > **Note:** Wildcards are supported in the search.

4. Accept the latest Recovery Point or click ✎ and select a different recovery point.
    - Snapshots: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
5. Click **Next: Recover Options**. The recovery details are displayed.
6. Under **Recover As** select **Replace Original** or **New Object**.
    - **Replace Original** - Replace the object(s) back to the source from where it was backed up.

- **New Object** - Specify the target for the recovered object. Use the toggle to **Overwrite existing objects with the same name**. A warning is displayed, enter **YES** to proceed.

7. If you select **New Object**, enter a suffix to append to all the selected objects.
8. Under **Recovery Options**, select the required options for the recovery. For more information, see Recovery Options.
9. Click **Start Recovery** to start the recovery operation.

Recover HBase Objects

To recover HBase objects you can search through the backups using the protection group, keyspaces or table names for the objects you plan to recover, and then start the recovery.

To recover HBase objects:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** on the top-right of the page and then select **Hadoop > HBase**. The **New Recovery** dialog is displayed.
3. Search for the HBase protection groups or objects that need to be recovered.

   > **Note:** Wildcards are supported in the search.

4. Accept the latest Recovery Point or click ✎ and select a different recovery point.
   - Snapshots: Select a different recovery point. You can change the time zone before or after selecting a new restore time.
5. Click **Next: Recover Options**. The recovery details are displayed.
6. Under **Recover As** select **Replace Original** or **New Object**.
   - **Replace Original** - Replace the object(s) back to the source from where it was backed up.
   - **New Object** - Specify the target for the recovered object. Use the toggle to **Overwrite existing objects with the same name**. A warning is displayed, enter **YES** to proceed.
7. If you select **New Object**, under **Recovery Options** in the **Rename** field, choose Bulk Rename to append a suffix to all the selected objects, or choose Rename Individual Objects.
8. Under **Recovery Options**, select the required options for the recovery. For more information, see Recovery Options.
9. Click **Start Recovery** to start the recovery operation.

Recovery Options

Field	Description
Concurrency	Maximum number of recovery streams that will be created to exchange data with the cluster. By default, maximum scaling is capped at 16.
Bandwidth Throttling	Maximum network bandwidth that each recovery stream can use for exchanging data with the cluster.
Task Name	Recovery task name.

Chapter 20: About the Remote Adapter

> **Note:** Cohesity Remote Adapter functionality is not supported if SmartFiles is not licensed on your Cohesity cluster.

Remote Adapter Protection Groups provide the ability to run scripts or executables on a remote AIX or Linux machine using protection schedules defined by a Cohesity Policy. Typically Remote Adapter Protection Groups invoke scripts that back up using third party tools such as Oracle's RMAN.

During the configuration of the Remote Adapter Job, a new View is created to store the results of the backup. After a script run, the Job captures a Snapshot of the View with the backup results. If the Policy associated with the Remote Adapter Job has a replication schedule and an archive schedule, the View Snapshot is copied to a remote Cluster and to an archive External Target.

For example, a Remote Adapter Job could invoke a script that backs up Oracle databases using Oracle's Recovery Manager (RMAN). The results of the backup are stored in a new View created during the configuration of the Remote Adapter Job. Cohesity provides an example Bash shell script that invokes an Oracle's RMAN to back up a database.

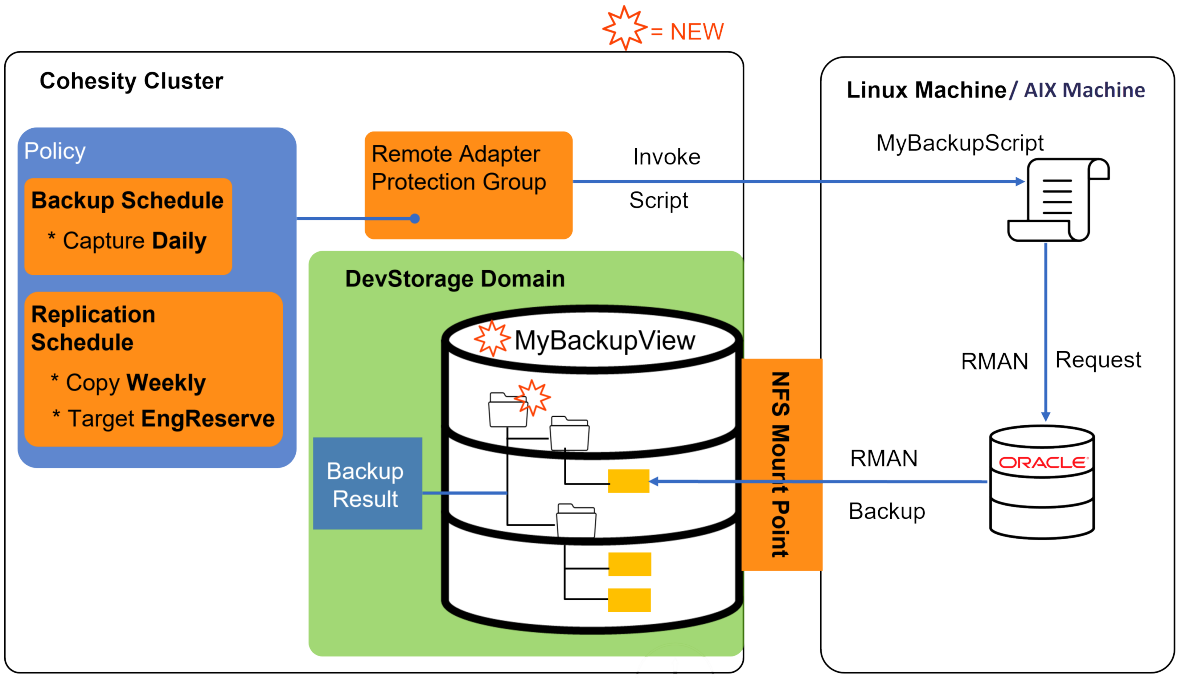

The Job captures a Snapshot of the View with the backup results.

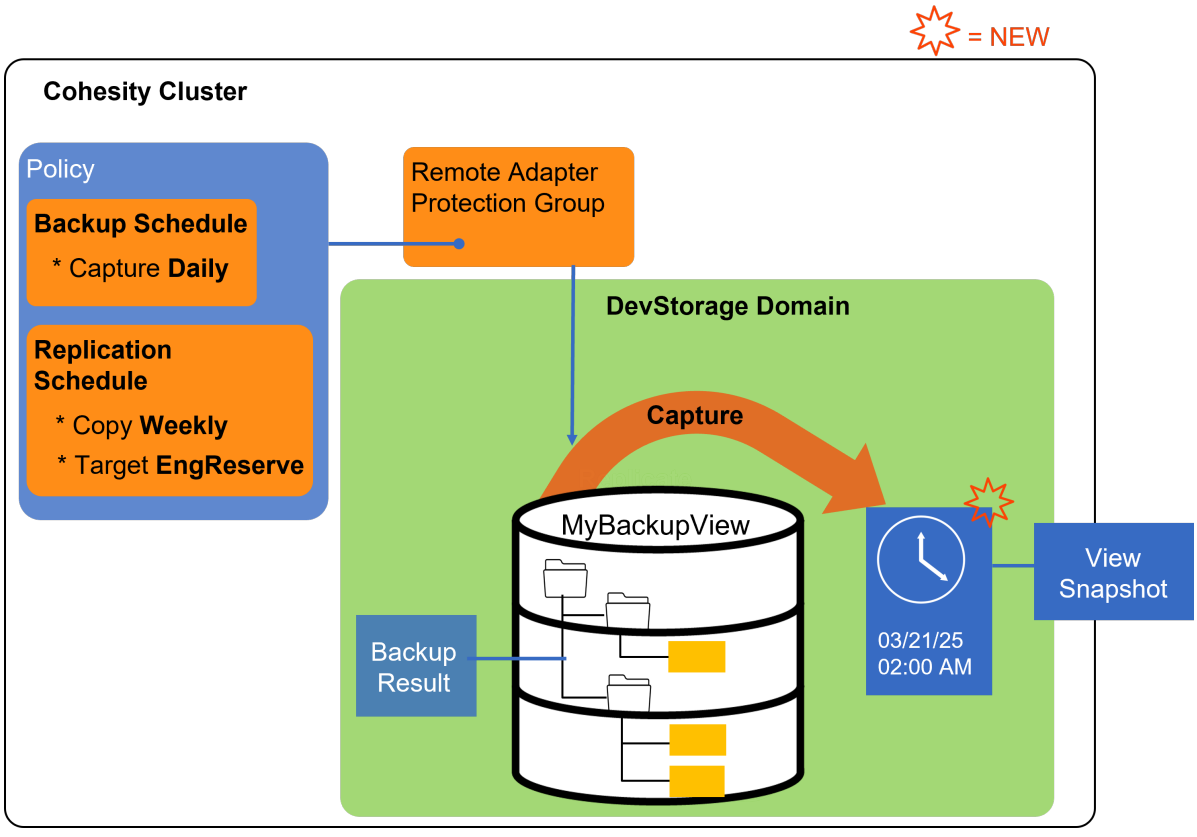

If the Policy associated with the Remote Adapter Job has a replication schedule, the Snapshot is copied to a remote Cluster ("EngReserve", below).

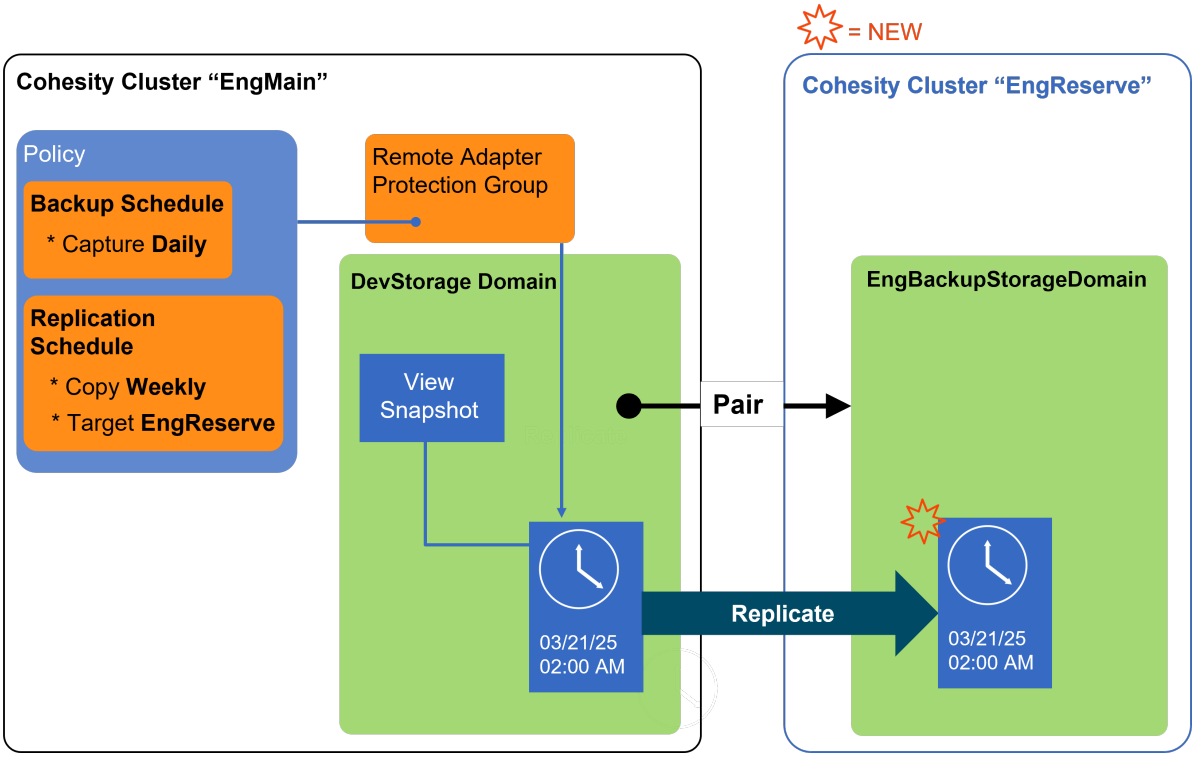

Replication occurs if the following conditions are met:

- The Policy associated with the Remote Adapter Protection Group has a replication schedule. In this example, the Policy defines a Weekly replication schedule.
- The Policy associated with the Remote Adapter Protection Group specifies a remote Cluster as the replication Target and there is a Storage Domain pairing between the capturing Cluster and this remote Cluster. In the example above, the Policy specifies the `EngReserve` Cluster as a Replication Target and there is Storage Domain pairing between `DevStorageDomain` on the `EngMain` Cluster and `EngBackupStorageDomain` on the `EngRerserve` Cluster.

If the Policy associated with the Remote Adapter Job has an archive schedule, the Snapshot is copied to an archive External Target.

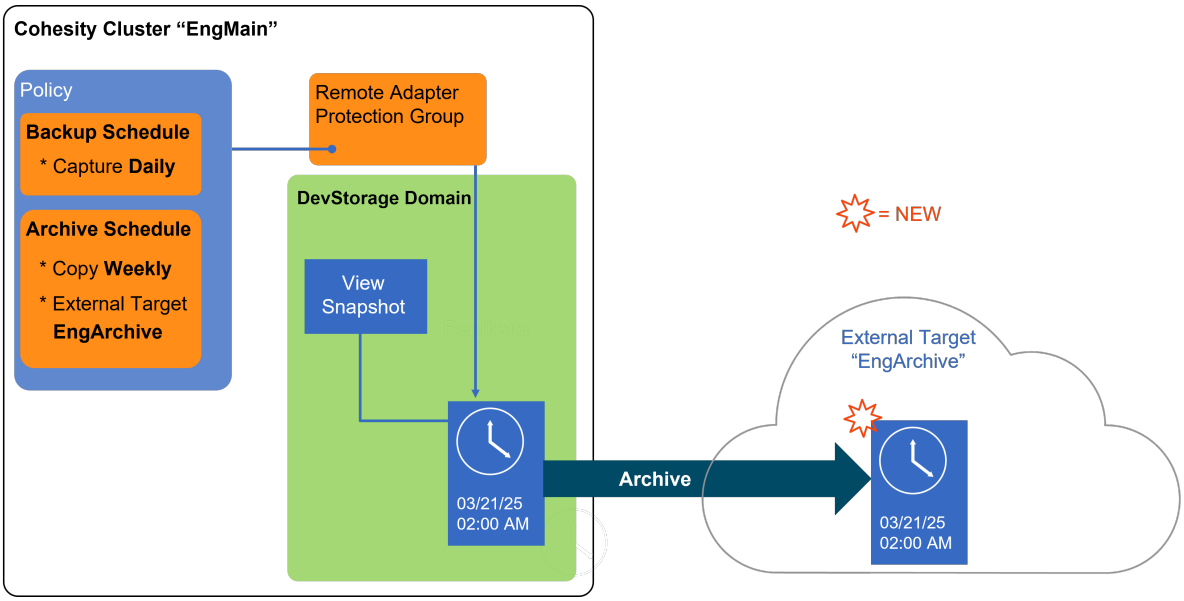

Archival occurs if the following conditions are met:

- The Policy associated with the Remote Adapter Protection Group has an archive schedule. In this example, the Policy defines a Monthly archive schedule.
- The Policy associated with the Remote Adapter Protection Group specifies an External Target as the archive Target. In this example, the Policy specifies the `EngArchive` External Target as an archive Target.

For more information about creating and configuring a Remote Adapter Protection Group, see Add or Edit a Remote Adapter Protection Group.

Recommendations for Dedupe-focused Backups with RMAN

For higher storage-side dedupe with RMAN backups, Cohesity recommends the following settings:

RMAN Parameter	Recommendation
FILESPERSET	1
Number of channels	Number of Cohesity nodes or greater
Compression	None
Encryption	None

Example of a full RMAN backup script

In the example below, four Oracle channels are created and each channel is directed to a separate Cohesity node leveraging the Cohesity dedupe:

```
DATABASE = “Name of DB” ORACLE_HOME=/u01/OracleHomes/db/product/dbhome_1
COHESITY_DIR_PATH =/mnt/oracle01 COHESITY_DIR_PATH2=/mnt/oracle02 COHESITY_
DIR_PATH3=/mnt/oracle03  COHESITY_DIR_PATH4=/mnt/oracle04

set ORACLE_SID=${DATABASE}
$ORACLE_HOME/bin/rman target=/ << EOF
RUN {
allocate channel C1 type disk FORMAT '$COHESITY_DIR_PATH/%U' MAXOPENFILES 1;
allocate channel C2 type disk FORMAT '$COHESITY_DIR_PATH2/%U' MAXOPENFILES
1;
allocate channel C3 type disk FORMAT '$COHESITY_DIR_PATH3/%U' MAXOPENFILES
1;
allocate channel C4 type disk FORMAT '$COHESITY_DIR_PATH4/%U' MAXOPENFILES
1;

BACKUP FULL DATABASE PLUS ARCHIVELOG FILESPERSET 1 TAG 'dedupe_focused';
RELEASE CHANNEL C1;
RELEASE CHANNEL C2;
RELEASE CHANNEL C3;
RELEASE CHANNEL C4;
}

EOF
```

Add or Edit a Remote Adapter Protection Group

Remote Adapter Protection Groups provide the ability to run scripts or executables on a remote AIX or Linux machine using protection schedules defined by a Cohesity Policy. Typically Remote Adapter Protection Groups invoke scripts that back up using third party tools such as Oracle's RMAN.

During the configuration of the Remote Adapter Job, a new View is created to store the results of the backup. After a script run, the Job captures a Snapshot of the View with the backup results. If the Policy associated with the Remote Adapter Job has a replication schedule and an archive schedule, the View Snapshot is copied to a remote Cluster and to an archive External Target.

The Protection Group captures snapshots of views and snapshots of views are required for the following tasks:

- creating a new view from a view snapshot
- replicating snapshots

- archiving snapshots

The Cohesity cluster supports backing up the following entities stored in the view:

- ordinary files
- directories
- symbolic links
- hard links

These entities can be captured in a view snapshot and restored to a new view. The backup and recovery of the following special files is supported: pipes, block devices, sockets and character devices.

Limitation:

If Remote Adapter jobs are running on a Cohesity cluster and the cluster is upgraded, the nodes are restarted and any scripts still running on the remote system are killed and restarted. This may happen repeatedly.

When you create a Remote Adapter Protection Group, you will create a View or use an existing View to store the result of the scripts. You must then mount the View on your system and configure your script to write the results to a directory in the mounted View. At the end of the Remote Adapter protection run, the Cohesity cluster captures a Snapshot of the View.

When you create a Remote Adapter Protection Group, you can select an existing Policy, Storage Domain and View. You can also create them while creating the Protection Group. However, you might find it easier to create them prior to creating the Protection Group, as described in the following topics.

- Create or Edit a Standard Policy
- Create or Edit Storage Domains
- Create Views Using a Predefined Template

To create a Remote Adapter Protection Group:

1. Navigate to **Data Protection > Protection** and do one of the following:
   - To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
   - To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.

   Here the Protection Group type is **Remote Adapter**.
2. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
3. **Linux Hostname or IP:** Specify the Hostname or the IP Address of the Linux system where the scripts are located and will be run.

4. **Username:** Specify the username on the Linux system that will be running the scripts.
5. **Cluster SSH Public Key:** Set up permissions on the Linux system to allow the Cohesity cluster to run the scripts remotely on the system.
   a. Click **Copy Key to Clipboard** to copy the SSH public key into the clipboard.
   b. Log in to the Linux system with the username specified earlier.
   c. In the home directory of the user, determine if there is a `.ssh` directory that contains a file called `authorized_keys`:

      ```
      ls -lr ~/.ssh
      ```

   d. Add the SSH key to the `authorized_keys` file located in the `~/.ssh` directory of the specified User:
      - If the directory and file do not exist, create the `.ssh` directory in the home directory.

        ```
        [oracle@localhost ~]$ cd ~
        [oracle@localhost ~]$ mkdir .ssh
        ```

      - Change directories into the `.ssh` directory, create the `authorized_keys` file and paste the SSH key into a new line of the file.

        ```
        [oracle@localhost ~]$ touch ~/.ssh/authorized_keys
        ```

      - If the `.ssh` directory exists but the file does not, change directories into the `.ssh` directory, create the `authorized_keys` file and paste the SSH key into a new line of the file.
      - If the `authorized_keys` file exists in the `.ssh` directory, edit the `authorized_keys` file and paste the SSH key into a new line of the file.
   e. Set the `.ssh` directory and `authorized_keys` file permissions:

      ```
      chmod 700 .ssh
      ```

      ```
      chmod 640 .ssh/authorized_keys
      ```

6. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.
7. **Script (with Full Path):** Enter the full path to the script, for example `/home/joe/scripts/myscript.sh`.
8. **Parameters:** If the script requires parameters, enter them in this field.

   Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

> **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

The selected Storage Domain will contain the new View where the script results are stored.

9. **NFS View:** Select an existing View create a new one by selecting **Add View**. For more information, see Create Views Using a Predefined Template. The View will contain the script results.
10. **NFS Mount Path:** Click **Copy NFS Path to Clipboard** for use later when you mount the View on the Linux system.
11. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
12. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
13. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
14. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:

- **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
- **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
- **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

For more information, see Alerts.

15. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
16. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

> **Note:**
> The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

17. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
18. **Description:** Optional. Click the plus sign and enter a description for the Protection Group.
19. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

About SAP HANA Database Backups

Cohesity's software-defined platform provides a simple, fast, and cost-effective backup, recovery, and data management solution for growing SAP HANA database environments. It provides an SAP HANA-BRINT 1.1 certified Cohesity BACKINT plugin for backups.

> **Tip:** To protect the SAP HANA database using the Cohesity Connector Agent, see SAP HANA.

When the Cohesity BACKINT plugin is installed and configured on the SAP HANA system, it manages the communication between SAP HANA and the Cohesity cluster.

The Cohesity BACKINT plugin is available for both x86-64 and IBM PowerPC systems. For data transfer between the SAP HANA node and the Cohesity cluster, the Cohesity BACKINT plugin leverages NFS or RPC for IBM PowerPC and RPC for x86-64 systems respectively.

> **Note:** Cohesity supports the recovery of SAP HANA 2.0 backups across both Intel x86-64 and Power PC platforms, allowing you to restore your critical SAP HANA data regardless of the underlying hardware architecture.

You can perform the following operations for data and logs using one of the following SAP HANA interfaces:

Interface	Operations
SAP HANA Studio	Backup Recovery Inquire Deletion
SAP HANA Cockpit	
HDBSQL	
DBA Cockpit	

Core Components

The following are the core components for protecting SAP HANA databases:

- SAP HANA appliance
- Cohesity BACKINT plugin
- Cohesity Data Cloud (Self-Managed)

As illustrated, you can integrate the Cohesity Cohesity Data Cloud (Self-Managed) with a new or existing SAP HANA infrastructure by installing the Cohesity BACKINT plugin on SAP HANA.

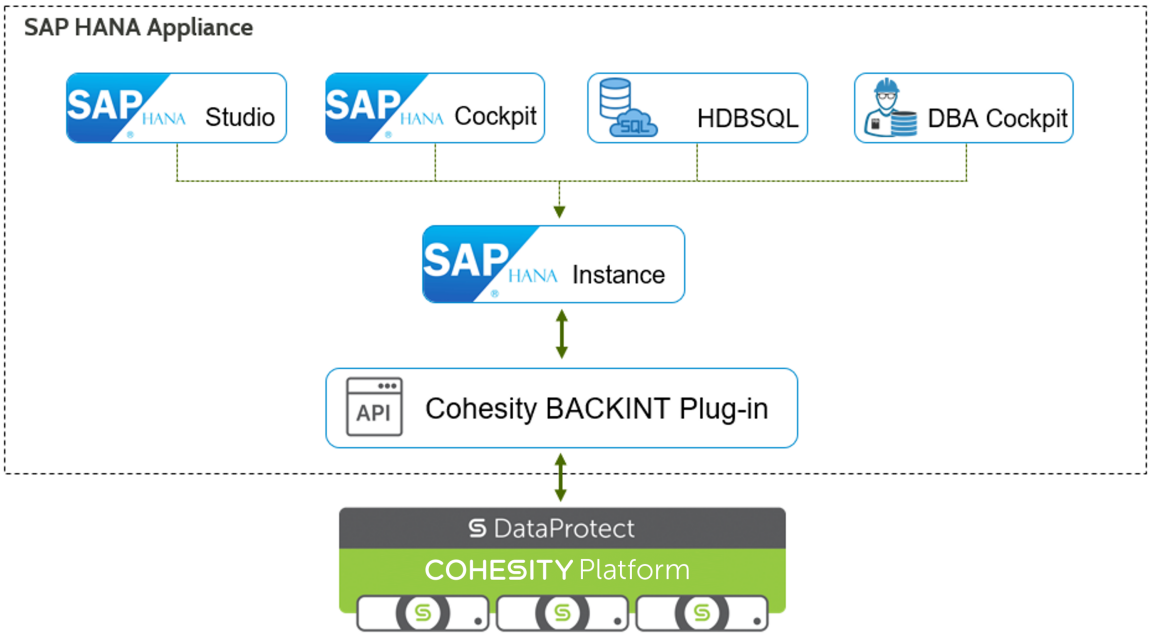

The Cohesity BACKINT plugin on SAP HANA enables you to perform the following functions:

Category	Supported Functions
Operations	Backup Restore Inquire Deletion
Supported Backup Types	Data Logs
Database Type	System Database Tenant Database
Configuration	Single Node SAP HANA Setup Multi-node SAP HANA Setup (Scale-out and Scale-up architecture)
Deployment	SAP HANA instance as a multi-tenant setup

Supported SAP HANA Releases and Environments

For more information on supported SAP on Oracle versions, see Supported Software versions.

Create a Cohesity View with IP Allowlist for SAP HANA node

Create a Cohesity view to provide a storage location in a storage domain on the Cohesity cluster for the SAP HANA database nodes. For more information on views, see Cohesity Views.

To provide effective authorization and prevent rogue access to the Cohesity cluster, the SAP HANA node IP address must be added in the Cohesity view configuration. For detailed instructions to create a view, see Create Views Using a Predefined Template.

Adding Custom Allowlist for the Cohesity View

While creating a Cohesity View, you have an option to add an allowlist for the View. Ensure that you add the SAP HANA node IP address in the Allowlist for the Cohesity View.

To add subnets to the allowlist of the Cohesity view:

1. In the **Create View** page, under the **Advanced** section, expand **Security**.
2. In the **Allowlist** section, perform the following:
    1. Select **Override Global Allowlist** and click **Add Allowlist**.
    2. In the **Subnet IP** and **Subnet Mask** fields respectively, enter the Subnet IP address and Subnet Mask to specify a range of systems that should have permissions to access and mount the Cohesity View.
    3. In the **NFS Permissions**, **SMB Permissions**, and **S3 Permissions** drop-down lists, select the NFS, SMB, and S3 permissions respectively.
    4. In the **NFS Root Squash** and **NFS All Squash** drop-down lists, specify whether to turn on NFS Root or NFS All squash respectively.
    5. In the **Description** field, enter a brief description.
    6. Click **Add**.
    7. Repeat these steps to add additional subnets to the allowlist of the Cohesity view.

> **Note:**
> - To use the IP allowlist feature the Cohesity cluster should be upgraded to 6.3 release or later.
> - If the Override Global Allowlist toggle is off, this View can only be accessed and mounted by systems specified by the Global Cohesity cluster Allowlist.
> - In case of a multinode HANA setup, specify all nodes individually in the global IP allowlist.

Install the Cohesity BACKINT plugin

The Cohesity BACKINT plugin is available for both Intel-based x86 and IBM PowerPC systems. Cohesity supports IPv6 for SAP HANA.

This topic covers instructions on how to:

- Install the Cohesity BACKINT plugin on x86-64 Platform
- Install the Cohesity BACKINT plugin on IBM PowerPC Platform

Install the Cohesity BACKINT plugin on x86-64 Platform

The Cohesity BACKINT plugin is an SAP HANA-BACKINT 1.1 certified service that enables you to establish communication between the Cohesity cluster and the SAP HANA nodes.

The BACKINT plugin is available as a script installer and also as an RPM-based installer for the x86-64 system. Based on your requirements, download the appropriate installer.

Prerequisites

Before you install the BACKINT plugin, make sure that you:

- Create a Cohesity view with the **DefaultStorageDomain** options:
  - In an SAP HANA single-node setup, install the BACKINT plugin in the appropriate SAP HANA directory.
  - For a multi-node setup, install the BACKINT plugin on a shared NFS path that all the nodes can access in the SAP HANA deployment.

  For more information, see Create a Cohesity View with IP Allowlist for SAP HANA node.
- Download and copy the appropriate BACKINT plugin installer for x86-64 systems. For more information, see Download the BACKINT plugin Installer.
- Check if the SAP HANA node has an `<SID>adm` user account, if not, then create an Operating System (OS) user account to own the backup files on the Cohesity view. The format of the user account must be `<SID>adm`.
- "bash" shell is required for BACKINT plugin installation, backup, and restore workflows.
- Run the following command on the SAP HANA node to synchronize time between the SAP HANA server and the Cohesity NTP server:

  ```
  timedatectl set-ntp true
  ```

- If you are using the RPM-based installer, ensure that you configure the following environment variables on the SAP HANA node(s):

Environment Variable	Description	Mandatory?	Command
SID	The <SID>adm user account that must own the backup files on the Cohesity view. If you have not created the user account, then create an Operating System (OS) user account in the following format: <SID>adm. Where <SID> is a unique three-character code called SAP System Identification (SID). It is available for every R/3 installation (SAP system) that consists of a database server and several application servers.	Yes	export SID=<VALUE>
SYSTEMDB_KEY	The SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB.	No	export SYSTEMDB_KEY=<HDBUSERSTORE key>
CLUSTER_PORT	The ports in the firewall to allow the Cohesity cluster to transmit and receive data. The supported values are 11111 or 11113. By default the port is set as 11117.	No	export CLUSTER_PORT=<Port_number>
WORKFLOW	The type of workflow you intend to install the BACKINT plugin. The supported values are uda or ra. The uda value is for Universal Data Adapter workflow and ra is for Remote Adapter workflow. Default is uda.	No	export WORKFLOW=ra
MULTI_TENANCY	The type of deployment. The supported values are vlan or hyx. Specify vlan for 'VLAN' and hyx for 'Single Homed NAT' deployments.	No	export MULTI_TENANCY=<VALUE>
STANDBY	Specify if you are installing the BACKINT plugin on a standby node. Supported values are 1 or 0.	No	export STANDBY=<VALUE>

Environment Variable	Description	Mandatory?	Command
HANA_ VERSION	The SAP HANA version. Supported values are 1 or 2.	No	`export HANA_ VERSION=<VALUE>`
DEBUG	Option to enable installer debug. Supported values are 1 or 0.	No	`export DEBUG =<VALUE>`

Download the BACKINT plugin Installer

Cohesity provides a BACKINT plugin installer to backup and restore SAP HANA databases. The BACKINT plugin is available as script-based and RPM-based installers on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective BACKINT plugin installer, and copy the installer to the SAP HANA host.

Install the BACKINT plugin using the Script Installer

The Cohesity BACKINT plugin is an SAP HANA-BACKINT 1.1 certified service that enables you to establish communication between the Cohesity cluster and the SAP HANA nodes.

> **Note:** For multi-node SAP HANA deployments, ensure that you install the Cohesity BACKINT plugin on the shared NFS path where all the SAP HANA nodes have access or install the plugin on all the nodes of the SAP HANA deployment as an SAP HANA admin user.

On the SAP HANA node, perform the following to install the plugin:

1. Log on to the SAP HANA node's bash CLI.
2. Make the plugin installer (`cohesity_secure_connector_service_<version>_sap_hana_installer`) executable using the following command:

   ```
   SAP_HANA_SERVER:/home/adminse # chmod +x cohesity_secure_connector_service_<version>_sap_hana_installer
   ```

3. Ensure that the `opt` folder in the `/usr/sap/<SID>/SYS/global/hdb/` directory, if not create the `opt` folder as follows:

   ```
   SAP_HANA_SERVER:/usr/sap/COH/SYS/global/hdb # mkdir opt
   ```

> **Note:** During the BACKINT plugin installation, by default, the script installer extracts the files to the `/tmp` directory. If you plan to extract the files to a different directory, you can define the directory path in the `TMPDIR` environment variable as follows: `export TMPDIR=<directory_path>`.

4. Install the plugin using the following command:

```
./cohesity_secure_connector_service_<version>_sap_hana_installer -- --install --install-dir=<backint install directory> --sid=<SID> --systemdb-key=<systemdb_key> --workflow=<workflow>
```

Where:

Options	Mandatory?	Description
`--install-dir`	No	The directory on the SAP HANA node on which you want to install the BACKINT plugin. **Note:** If `--install-dir` is not specified, the BACKINT plugin is installed in the `/usr/sap/${SAPSYSTEMNAME}/home` directory.
`--systemdb-key`	NO	The SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB. • Use the option `--systemdb-key <user_store_key>` only if you want to edit the global.ini file and create a symbolic link from `HDBBACKINT` to the BACKINT plugin. • If you plan to back up your SAP HANA databases using the BACKINT plugin, ensure that you use the option `--systemdb-key <user_store_key>`. • If you have not used the option `--systemdb-key <user_store_key>` while installing the BACKINT plugin, specify the option `--systemdb-key <user_store_key>` while upgrading the BACKINT plugin.

Options	Mandatory?	Description
`--cluster-port`	No	The ports in the firewall to allow the Cohesity cluster to transmit and receive data. The supported values are 11111 or 11113. By default the port is set as `11117`.
`--workflow`	No	The type of workflow you intend to install the BACKINT plugin. The supported values are `uda` or `ra`. The `uda` value is for Universal Data Adapter workflow and `ra` is for Remote Adapter workflow.
`--multi-tenancy`	No	The type of deployment. The supported values are `vlan` or `hyx`. Specify `vlan` for 'VLAN' and `hyx` for 'Single Homed NAT' deployments.
`--standby`	No	Specify if you are installing the BACKINT plugin on a standby node.
`--hana-version`	No	The SAP HANA version. Supported values are 1 or 2. By default, the value is set to 1.
`--debug`	No	Option to enable installer debug.

Install the BACKINT plugin using the RPM-based Installer

On the SAP HANA node, perform the following to install the BACKINT plugin using an RPM-based installer, `cohesity-hana-backint-<version>.1-1.x86_64.rpm`. You can install the BACKINT plugin using a `root` user account or a user account with `sudo` privileges.

To install the BACKINT plugin using a user account with sudo privileges:

1. Log in to the SAP HANA node's bash CLI.
2. Navigate to the directory where you have copied the RPM-based installer.
3. Run the following command to install the BACKINT plugin:

```
sudo -E rpm -ivh  cohesity-hana-backint-<version>.x86_64.rpm
```

To install the BACKINT plugin using the `root` user account:

1. Log on to the SAP HANA node's bash CLI as a root user.
2. Navigate to the directory where you have copied the RPM-based installer.
3. Run the following command to install the BACKINT plugin:

```
rpm -ivh cohesity-hana-backint-<version>.x86_64.rpm
```

Verify the BACKINT plugin Installation

To verify if the RPM-based BACKINT plugin installation is successful, run the following command:

```
rpm -qa|grep  cohesity-hana-backint
```

The following is a sample output of the command:

```
sap-hana-1.eng:qqqadm> rpm -qa|grep  cohesity-hana-backint
cohesity-hana-backint-<version>.1-1.x86_64
```

Update the Cohesity Parameter File

As part of the Cohesity BACKINT plugin installation, Cohesity provides a parameter file (`cohesity_param_file`) that enables SAP HANA to use the Cohesity BACKINT plugin for a backup job.

The parameter file is available in the <`backint_install_directory`>. Edit the parameter file and enter the values for the following parameters:

Parameter	Description
`cluster_ips`	List of IP addresses of the Cohesity Cluster. **Note:** Ensure that you provide only the virtual IP addresses.
`cluster_port`	The port number of the Cohesity cluster nodes that the plugin uses to establish communication between SAP HANA and the Cohesity cluster. The default port number is 11113 (gRPC). **Note:** You can also use the 11111 (SunRPC) port to establish communication between SAP HANA and the Cohesity cluster. Ensure that you open the 11111 (SunRPC) port on the Cohesity cluster.
`view_name`	The name of the Cohesity view used to store the backup data. **Note:** Ensure that the Cohesity view name does not contain the path of the Cohesity view or Cluster Name or IP address.
`io_thread_count`	The number of IO threads for reading or writing data on the Cohesity view.

`max_outstanding_io`	The number of pending IO threads. The range is 1-128 and the recommended value is 32.
`max_io_size_bytes`	Maximum size of each IO in bytes. The range is 1048576-4194304 bytes and the recommended value is 4194304 bytes.
`max_cached_io`	The maximum number of IO cached/prefetched. The range is 1-128 and the recommended value is 64.
`rpc_timeout_msecs`	Timeout in milliseconds for RPC requests.
`inquire_period_days`	The number of days for which the inquire information is returned. The default value is 180 days.
`log_dir`	Location on the SAP nodes where the Cohesity BACKINT plugin creates the debug logs. **Note:** The Backint logs include details about secure communication during backup and restore operations. When secure communication is enabled, the logs record the negotiated TLS version and cipher suite between the SAP HANA host and the Cohesity cluster in the following format: `TLSv1.2:AES256-GCM-SHA384`
`metafiles_read_concurrency`	Tune the metadata read performance for inquiry and restoration. The default value is 100.
`max_backup_streams`	The number of pipes that are backed up in parallel at any time by Cohesity BACKINT plugin during the backup process. It is recommended that you provide the number of IP addresses of the Cohesity Cluster that you define in the `cluster_ips` parameter.
`enable_dedup_write`	To enable source-side deduplication during backup, set this parameter value to true. **Note:** By default, source-side deduplication is disabled during the backup process.
`use_tcp_proxy`	To use anycast IP address and stream data through multiple TCP connections. Set this parameter value to `true`, if you want to use anycast IP load balancing in your deployment.

`proxy_port_range`	Use the `proxy_port_range` parameter only if you have set the `use_tcp_proxy` parameter to true. Enabling the `use_tcp_proxy` parameter opens certain ports on the localhost interface on the SAP HANA source. However, if you have a firewall enabled that blocks these ports, then you need to allowlist a range of ports and specify that range of ports using the `proxy_port_range` parameter. For example, `'proxy_port_range : "32001-35000"'`

Schedule and Run an SAP HANA Backup Job

Using external applications like SAP HANA Cockpit, SAP HANA Studio or using native SQL, schedule and run the backup job.

Uninstall the Cohesity BACKINT plugin

To uninstall the Cohesity BACKINT plugin, run the following command on the SAP HANA node's bash CLI:

```
cohesity_secure_connector_service_<version>_sap_hana_installer -- -u -d
<backint_install_directory> -s <SID>
```

If you have installed BACKINT plugin using the RPM-based installer, you can uninstall the service using the following steps.

To uninstall the BACKINT plugin using the `<SID>adm` user account, use the following command:

```
sudo -E rpm -evh cohesity-hana-backint-<version>.1-1.x86_64
```

To uninstall the BACKINT plugin using the `root` user account, use the following command:

```
rpm -evh cohesity-hana-backint-<version>.1-1.x86_64
```

The following is a sample output of the command:

```
[root@sap-hana-1 home]# rpm -evh cohesity-hana-backint-<version>.1-1.x86_64
```

Upgrade the Cohesity BACKINT plugin

To upgrade the BACKINT plugin, run the following command on all SAP HANA nodes:

```
cohesity_secure_connector_service_<version>_sap_hana_installer -- -n -d
<backint_install_directory> -s <SID> -k <systemdb_key>
```

Install the Cohesity BACKINT plugin on IBM PowerPC Platform

This topic covers instructions on how to install the BACKINT plugin for SAP HANA on IBM PowerPC Platform.

Prerequisites

Before you install the BACKINT plugin, make sure that you:

- Create a Cohesity view with the **DefaultStorageDomain** options:
  - In an SAP HANA single-node setup, install the BACKINT plugin in the appropriate SAP HANA directory.
  - For a multi-node setup, install the BACKINT plugin on each node of the SAP HANA deployment.
    For more information, see Create a Cohesity View with IP Allowlist for SAP HANA node.
- SAP HANA node(s) must be using JDK version 1.8 or higher.
- Install NFS common utility packages on the SAP HANA node(s) if you want to use NFS for data transfer between the SAP HANA node and the Cohesity cluster.
- Download and copy the appropriate BACKINT plugin installer for IBM PowerPC systems to the node or nodes of your SAP HANA deployment. For more information, see Download the BACKINT Plugin Installer.
- Check if the SAP HANA node has a `<SAP_HANA_system_ID>adm` user account, if not, then create an Operating System (OS) user account to own the backup files on the Cohesity view. The format of the user account must be `<SAP_HANA_system_ID>adm`.
- Use "bash" shell to install the BACKINT plugin and to run backups and restore workflows of SAP HANA databases.
- Create SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB. You can create the key using the `hdbuserstore` command. For detailed instructions, see hdbuserstore Commands in SAP HANA documentation.
- By default, SAP HANA uses BACKINT protocol version 1.5 on the IBM PowerPC Platform, starting from SAP HANA 2.0 SPS 05. However, if you have an earlier version of SAP HANA databases, you can enable it to use BACKINT protocol version 1.5 by configuring the `global.ini` file with the following command:

```
hdbsql -AU <USER_STORE_KEY> "alter system alter configuration
('global.ini', 'SYSTEM') set ('backup', 'backint_protocol_version') =
'1.5' with reconfigure"
```

> **Note:** Ensure that the Cohesity BACKINT plugin and SAP HANA use the same BACKINT protocol version in your deployment. You can specify the BACKINT protocol version that the Cohesity BACKINT plugin uses using the Cohesity Parameter File.

Download the BACKINT Plugin Installer

Cohesity provides a BACKINT plugin installer to backup and restore SAP HANA databases. The BACKINT plugin is available as script-based and RPM-based installers on the Cohesity Download portal. Log on to the Cohesity Download portal, download the respective BACKINT plugin installer, and copy the installer to the SAP HANA host.

Install the BACKINT Plugin Using Script Based Installer

Cohesity BACKINT plugin is an SAP HANA-BRINT 1.1 certified service that enables you to establish communication between the Cohesity cluster and the SAP HANA nodes.

The BACKINT plugin can be installed only using `<sid>adm` user.

> **Note:** For multi-node SAP HANA deployments, ensure that you install the Cohesity BACKINT plugin on the shared NFS path where all the SAP HANA nodes have access or install the plugin on all the nodes of the SAP HANA deployment as an SAP HANA admin user.

On the SAP HANA node, perform the following to install the BACKINT plugin:

1. Log on to the SAP HANA node's bash CLI.
2. Make the BACKINT plugin installer (`cohesity_secure_connector_service_hana_<version>_sap_hana_java_installer`) executable using the following command:

   ```
   SAP_HANA_SERVER:/home/adminse # chmod +x cohesity_secure_connector_service_hana_<version>_sap_hana_java_installer
   ```

3. Ensure that the `/opt` folder in the `/usr/sap/<SID>/SYS/global/hdb/` directory, if not create the `/opt` folder as follows:

   ```
   SAP_HANA_SERVER:/usr/sap/COH/SYS/global/hdb # mkdir opt
   ```

   > **Note:** During the BACKINT plugin installation, by default, the script installer extracts the files to the `/tmp` directory. If you plan to extract the files to a different directory, you can define the directory path in the `TMPDIR` environment variable as follows: `export TMPDIR=<directory_path>`.

4. Run the following command to install the BACKINT plugin:

   ```
   ./cohesity_secure_connector_service_<version>_sap_hana_java_installer -- --install=<backint install directory> --sid=<SID> --systemdb-key=<systemdb_key> --workflow=<workflow>
   ```

   > **Note:** You can view the details of BACKINT plugin installer options using the `--help` option. For example: `./cohesity_secure_connector_service_sap_hana_installer --help`After the BACKINT plugin is installed, the BACKINT plugin agent is available at: `<backint_install_directory>/cohesity_backint_plugin/uda_scripts`

Where:

Options	Mandatory?	Description
`--install-dir`	No	The directory on the SAP HANA node on which you want to install the BACKINT plugin. **Note:** If `--install-dir` is not specified, the BACKINT plugin is installed in the `/usr/sap/${SAPSYSTEMNAME}/home` directory.
`--systemdb-key`	NO	The SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB. • Use the option `--systemdb-key-k <user_store_key>` only if you want to edit the global.ini file and create a symbolic link from `HDBBACKINT` to the BACKINT plugin. • If you plan to backup your SAP HANA databases using the BACKINT plugin, ensure that you use the option `--systemdb-key <user_store_key>`. • If you have not used the option `--systemdb-key <user_store_key>` while installing the BACKINT plugin, specify the option `--systemdb-key <user_store_key>` while upgrading the BACKINT plugin.
`--cluster-port`	No	The ports in the firewall to allow the Cohesity cluster to transmit and receive data. The supported values are 11111 or 11113. By default the port is set as `11117`.
`--workflow`	No	The type of workflow you intend to install the BACKINT plugin. The supported values are `uda` or `ra`. The `uda` value is for Universal Data Adapter workflow and `ra` is for Remote Adapter workflow.
`--multi-tenancy`	No	The type of deployment. The supported values are `vlan` or `hyx`. Specify `vlan` for 'VLAN' and `hyx` for 'Single Homed NAT' deployments.

Options	Mandatory?	Description
`--standby`	No	Specify if you are installing the BACKINT plugin on a standby node.
`--hana-version`	No	The SAP HANA version. Supported values are 1 or 2. By default, the value is set to 1.
`--debug`	No	Option to enable installer debug.
`--catalog-view`	No	The name of the catalog view where the backup metadata resides. **Note:** The catalog view name should not include the view path or cluster name or IP address. Applicable only for Powerpc

Install the BACKINT Plugin Using RPM Based Installer

Before you install the BACKINT plugin Using RPM-based installer, ensure that you configure the following environment variables on the SAP HANA machine:

Environment Variable	Description	Mandatory?	Command
`SID`	The <SID>adm user account that must own the backup files on the Cohesity view. If you have not created the user account, then create an Operating System (OS) user account in the following format: <SID>adm. Where <SID> is a unique three-character code called SAP System Identification (SID). It is available for every R/3 installation (SAP system) that consists of a database server and several application servers.	Yes	`export SID=<VALUE>`
`SYSTEMDB_KEY`	The SAP HANA HDBUSERSTORE key to connect to the SYSTEM DB.	No	`export SYSTEMDB_KEY=<HDBUSERSTORE key>`

Environment Variable	Description	Mandatory?	Command
CLUSTER_PORT	The ports in the firewall to allow the Cohesity cluster to transmit and receive data. The supported values are 11111 or 11113. By default the port is set as 11117.	No	export CLUSTER_PORT=<Port_number>
WORKFLOW	The type of workflow you intend to install the BACKINT plugin. The supported values are uda or ra. The uda value is for Universal Data Adapter workflow and ra is for Remote Adapter workflow. Default is uda.	No	export WORKFLOW=ra
MULTI_TENANCY	The type of deployment. The supported values are vlan or hyx. Specify vlan for 'VLAN' and hyx for 'Single Homed NAT' deployments.	No	export MULTI_TENANCY=<VALUE>
STANDBY	Specify if you are installing the BACKINT plugin on a standby node. Supported values are 1 or 0.	No	export STANDBY=<VALUE>
HANA_VERSION	The SAP HANA version. Supported values are 1 or 2.	No	export HANA_VERSION=<VALUE>
DEBUG	Option to enable installer debug. Supported values are 1 or 0.	No	export DEBUG=<VALUE>
CATALOG_VIEW	The name of the catalog view where the backup metadata resides. **Note:** The catalog view name should not include the view path or cluster name or IP address.	No	export CATALOG_VIEW=<VALUE>

To install the BACKINT plugin using rpm-based installer:

1. Log on to the SAP HANA node's bash CLI using <SID>adm user.
2. Make the BACKINT plugin installer (`cohesity-hana-java-backint-<version>.noarch.rpm`) executable using the following command:

```
> chmod +x cohesity-hana-java-backint-<version>.noarch.rpm
```

3. Install the BACKINT plugin using the following command:

```
> sudo -E rpm -ivh cohesity-hana-java-backint-<version>.noarch.rpm
```

Update the Cohesity Parameter File

As part of the Cohesity BACKINT plugin installation, Cohesity provides a parameter file (`cohesity_param_file`) that enables SAP HANA to use BACKINT plugin for a backup job.

The parameter file is available in the <`backint_install_directory`>. Edit the parameter file and enter the values for the following parameters:

Parameter	Description
`cluster_ips`	List of IP addresses of the Cohesity Cluster. **Note:** Ensure that you provide only the virtual IP addresses.
`cluster_port`	The port number of the Cohesity cluster nodes that the plugin uses to establish communication between SAP HANA and the Cohesity cluster. The default port number is 11113 (gRPC). **Note:** You can also use the 11111 (SunRPC) port to establish communication between SAP HANA and the Cohesity cluster. Ensure that you open the 11111 (SunRPC) port on the Cohesity cluster.
`view_name`	The name of the Cohesity view used to store the backup data. **Note:** Ensure that the Cohesity view name does not contain the path of the Cohesity view or Cluster Name or IP address.
`sap_nfs_mount_path`	The NFS mount path.
`sap_nfs_mount_options`	The NFS mount options.
`sap_max_read_write_retries`	The Maximum number of retries during reads or writes.
`sap_max_io_size_in_bytes`	Maximum size of each IO in bytes. The range is 1048576-4194304 bytes and the recommended value is 4194304 bytes.

`sap_max_data_queue_size`	Maximum number of IO cached or prefetched. The range is 1-128 and the recommended value is 64.
`sap_max_num_inflight_io`	Maximum number of inflight reads or writes at a given time. The range is 1-128.
`sap_max_io_asynchronous_callback_concurrency`	Maximum number of asynchronous callback threads for NFS read-write.
`sap_max_backup_streams`	The number of pipes that are backed up in parallel at any given time by BACKINT plugin during the backup process. It is recommended that you provide the number of IP addresses of the Cohesity Cluster that you define in the `cluster_ips` parameter.
`sap_max_restore_streams`	The number of pipes that are restored in parallel at any given time by BACKINT plugin during the restore process. It is recommended that you provide the number of IP addresses of the Cohesity Cluster that you define in the `cluster_ips` parameter.
`sap_max_metafiles_read_concurrency`	Tune the metadata read performance for inquiry and restore. The default value is 4.
`sap_max_inquire_period_days`	The number of days for which the inquire information is returned. The default value is 180 days.
`backint_log_dir`	Location on the SAP nodes where the BACKINT plugin creates the debug logs. **Note:** The BACKINT logs (configured using the backint_log_dir parameter) include details about secure communication between the SAP HANA host and the Cohesity cluster. In addition to the TLS version, the logs also record the negotiated cipher suite in the following format: `tls_version=TLSv1.2 cipher=TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384` Use the BACKINT logs to verify that TLS is enabled and that a valid cipher suite is negotiated. If this information is not present, ensure that secure communication is enabled between SAP HANA host and the Cohesity cluster. In some cases, this information may also appear in pulse logs.

`java_log_level`	The java log level. The supported values are: • all • trace • debug • info • warn • error • fatal
`service_type`	The type of service. The supported values are NFS and RPC.
`backint_ protocol_ version`	The BACKINT protocol version that you want the BACKINT plugin to use. The BACKINT protocol version 1.5 is supported by default from SAP HANA 2.0 SPS 05 or later versions on the IBM PowerPC Platform. By default, the BACKINT plugin supports BACKINT protocol version 1.0. The supported values are 1.0 and 1.5.

Schedule and Run an SAP HANA Backup Job

Using external applications like, SAP HANA Cockpit, SAP HANA Studio or using native SQL, schedule and run the backup job.

Uninstall the BACKINT Plugin

To uninstall the BACKINT plugin using the root user account, use the following command:

Using the Script installer:

```
./cohesity_secure_connector_service_<version>_sap_hana_java_installer -- --
uninstall --install-dir=<backint_install_directory> --workflow=<workflow>
```

Using the RPM installer:

```
> sudo -E rpm -evh cohesity-hana-java-backint-<version>.noarch.rpm
```

Upgrade the BACKINT Plugin

To upgrade the BACKINT plugin, run the following command on all SAP HANA nodes.

Using the Script installer:

```
./cohesity_secure_connector_service_<version>_sap_hana_java_installer -- --
upgrade --install-dir=<backint_install_directory> --workflow=<workflow>
```

Using the RPM installer:

```
sudo -E rpm -Uvh cohesity-hana-java-backint-<version>.noarch.rpm
```

Create a Remote Adapter Job for SAP HANA Database

You can create backup jobs and purge jobs for SAP HANA databases using the Remote Adapters on Cohesity cluster.

To Create a Protection Group for SAP HANA Database

1. Run the following command on the SAP HANA node's bash CLI as an SAP HANA admin user.

   For example, if the SID is `DEV` and the secure connector service installation location is `/usr/sap/DEV/home/install`.

   ```
   /usr/sap/DEV/home/install/backup_job.sh
   ```

2. Navigate to **Data Protection > Protection** and do one of the following:
   - To create a Protection Group—Click **Protect** on the top-right of the page and then select the type of Protection Group.
   - To edit a Protection Group—Click the action icon (⋮) next to the Protection Group and select **Edit**.

   Here the *Job Type* is **Remote Adapter**.
3. **Name:** Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
4. **Linux Hostname or IP:** Specify the Hostname or the IP Address of the Linux system where the scripts are located and will be run.
5. **Username:** Specify the username on the Linux system that will be running the scripts.
6. **Cluster SSH Public Key:** Set up permissions on the Linux system to allow the Cohesity cluster to run the scripts remotely on the system.
   a. Click **Copy Key to Clipboard** to copy the SSH public key into the clipboard.
   b. Log in to the Linux system with the username specified earlier.
   c. In the home directory of the user, determine if there is a `.ssh` directory that contains a file called `authorized_keys`:

      ```
      ls -lr ~/.ssh
      ```

   d. Add the SSH key to the `authorized_keys` file located in the `~/.ssh` directory of the specified User:
      - If the directory and file do not exist, create the `.ssh` directory in the home directory, change directories into the `.ssh` directory, create the `authorized_keys` file and paste the SSH key into a new line of the file.
      - If the `.ssh` directory exists but the file does not, change directories into the `.ssh` directory, create the `authorized_keys` file and paste the SSH key into a new line of the file.

- If the `authorized_keys` file exists in the `.ssh` directory, edit the `authorized_keys` file and paste the SSH key into a new line of the file.

e. Set the `.ssh` directory and `authorized_keys` file permissions:

```
chmod 700 .ssh
chmod 640 .ssh/authorized_keys
```

For example:

```
orcadm@rhel-sap-hana-002:/usr/sap/ORC/home> ls -lda .ssh
drwx------ 2 orcadm sapsys 48 Nov 27  2020 .ssh
orcadm@rhel-sap-hana-002:/usr/sap/ORC/home> cd .ssh/
orcadm@rhel-sap-hana-002:/usr/sap/ORC/home/.ssh> ls -ltr
total 8
-rw-r----- 1 orcadm sapsys 837 Nov 14  2020 authorized_keys
-rw-r--r-- 1 orcadm sapsys 868 May 30 23:46 known_hosts
orcadm@rhel-sap-hana-002:/usr/sap/ORC/home/.ssh>
```

7. **Policy:** Select an existing protection policy or create a new one by selecting **New Policy**.
8. **Script (with Full Path):** Enter the full path to the script, for example: `/usr/sap/DEV/home/install/backup_job.sh`.
9. **Parameters:** If the script requires parameters, enter them in this field.

   For example: `--backup-delta=INCREMENTAL --user-store-key=SYSTEMDB --tenant-database-names=BHE,BWE --mode=1`

> **Note:**
> You can define the values for the **backup_job.sh** script parameters either using the long or short options. Ensure that:
>
> - If you are using the long options, you need to define the value for the parameter using the "=" character. For example, `--tenant-database-names=BHE`.
> - If you are using the short options, you need to define the value for the parameter using a single whitespace. For example, `-t BHE`.

Parameter	Mandatory?	Description
`-t, --tenant-database-names`	Yes	Specify a comma-separated list of tenant database names that need to be backed up. The default value is `SYSTEM DB`. **Note:** Do not include space between the tenant database names. For example: `--tenant-database-names=ORC,PRD,QTY` or `-t ORC,PRD,QTY`.

Parameter	Mandatory?	Description
`-b, --backup-delta`	No	Specify the backup type. Allowed values are `INCREMENTAL` or `DIFFERENTIAL`. The default value is `COMPLETE`. **Note:** If you have not specified the backup type in the parameter, then a full backup (`COMPLETE`) is performed.
`-k, --user-store-key`	Yes	The user store key to connect to the `SYSTEM DB`.

Parameter	Mandatory?	Description
`-m, --mode [0|1]`	No	The backup order of the databases: • If `--mode=1`, then the backup is performed in the order of `SYSTEM DB`, all tenant databases, and `SYSTEM DB` as per SAP standard backup. • If `--mode=0`, then the backup is performed only for the provided tenant databases. Allowed values are 0 or 1. The default value is 0.

Parameter	Mandatory?	Description
`-l, --backup-log-file [file]`	No	Define the path and filename of the log file to store the backup script logs. By default, the log filename is set as: `/tmp/cohesity_hana_backup_logs/cohesity_hana_backup_<time_in_seconds>_<pid>.log`
`-d, --debug`	No	Allows you to enable shell script debugging. By default, shell script debugging is disabled.
`-e, --enable-safepurge`	No	Enable to delete the failed backups or the backups not present in the catalog from the View. By default, the option is disabled

Parameter	Mandatory?	Description
`-r, --retention-days`	No	You can use the same Protection Group to run both the backup Remote Adapter job (`backup_job.sh`) and the purge Remote Adapter job (`purge_job.sh`). To run the purge job, you need to specify the number of days for the data to be purged from the catalog. This means that any data that is older than the specified number of days will be removed from the catalog. **Note:** It is recommended to use separate Protection Groups to run the backup Remote Adapter job (`backup_job.sh`) and the purge Remote Adapter job (`purge_job.sh`)

Storage Domain: Select an existing storage domain or create a new one by selecting **New Storage Domain**.

> **Note:** Once you create a Protection Group, you cannot change the Storage Domain you selected for that job.

The selected Storage Domain will contain the new View where the script results are stored.

10. **NFS View:** Select an existing View create a new one by selecting **Add View**. For more information, see Create Views Using a Predefined Template. The View will contain the script results.
11. **NFS Mount Path:** Click **Copy NFS Path to Clipboard** for use later when you mount the View on the Linux system.
12. **Start Time:** Indicates when the Protection Group should run. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting.

    The default time zone is the browser's time zone. You can change the time zone by selecting a different time zone.
13. **End Date:** Optional. Toggle on and select the date when the Protection Group stops capturing snapshots. A protection run that starts prior to this date runs until completion even if it completes after this date.
14. **QoS Policy:** Select an appropriate quality of service (QoS) policy.

    > **Note:** Cohesity recommends specifying **Backup HDD**, which is the default.

    - **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group.
    - **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups.
    - **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy.
    - **TestAndDev High:** The Cohesity cluster writes the data to an SSD drive for this Protection Group. The I/Os with this QoS policy are given higher priority compared to other QoS policies. Use this policy for workloads that require lower I/O latency.
15. **Alerts:** Optional. Select one or more of the following settings if you want alerts to be created for the following triggers:

- **Success:** Create an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created.
- **Failure:** Create a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created.
- **SLA Violation:** Create a warning alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created.

For more information, see Alerts.

16. **Priority:** Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Protection Groups exceeds the ability to process them, those with **High** priority are given the highest priority to execute, **Medium** the second-highest priority, and **Low** the lowest priority.
17. **Email Recipients**: You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the Protection Group.

> **Note:**
> The connection to the SMTP Server must be enabled to send email notification. You can enable the SMTP Server connection in the configuration settings page (Cohesity UI > **Settings** > **Summary** > **Configure** > **Enable SMTP Server**).

18. **SLA:** The service-level agreement (SLA) defines the expected time to complete either a full backup or incremental backup. The backup completion time depends on the Windows volume size on AD, objects changed between incremental backups, network speed, and other Protection Groups running on the Cohesity cluster.
    - **Incremental:** Enter the number of minutes you expect an incremental protection run to complete. An incremental backup captures only the differences (changed blocks) since the last protection run.
    - **Full:** Enter the number of minutes you expect a full backup protection run to complete. A full backup captures the entire object (all blocks).
19. Click **Protect**.

    The Protection page is displayed and includes the new Protection Group.

To Create a Purge Job for SAP HANA Database

The purge Remote Adapter job removes the expired backups (decided by the retention period) from the Cohesity view. The full backup is independent, and the incremental, log, and differential backups depend on the full backup. Hence, we retain the last successful full backup on or before the retention period.

For example, suppose the purge Remote Adapter job is triggered on November 26th, and the retention period is configured as 14 days. In that case, the last successful full backup on

or before November 12th is retained, and any backups older than the last successful full backups are deleted.

> **Note:** DataLock is not supported when using the Remote Adapter workflow.

1. Run the following command on the SAP HANA node's bash CLI as an SAP HANA admin user.

   ID is `DEV` and secure connector service installation location is `/usr/sap/DEV/home/install`.

   ```
   /usr/sap/DEV/home/install/purge_job.sh
   ```

2. The steps for creating a purge job is similar to creating a backup job for the SAP HANA database. Follow the same steps listed in the Create a Remote Adapter Job for SAP HANA Database section.

However, you must note the following:

- For **Script (with Full Path):** Enter the full path to the script.

  ```
  /usr/sap/DEV/home/install/purge_job.sh
  ```

- For **Parameters:**, enter the following parameters.

  > **Note:**
  > You can define the values for the **purge_job.sh** script parameters either using the long or short options. Ensure that:
  >
  > - If you are using the long options, you need to define the value for the parameter using the `"="` character. For example, `--tenant-database-names=BHE`.
  > - If you are using the short options, you need to define the value for the parameter using a single whitespace. For example, `-t BHE`.

Parameter	Mandatory?	Description
`-t, --tenant-database-names`	Yes	The names of the tenant database for which backups are purged. **Note:** Ensure that you add all the tenant DBs for the SID in a comma separated list, for example, `./purge_job.sh -r 7 -t AOV,H0V,S0V,COCKPITDB -k COHESITY -e`

Parameter	Mandatory?	Description
`-k, --user-store-key`	Yes	The user store key to connect to the `SYSTEM DB`.
`-l, --purge-log-file [file]`	No	Define the log file's path and filename to store the purge script logs. By default, the log filename is set as: `/tmp/cohesity_hana_purge_logs/cohesity_hana_purge_<time_in_seconds>_<pid>.log`
`-d, --debug`	No	Allows you to enable shell script debugging. The default value is `off`.
`-y, --yes-all`	No	Enables you to answer '`yes`' to all the questions in the purge script. The default value is `off`.
`-r, --retention-days`	No	Define to purge any data from the catalog older than the number (#) of days passed.
`-e, --enable-safe-purge`	No	Enable to delete the failed backups or the backups not present in the catalog from the Cohesity View. By default, the option is disabled.

SAP HANA Certificate Authentication

The SAP HANA Certification provides the necessary authentication and also verifies whether the SAP HANA database, which performs backups to the Cohesity cluster, is a trusted entity. This certification provides the capability to set verification which allows only trusted sources to write to the Cohesity cluster through the SpanFS.

Create an SAP HANA Authentication Certificate

This section lists the steps for creating an SAP HANA authentication certificate.

1. Start the Cohesity CLI locally and log in to the Cohesity cluster using the ssh command with the `support` user account and specify the IP address. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

   ```
   ssh support@198.51.100.12
   ```

2. If prompted, enter the password.

   ```
   [support@198.51.100.12's password:
   ```

   The cohesity_shell prompt is displayed.

3. Log in to the CLI:

   ```
   cohesity shell# iris_cli
   ```

The admin prompt is displayed.

4. Optionally, run the help command for a list of CLI commands.

```
admin@198.51.100.12> help
```

5. Run the cluster gen-src-cert command to generate a new SAP HANA certificate.

```
admin@198.51.100.12 cluster gen-src-cert cert-name="newdemocert" cert-validity=100 host-name="10.2.146.175" host-target="/tmp"
Success: Certificate "newdemocert" generation and deployment successful.
```

Parameter Name	Description
cert-name	Name of the certificate file. Its naming format is ${cert-name}_<random_number>.cfg.
cert-validity	The validity of the source certificate generated. The time is displayed in days.
host-name	The IP address/host name of the target system where you want to move the certificate. This is a mandatory parameter because the certificate generation process needs details of the host machine to generate public and private keys.
host-target	The file directory path on the Cohesity cluster where the certificate is generated.

> **Note:** After the SAP HANA certificate is generated on the Cohesity cluster, manually copy the certificate to your SAP HANA machine.

6. Update the following parameters in the parameter file.

Parameter Name	Description
cluster_port	Port number on Cohesity cluster nodes to which the BACKINT executable connects. Use the port 11117 for secure connection (certificate config path should be a valid path). Use the port 11111 for an insecure connection (certificate config path should be an empty string). Contact Cohesity Support to enable these ports.

Parameter Name	Description
certificate_ config_path	The path to which the certificate is generated and deployed from the Cohesity Dashboard to the SAP HANA Node. You can use this path as the default config path: `/usr/sap/ORC/home/install/newdemocert.cfg`

7. Run the `list-source-certificates` command to check all the global certificates deployed from the cluster.

Deploy an SAP HANA Certificate in a Multi-node Setup

This section lists the steps to deploy an SAP HANA authentication certificate in a multi-node SAP HANA setup.

1. Generate and deploy the SAP HANA authentication certificate to any one of the nodes in a multi-node SAP HANA system using the instructions provided in the previous section.
2. Use SCP to securely transfer the certificate to a target location on all the remaining (N-1) nodes.

   ```
   scp <certificate_config_path> ip_address:<certificate_config_path>
   ```

   Example

   ```
   scp /usr/sap/ORC/home/install/cert.cfg 10.2.146.72:
   /usr/sap/ORC/home/install/cert.cfg
   ```

3. Change the `cluster_port parameter` value to 11117 and update the `certificate_ config_path` parameter value to the correct path of the certificate in the respective paramfile of each of the N nodes.

Source-Side Deduplication for SAP HANA

The Cohesity source-side deduplication provides storage efficient deduplication that enhances your SAP HANA environment with faster backups and reduced network traffic between SAP HANA and Cohesity cluster.

Enabling this functionality shifts the deduplication processes to the SAP HANA node and reduces the CPU consumption on the Cohesity cluster thus allowing further scaling to support more workloads. Additionally, source-side deduplication plugin integrates with Cohesity's built-in deduplication and minimizes data traffic by moving only the new data blocks.

Advantages of the source-side deduplication plugin:

- Faster backups.
- Reduces network traffic between SAP HANA nodes and Cohesity cluster.
- Enables more frequent backups resulting in higher RPOs.
- Offloading deduplication reduces the CPU consumption and enables the Cohesity cluster to run more backup workload sessions.

You can enable source-side deduplication during a backup or recovery process.

To enable source-side deduplication during backup:

Edit the `cohesity_param_file` located at `<backint_install_directory>`. Set the `enable_dedup_write` parameter to `true`.

> **Note:** By default, source-side deduplication is disabled during the backup process.

Alternate Restore of SAP HANA Database

Using the SAP HANA Studio, you can perform an alternate restore of the SAP HANA database backed up on the Cohesity cluster using the Cohesity BACKINT Secure Connector Service. For example, you can restore the SAP HANA database to a new or active database on the same or different SAP HANA system.

Restore to a New or Active Database on the Same SAP HANA System

You can restore the SAP HANA database backed up on the Cohesity cluster to a new or active database on the same SAP HANA system.

Prerequisites

Before you restore the SAP HANA database, ensure that:

- You back up the database using either the SAP HANA Studio, hdbsql, or the Cohesity Remote Adapter. For instructions to back up the database using the Cohesity Remote Adapter, see Create a Remote Adapter Job for SAP HANA Database.

> **Note:** In this topic, for ease of understanding, the database and the table backed up in the Cohesity cluster are referred to as `ORCT2` and `DEMO_TEXTDB`, respectively.
>
> 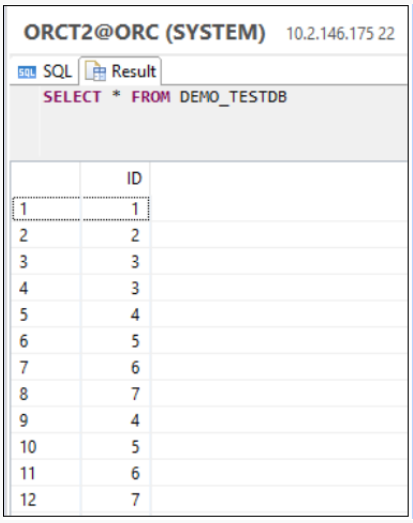
>

- If you plan to restore the database to a new database, create a new database, for example, `ORCT3`.

  Sample Command:

  ```
  hdbsql SYSTEMDB=> create database ORCT3 system user password
  Fr8shst8rt;
  0 rows affected (overall time 80.351254 sec; server time 80.347960 sec)
  ```

To restore the database to a new or active database on the same SAP HANA system:

1. Log on to the SAP HANA system..
2. Verify the status of the source and target databases (for example, `ORCT3`) using the following command:

   > **Note:** The target database can be a new or active database based on your requirement.

   ```
   hdbsql SYSTEMDB=> select * from "SYS"."M_DATABASES";
   | DATABASE | DESCRIPTION     | ACT | ACTIVE_S | OS_USER  | OS_GROUP |
   RESTART | F |
   | -------- | --------------- | --- | -------- | -------- | -------- | -
   ------ | - |
   ```

```
| SYSTEMDB  | SystemDB-ORC-22 | YES |          |          |          |
DEFAULT | ? |
| ORCT1     | ORC-22          | YES |          |          |          |
DEFAULT | ? |
| ORCT2     |                 | YES |          |          |          |
DEFAULT | ? |
| ORCT3     |                 | YES |          |          |          |
DEFAULT | ? |
```

Ensure that the target database shows the **ACTIVE** status as **Yes**.

3. Log on to the SAP HANA Studio.
4. Right-click on the source database (`SYSTEMDB@ORC (SYSTEM)`), select **Backup and Recovery** > **Recover Tenant Database**.

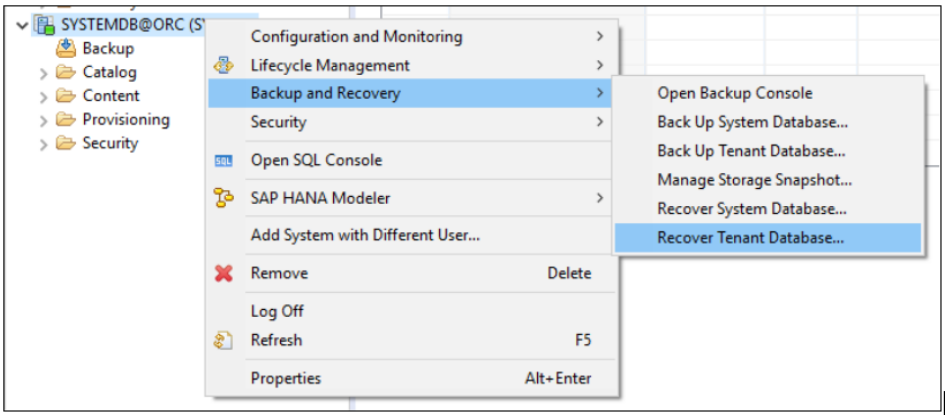

5. Select the target database (for example, `ORCT3`) on the **Specify tenant database** window and click **Next**.

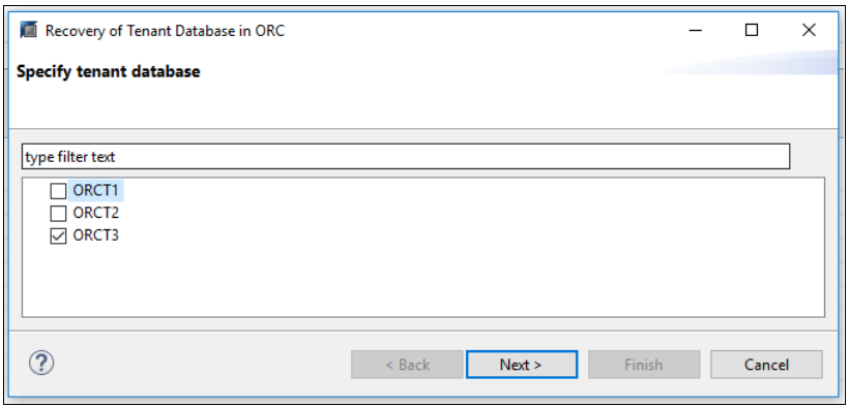

6. Select **Recover the database to a specific data backup** and click **Next** on the **Specify Recovery Type** window.

7. On the **Specify Backup Location** window, select the **Backint System Copy** checkbox, and in the **Source System** field, enter the source database name and the SAP HANA SID in the `DBNAME@SID` format (for example, `ORCT2@ORC`). Then, click **Next**.

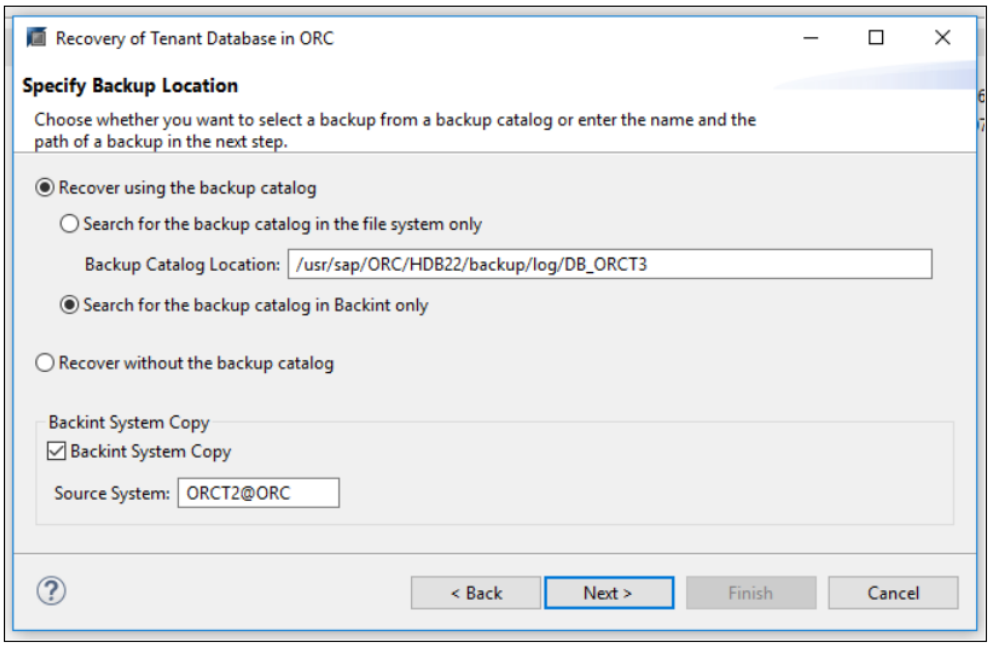

8. On the **Select a Backup** window, select a backup from the backup catalog list that you want to recover to the SAP HANA database and click **Check Availablity**.

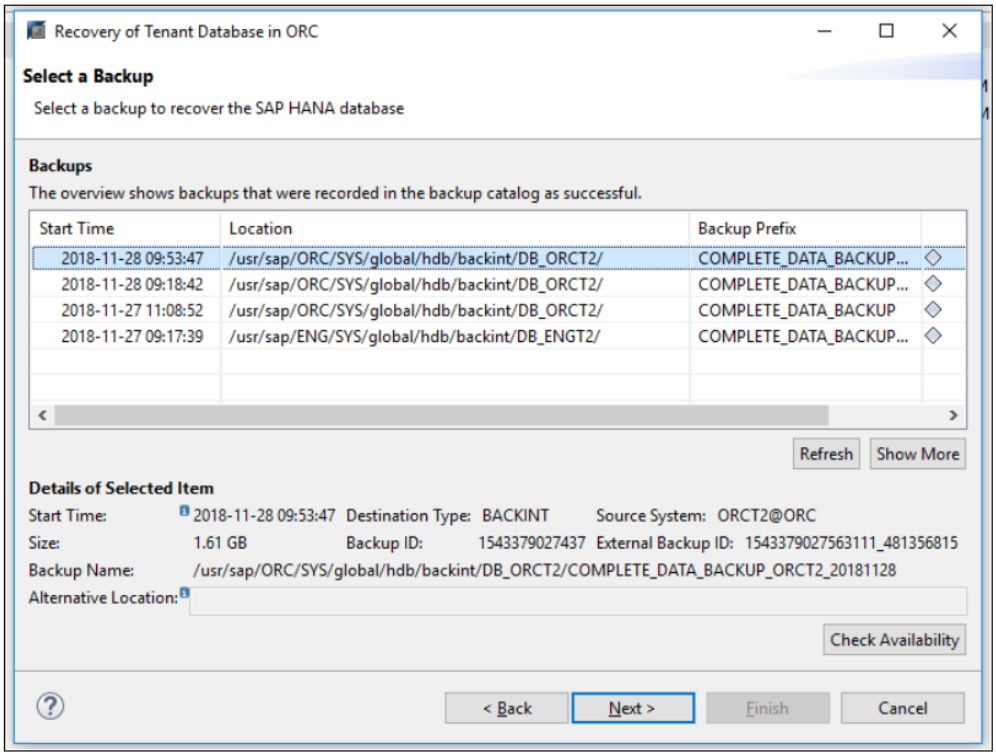

Once the selected backup is **Available**, click **Next**.

9. On the **Other Settings** window, select the required options based on your requirements and click **Next**.

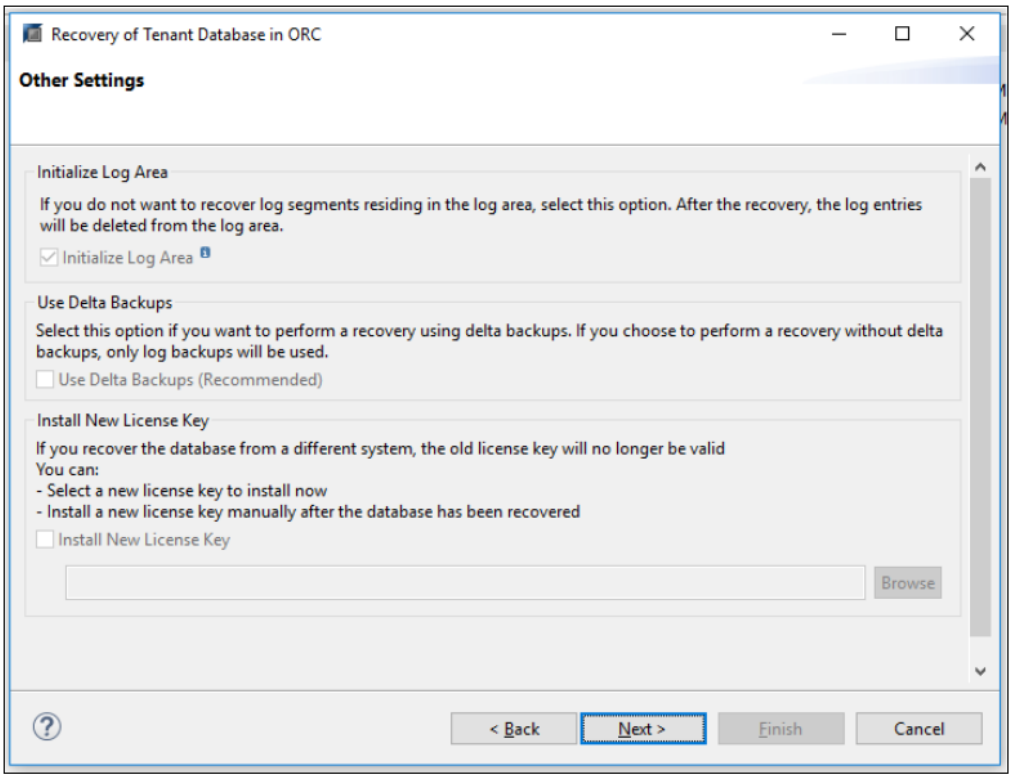

10. On the **Review Recovery Settings** window, review the recovery settings and click **Finish** to start the recovery.

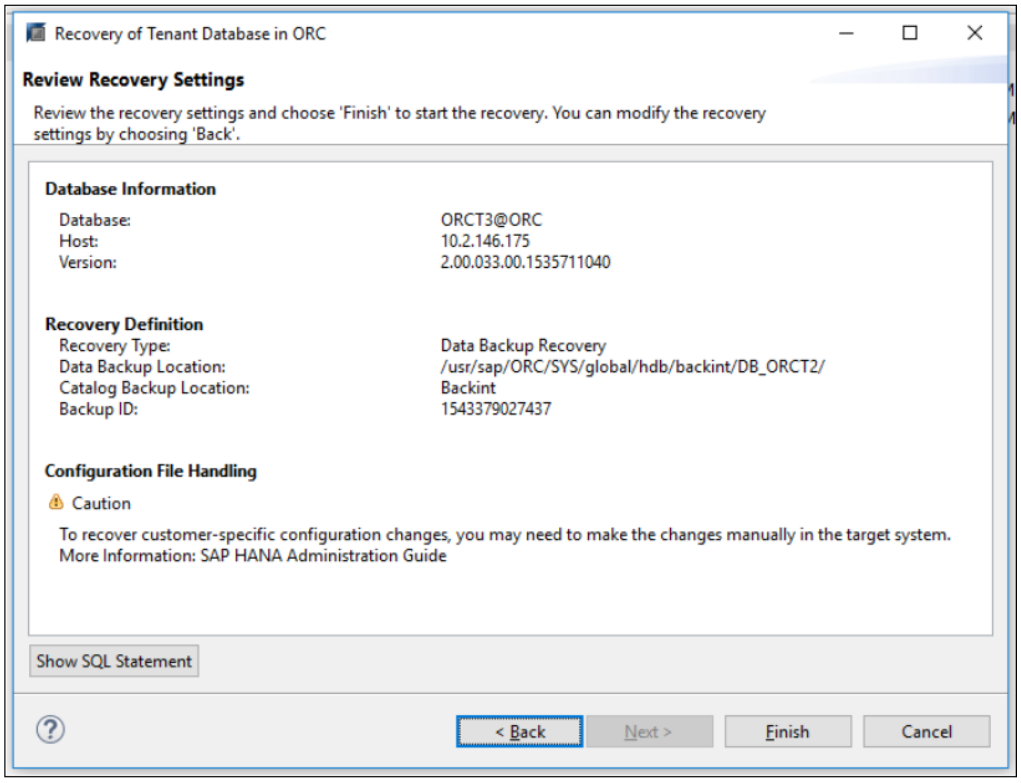

11. On the **Recovery Progress Information** window, verify the recovery progress and click **close** the recovery is completed.

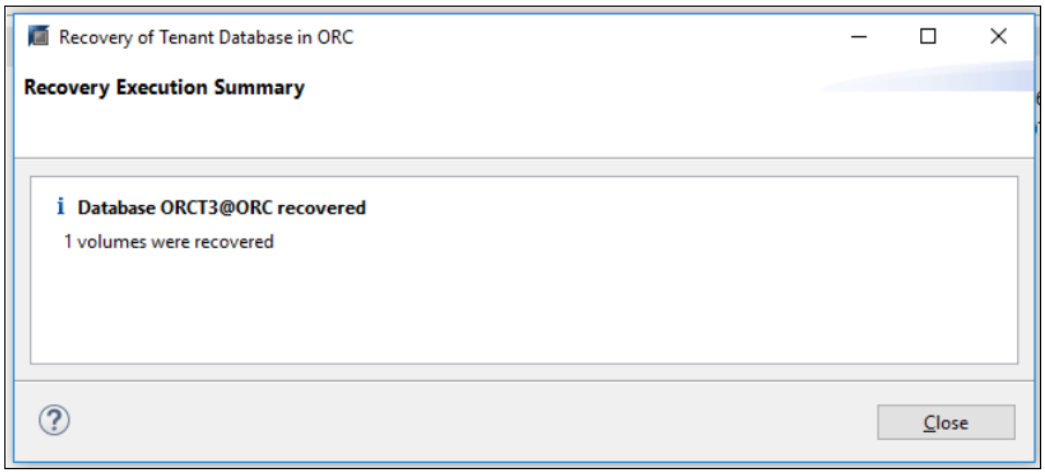

The data is recovered on the target database as shown below:

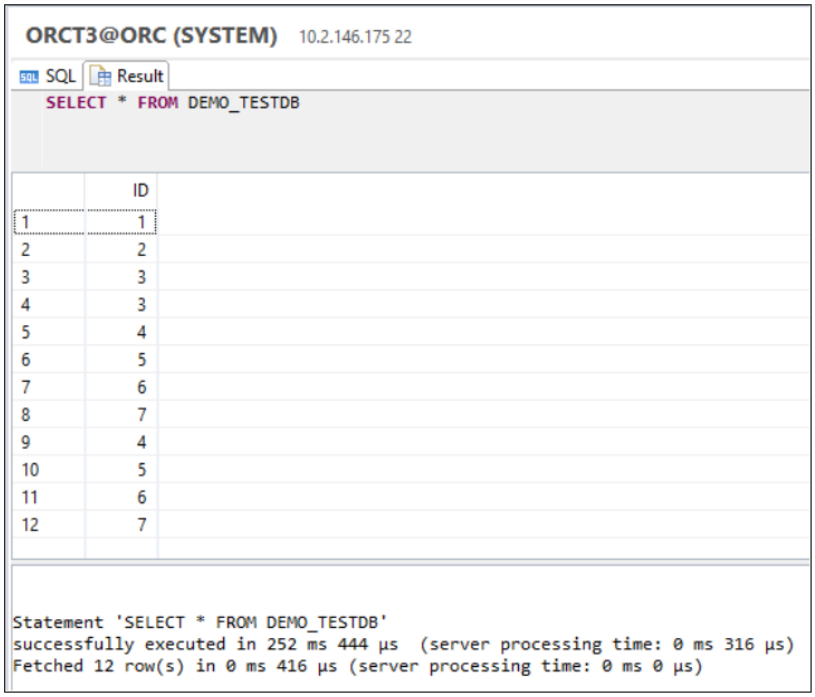

Restore to a New or Active Database on a Different SAP HANA System

You can restore the SAP HANA database backed up on the Cohesity cluster to a new or active database on a different SAP HANA system.

Prerequisites

Before you restore the SAP HANA database, ensure that:

- Install the Cohesity BACKINT Secure Connector service on the target SAP HANA server. For more information, see Install the Cohesity BACKINT plugin.
- From the Cohesity UI, add the IP address of the target SAP HANA node to the custom allowlist of your Cohesity view configuration. For more information, see Create a Cohesity View with IP Allowlist for SAP HANA node.
- If you regularly perform scheduled SAP HANA backups using Cohesity Remote Adapter then perform a backup now or use a recent backup for restore. For more information, see Create a Remote Adapter Job for SAP HANA Database. Otherwise you can use SAP HANA Studio, SAP HANA Cockpit or hdbsql to perform a backup.
- From the SAP HANA Cockpit, take a backup of the source database to the Cohesity cluster to ensure you have correctly installed the Cohesity BACKINT Secure Connector service. In this topic, for ease of understanding, the database and the table backed up on the Cohesity cluster are referred to as ENGT2 and DEMO_TESTDB, respectively.

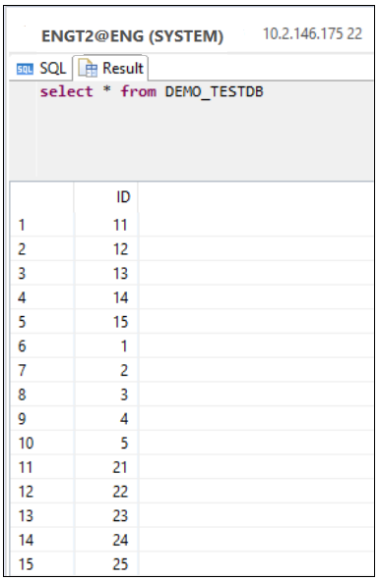

	ID
1	11
2	12
3	13
4	14
5	15
6	1
7	2
8	3
9	4
10	5
11	21
12	22
13	23
14	24
15	25

- If you plan to restore the database to a new database, create a new database, for example, `ORCT3`.

  Sample Command:

```
hdbsql SYSTEMDB=> create database ORCT3 system user password
Fr8shst8rt;
0 rows affected (overall time 80.351254 sec; server time 80.347960 sec)
```

To restore the database to a new or active database on a different SAP HANA system:

1. Log on to the SAP HANA system.
2. Verify the status of the source and target databases (for example, `ORCT3`) using the following command:

   > **Note:** The target database can be a new or active database based on your requirement.

```
hdbsql SYSTEMDB=> select * from "SYS"."M_DATABASES";
| DATABASE | DESCRIPTION     | ACT | ACTIVE_S | OS_USER  | OS_GROUP |
RESTART | F |
| -------- | --------------- | --- | -------- | -------- | -------- | -
------ | - |
| SYSTEMDB | SystemDB-ORC-22 | YES |          |          |          |
DEFAULT | ? |
| ORCT1    | ORC-22          | YES |          |          |          |
DEFAULT | ? |
```

```
| ORCT2    |                  | YES |           |           |           |
DEFAULT | ? |
| ORCT3    |                  | YES |           |           |           |
DEFAULT | ? |
```

Ensure that the target database shows the **ACTIVE** status as **Yes**.

3. Configure the source database's parameter file in the target database (`ORCT3`). Perform the following:
   a. Copy the parameter file (`cohesity_param_file`) from the source SAP HANA system to the target SAP HANA system.
   b. Update the **global.ini** file in the target SAP HANA system. You must update the following parameter files using the mentioned commands:
      - Update the **catalog_backup_parameter_file** file. Run the following command:

        ```
        hdbsql ORCT3=> ALTER SYSTEM ALTER CONFIGURATION
        ('global.ini','SYSTEM') SET ('backup','catalog_backup_
        parameter_file') = '/usr/sap/ORC/home/install/cohesity_param_
        file_$(SAPSYSTEMNAME)' WITH RECONFIGURE;
        ```

      - Update the **data_backup_parameter_file** file. Run the following command:

        ```
        hdbsql ORCT3=> ALTER SYSTEM ALTER CONFIGURATION
        ('global.ini','SYSTEM') SET ('backup','data_backup_parameter_
        file') = '/usr/sap/ORC/home/install/cohesity_param_file_
        $(SAPSYSTEMNAME)' WITH RECONFIGURE;
        ```

      - Update the **log_backup_parameter_file** file. Run the following command:

        ```
        hdbsql ORCT3=> ALTER SYSTEM ALTER CONFIGURATION
        ('global.ini','SYSTEM') SET ('backup','log_backup_parameter_
        file') = '/usr/sap/ORC/home/install/cohesity_param_file_
        $(SAPSYSTEMNAME)' WITH RECONFIGURE;
        ```

      - Update the **catalog_backup_using_backint** file. Run the following command:

        ```
        hdbsql ORCT3=> ALTER SYSTEM ALTER CONFIGURATION
        ('global.ini','SYSTEM') SET ('backup',''catalog_backup_using_
        backint') = '/usr/sap/ORC/home/install/cohesity_param_file_
        $(SAPSYSTEMNAME)' WITH RECONFIGURE;
        ```

- Update the **log_backup_using_backint** file. Run the following command:

```
hdbsql ORCT3=> ALTER SYSTEM ALTER CONFIGURATION
('global.ini','SYSTEM') SET ('backup',''log_backup_using_
backint') = '/usr/sap/ORC/home/install/cohesity_param_file_
$(SAPSYSTEMNAME)' WITH RECONFIGURE;
```

Using the SAP HANA Cockpit, verify if the default value is replaced with `usr/sap/ORC/home/install/cohesity_param_file_$(SAPSYSTEMNAME)`.

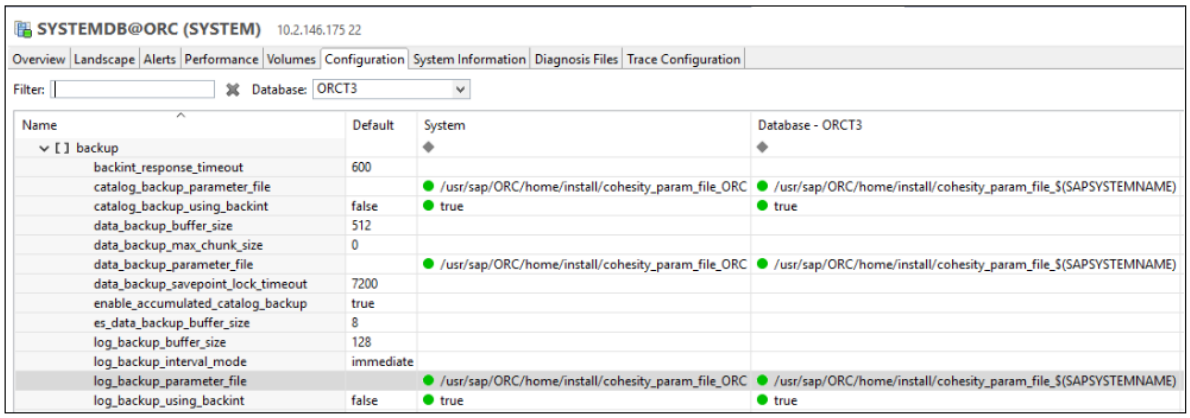

4. In SAP HANA Studio, right-click on the source database and click **Backup and Recovery** > **Recover Tenant Database**.
5. Select the target database (for example, `ORCT3`) on the **Specify tenant database** window and click **Next**.

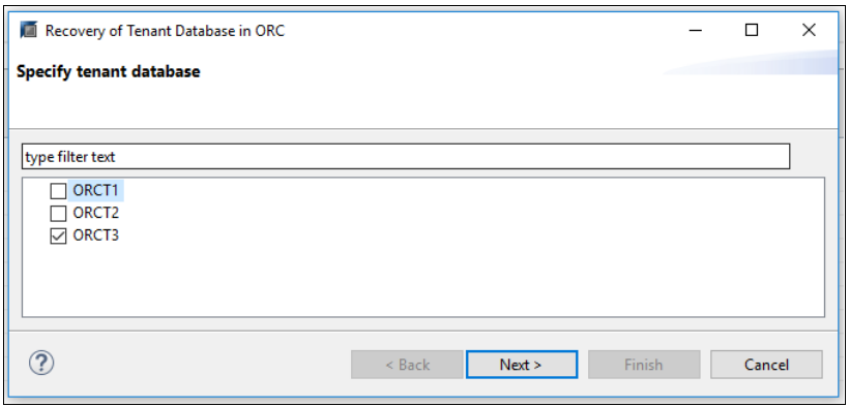

6. Select **Recover the database to a specific data backup** and click **Next** on the **Specify Recovery Type** window.

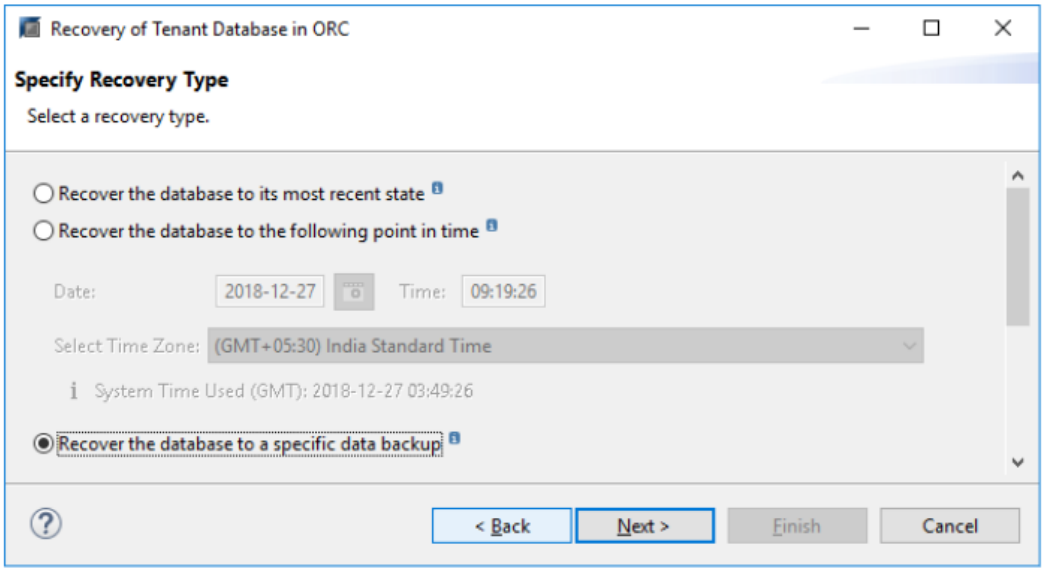

7. On the **Specify Backup Location** window, select the **Backint System Copy** checkbox and in the **Source System** field, enter the source database name and the SAP HANA SID in the `DBNAME@SID` format (for example, `ENGt2@ENG`). Then, click **Next**.

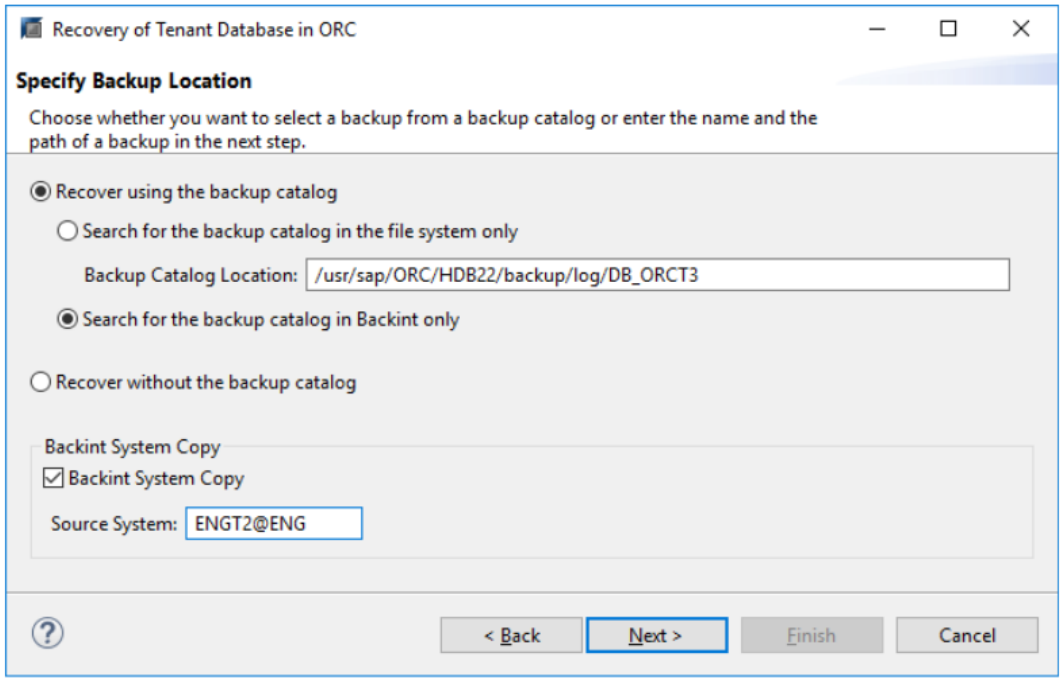

8. On the **Select a Backup** window, select a backup from the backup catalog list that you want to recover to the SAP HANA database and click **Check Availablity**.

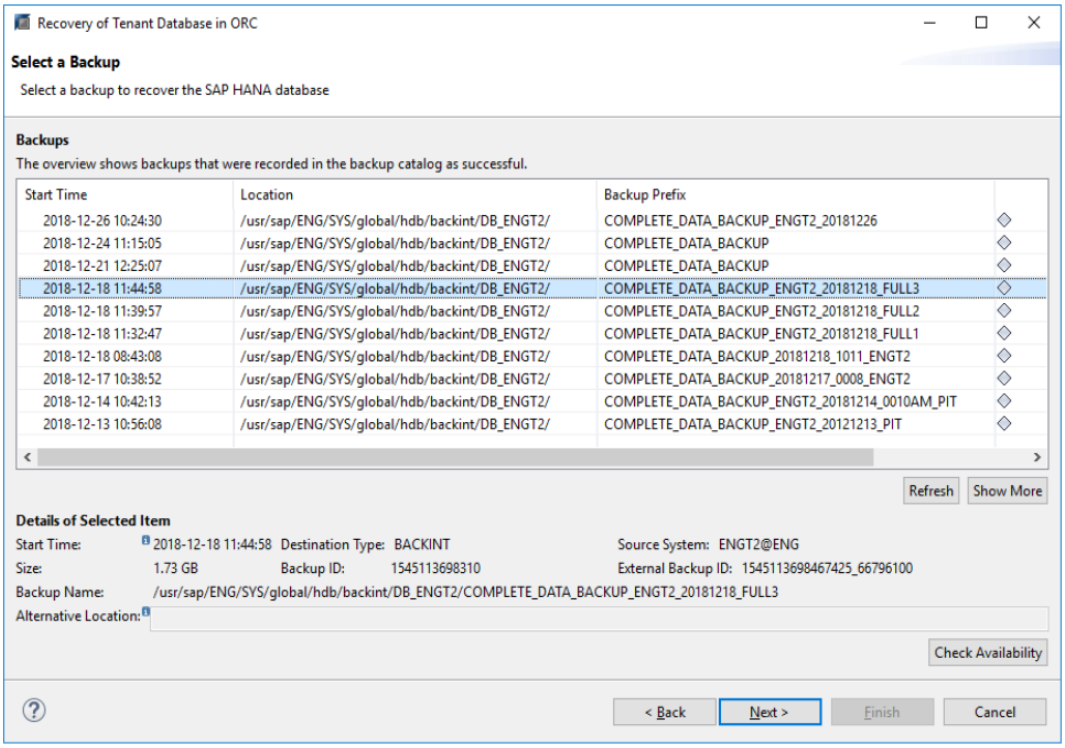

Once the selected backup is **Available**, click **Next**.

9. On the **Other Settings** window, select the required options based on your requirements and click **Next**.

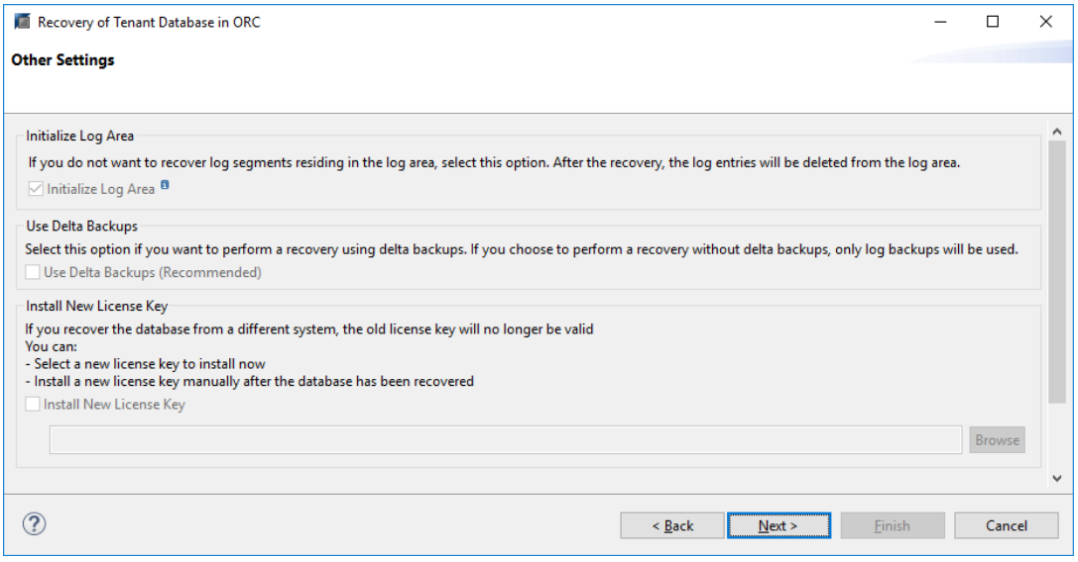

10. On the **Review Recovery Settings** window, review the recovery settings and click **Finish** to start the recovery.

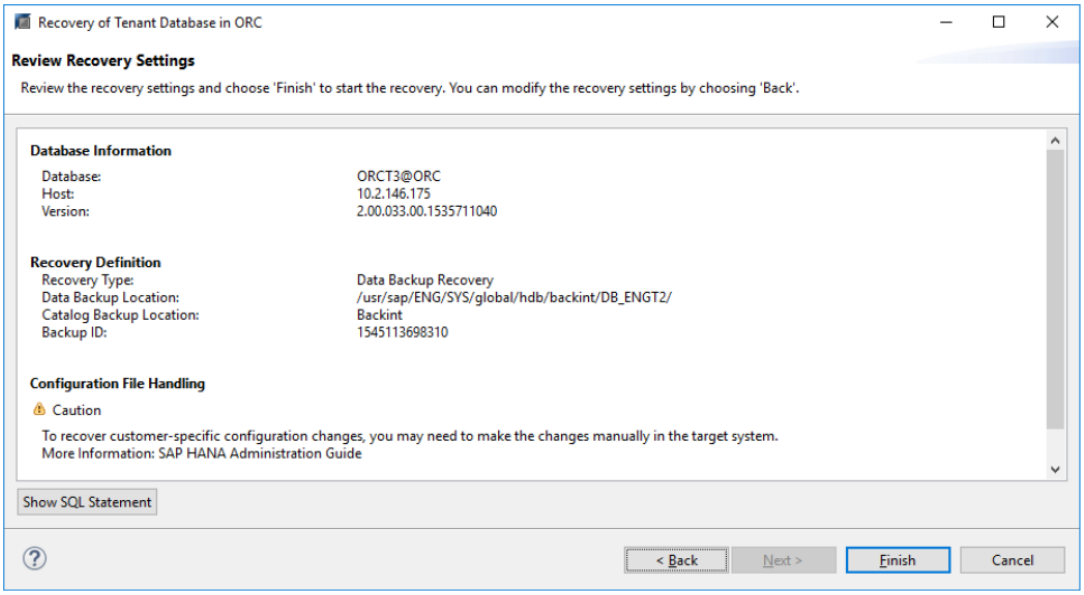

11. On the **Recovery Progress Information** window, verify the recovery progress and click **close** when the recovery is completed.

The data is recovered on the target database as shown below:

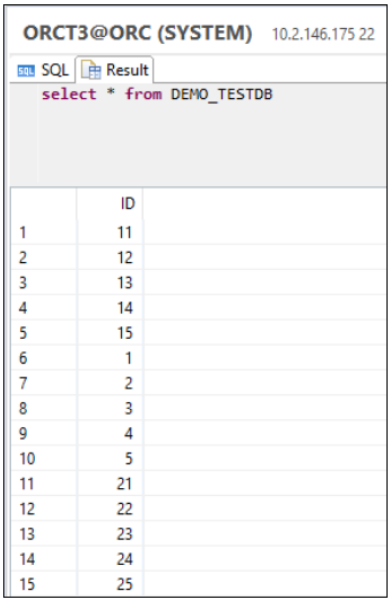

	ID
1	11
2	12
3	13
4	14
5	15
6	1
7	2
8	3
9	4
10	5
11	21
12	22
13	23
14	24
15	25

Replicate View and Restore SAP HANA Database

You can replicate the SAP HANA database data captured on a Cohesity cluster to a remote Cluster by performing the following tasks:

> **Note:** To restore from archival or replication, one needs to disable auto log backups.

1. From the Cohesity cluster (source) where the SAP HANA data is captured, create a connection to the receiving (target) cluster where the view will be replicated.

   For information about registering a connection to the target cluster, see Create a Connection to the Remote Cohesity Cluster.
2. From the target cluster, create a connection back to the source cluster.
3. From the source cluster, create a Remote Adapter Protection Group with replication enabled in the policy.

   For information about creating Remote Adapter job with replication enabled, see Add or Edit a Remote Adapter Protection Group.

Restore SAP HANA DB from the Replicated View

To restore the replicated data back to a SAP HANA Database, perform the following steps:

1. On the target cluster, create a clone of the replicated View. For information about cloning a View, see Clone a View.
2. Log in to the HANA database where you want to restore the replicated data. In the parameter file, replace the existing cluster IP address (source cluster IP address) with the target cluster IP address
3. Perform the restore operation. For more information about SAP HANA restore, see Alternate Restore of SAP HANA Database.

Archive View and Restore SAP HANA Database

You can archive the Snapshots of SAP HANA databases captured on a Cohesity cluster to an external target by performing the following tasks:

> **Note:** To restore from archival or replication, one needs to disable auto-log backups.

1. From the Cohesity cluster where the SAP HANA data is captured, register an external target. For more information, see Register an External Target.
2. Create a remote Adapter protection with archival enabled in the policy. For information about creating Remote Adapter job with archival enabled, see Add or Edit a Remote Adapter Protection Group.

Retrieve Archived Data to an Alternate Cohesity cluster

CloudRetrieve provides the ability to download data that was archived from a Cohesity cluster to an alternate Cohesity cluster. For more information see, About CloudRetrieve.

To retrieve the SAP HANA data from the external target to another Cohesity cluster, perform the following steps:

1. Register the external target that contains the data you want to download to the alternate Cohesity cluster.
2. Submit a search for Protection Groups that contain the data you want to download from the cloud external target. The search may take a few hours to complete. If the job to be downloaded is from a particular Cohesity cluster, specify the job name and cluster name to narrow the results.
2. If the external target's encryption key file is managed manually, upload the key file to the alternate Cohesity cluster before triggering the search. See Create a CloudRetrieve Search Task for more details.
3. Download a snapshot from any of the Protection Groups in the search results. You can also build a local index of Protection Groups that ran during a specified date range. Downloading can take a few hours to complete. See Create a CloudRetrieve Download Task for more details.
4. Recover or clone any downloaded objects using the Cohesity recovery or clone workflows. See CloudRetrieve Downloaded Protection Groups for more details.

Restore HANA DB from the Archived View

You can restore the retrieved data from the remote Cohesity cluster to the SAP HANA database by performing the following steps:

To restore the replicated data back to a SAP HANA Database, perform the following steps:

1. On the remote Cohesity cluster, create a clone of the retrieved View. For information about cloning a View, see Clone a View.
2. Login to the SAP HANA database where you want to restore the retrieved data. In the parameter file, replace the existing Cohesity cluster IP address (primary clusterCohesity cluster IP address) with the remote Cohesity cluster IP address.
3. Perform the restore operation. For more information about SAP HANA restore, see Alternate Restore of SAP HANA Database.

SAP HANA Database FAQ

1. Which SAP HANA versions are supported by Cohesity SAP HANA BACKINT plugin on 6.2 DataProtect?
   SAP HANA Platform Edition 2.0.
2. Is Cohesity SAP HANA Certified?
   Yes, Cohesity is certified on SAP HANA Platform Edition 2.0.

3. Which SAP HANA Environments are supported by the Cohesity SAP HANA BACKINT plugin on 6.2 DataProtect onwards?
   Single Tenancy and Multitenancy environments are supported.
4. Which OS version is supported with the Cohesity SAP HANA BACKINT plugin on 6.2 DataProtect onwards?
   All versions that are supported by SAP HANA Platform Edition (Intel platforms).
5. What is Cohesity SAP HANA BACKINT plugin?
   It is an SAP HANA agent installed on the SAP HANA appliance.
6. Does the Cohesity SAP HANA BACKINT plugin leverages BACKINT?
   Yes, the Cohesity SAP HANA BACKINT plugin leverages SAP HANA BACKINT to protect the SAP HANA appliance.
7. Does the Cohesity SAP HANA BACKINT plugin install requires a reboot of the SAP HANA nodes?
   No, the Cohesity SAP HANA BACKINT plugin (Cohesity SAP HANA agent) does not require a reboot of the SAP HANA nodes or restart of the SAP HANA instances.
8. Does the Cohesity SAP HANA BACKINT plugin support client-side dedupe?
   Yes, it is supported from the 6.4.1x release onwards.
9. Is the Cohesity SAP HANA BACKINT plugin leverages NFS?
   Yes. The Cohesity SAP HANA BACKINT plugin for IBM PowerPC systems leverages NFS. But, the Cohesity SAP HANA BACKINT plugin for Intel x86 based systems does not leverage the NFS protocol. It uses remote snapFS.
10. If I have a Multi-node SAP HANA system, do I need to install the Cohesity SAP HANA BACKINT plugin in all the nodes?
    Yes, this would allow the Cohesity SAP HANA BACKINT plugin to take advantage of the high availability functionality when any of the SAP HANA nodes are down to service the Cohesity SAP HANA BACKINT plugin.
11. Does the SAP HANA Studio require a restart if the `global.ini` was modified via the CLI?
    Yes. If you have modified `global.ini` for the Cohesity SAP HANA BACKINT plugin for Intel x86 based systems then the SAP HANA Studio requires a restart to apply the changes.
12. Does SAP HANA have any dependencies with Cohesity Linux or Oracle adapters?

    No, Cohesity SAP HANA BACKINT plugin is independent of the Linux file/block or Oracle adapters. SAP HANA BACKINT plugin does not need the Cohesity Linux agent service to be installed for SAP HANA protection. However, both the SAP HANA BACKINT plugin and the Linux agent can coexist together for their respective functionalities.
13. Do we need to create a folder, or use the SAP HANA View created for the Remote adaptor job?

No. Do not manually create any folders inside the View dedicated to the Remote adaptor job, as the Purge script will fail if any folders or files are created manually within the SAP HANA View.

SAP on Oracle

Cohesity provides a simple, fast, and cost-effective backup, recovery, and data management solution for Oracle databases running on the large-scale SAP ERP on Linux operating systems.

SAP uses backup BR*TOOLS to back up Oracle databases. Cohesity provides the following methods to backup Oracle database leveraging the BR*TOOLS:

- **Backup using BACKINT Secure Connector Service for Oracle** - Cohesity provides an SAP-certified BACKINT Secure Connector Service for Oracle that leverages BR*TOOLs for protecting Oracle databases.
- **Backup using custom Oracle SBT (Secure Backup to Tapes) and BACKINT Secure Connector Service for Oracle** - Cohesity provides a custom SBT extension that you can use along with BACKINT Secure Connector Service to backup Oracle databases.

Using the backup methods you can back up and restore the following for Oracle databases running in an SAP system:

- Oracle database files such as data files (***.dbf**) and control files (***.ctl**)
- Archived redo logs
- Parameter files such as PFILE or SPFILE
- SAP BR*TOOLS files
- SAP files

Core Components

The following are the core components for protecting Oracle databases running on an SAP system:

- Oracle Database
- SAP BR*TOOLS
- Cohesity BACKINT Secure Connector Service for Oracle
- Cohesity Data Cloud (Self-Managed)

As illustrated, you can integrate the Cohesity Data Cloud (Self-Managed) with new or existing Oracle database deployments in SAP by installing the Cohesity BACKINT Secure Connector Service for Oracle.

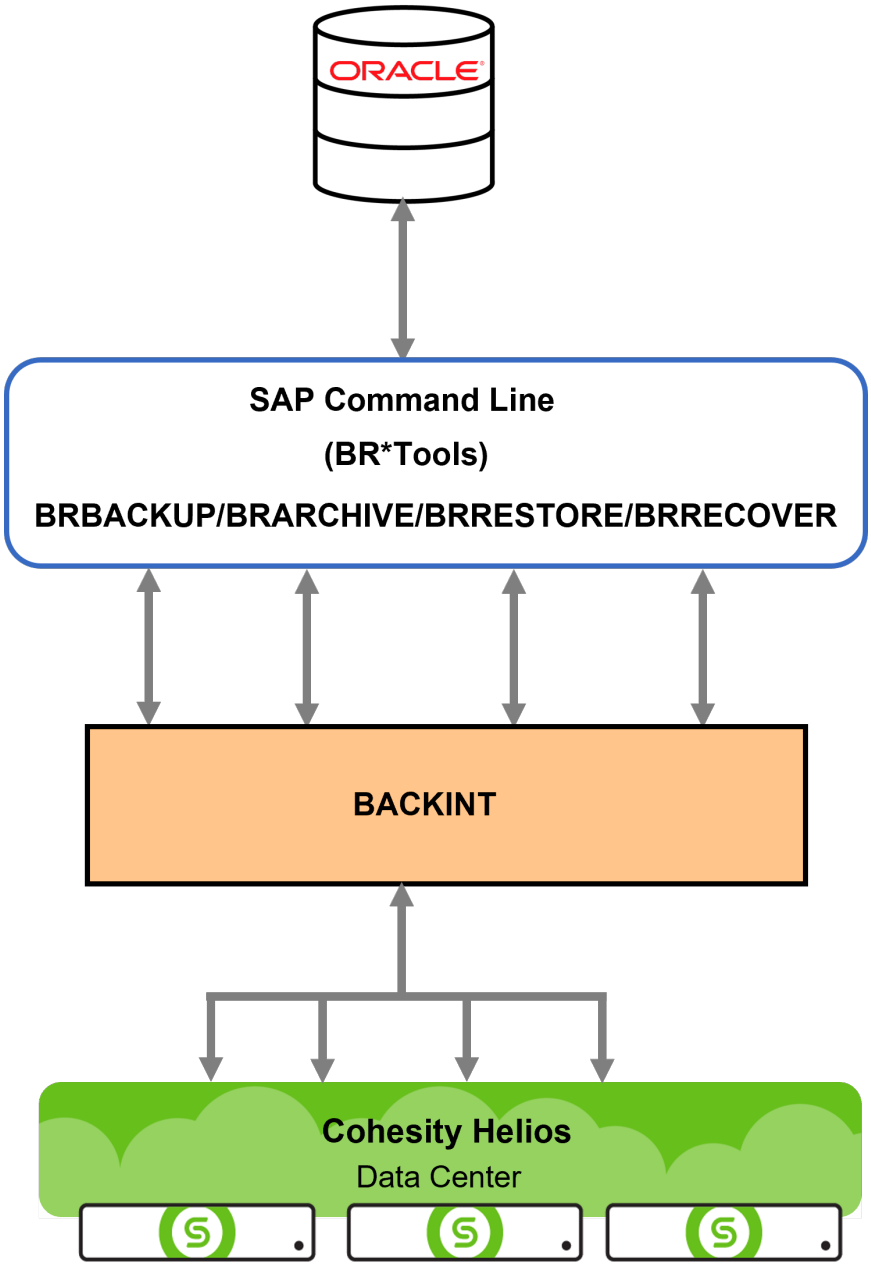

Plan and Prepare

This topic provides information on supported Oracle databases on SAP, and prerequisites for installing the Cohesity BACKINT Secure Connector Service for Oracle.

Supported Versions

For more information on supported SAP on Oracle versions, see Supported Software versions.

> **Note:** Backing up SAP on Oracle databases using custom Oracle SBT and BACKINT Secure Connector Service is supported only on SAP on Oracle databases running on Linux (x86-64) system.

Prerequisites

Make sure that you:

- Install OpenJDK JRE 1.8 or later, on the SAP system.
- Create a Cohesity view with the **DefaultStorageDomain** options. For more information, see Create a Cohesity View with IP Allowlist for the SAP Node.
- Download and copy the Cohesity BACKINT Secure Connector Service for Oracle to the SAP system. For more information, see Download the BACKINT Secure Connector Service for Oracle Installer
- Check if the SAP node has an `<SID>adm` user account; if not, create an Operating System (OS) user account to own the backup files on the Cohesity view. The format of the user account must be `<SID>adm`, where `SID` represents the SAP System Identification. It is a unique three-character code available for every R/3 installation (SAP system) that consists of a database server and several application servers.
- Use "bash" shell to install the Cohesity BACKINT Secure Connector Service for Oracle and to run backups and restore workflows of SAP on Oracle databases.
- Ensure that the NTP time is in sync between the client and the Cohesity cluster. Perform the following on the client:
  - Verify if NTP is enabled and synchronized in the client using the following command:

    ```
    timedatectl
    ```

  - If the NTP time is not synchronized, use the following command to synchronize the NTP time

    ```
    timedatectl set-ntp true
    ```

- The BACKINT Secure Connector Service for Oracle can be installed using a `root` user account or using a user account with `sudo` user privileges. If you are installing the BACKINT Secure Connector Service for Oracle using a `sudo` user account, for example, `oracle`. Ensure that you add the following entries in the **sudoers** file of the `oracle` account:

  ```
  Defaults:oracle !requiretty

  oracle ALL=(ALL) NOPASSWD:ALL
  ```

- If you are using the RPM-based installer, ensure that you configure the following environment variables on the SAP on Oracle node(s):

Environment Variable	Description	Command
ORACLE_HOME	The path to the Oracle database installation directory.	export ORACLE_HOME=<Oracle-home-Directory-Path>
SAPDATA_HOME	The directory of the database files.	export SAPDATA_HOME=<SAP-Home-Directory-Path>
SERVICE_USER	The <SID>adm user account that must own the backup files on the Cohesity view. If you have not created the user account, then create an Operating System (OS) user account in the following format: <SID>adm. Where <SID> is a unique three-character code called SAP System Identification (SID). It is available for every R/3 installation (SAP system) that consists of a database server and several application servers.	export SERVICE_USER=qqqadm Where qqqadm is the <SID>adm user account and qqq is the SAP System Identification (SID).
ORACLE_SID	The System ID (SID) of the Oracle database instance.	export ORACLE_SID=<Oracle_SID>
CLUSTER_IPS	A comma-separated list of virtual IP addresses of the servers in the Cohesity cluster. **Note:** Ensure that you use virtual IP addresses (VIPs) instead of physical IP addresses (node IP addresses). The physical IP address is bound to a node, and if the node goes down, then the IP address is no longer accessible and causes the backup to fail. Hence, Cohesity recommends using VIP because of the resilience; if the node fails, it won't cause backup failures.	export CLUSTER_IPS=<IP1>,<IP2>,<IP3>,...
VIEW_NAME	The name of the Cohesity view that you have created for SAP on Oracle. For more information, see Create a Cohesity View with IP Allowlist for the SAP Node. **Note:** Ensure that the Cohesity view name does not contain the path of the Cohesity view or Cluster name or IP address.	export VIEW_NAME=<view name>

Environment Variable	Description	Command
SERVICE_TYPE	The type of service. By default set to REMOTESNAPFS. Allowed values are: REMOTESNAPFS or NFS.	

Create a Cohesity View with IP Allowlist for the SAP Node

Create a CohesityView to provide a storage location in a storage domain on the Cohesity cluster for the SAP on Oracle nodes. For more information on views, see Cohesity Views.

To provide effective authorization and prevent rogue access to the Cohesity cluster, the SAP node IP address must be added to the allowlist in the Cohesity view configuration.

Adding Custom Allowlist for the Cohesity View

While creating a Cohesity View, you have an option to add an allowlist for the View. Ensure that you add the SAP HANA node IP address to the allowlist for the Cohesity View.

To add subnets to the allowlist of the Cohesity view:

1. In the **Create View** page, under the **More Options**, expand **Security**.
2. In the **IP Allowlist** section, perform the following:
    1. Select **Override Global Allowlist** and click **Add Allowlist**.
    2. In the **Subnet IP** and **Subnet Mask** fields respectively, enter the Subnet IP address and Subnet Mask to specify a range of systems that should have permissions to access and mount the Cohesity View.
    3. In the **NFS Permissions**, **SMB Permissions**, and **S3 Permissions** drop-down lists, select the NFS, SMB, and S3 permissions respectively.
    4. In the **NFS Root Squash** and **NFS All Squash** drop-down lists, specify whether to turn on NFS Root or NFS All squash respectively.
    5. In the **Description** field, enter a brief description.
    6. Click **Add**.
    7. Repeat these steps to add additional subnets to the allowlist of the Cohesity view.

Note:

- To use the IP allowlist feature, you need to upgrade the Cohesity cluster to 6.3 release or later.
- If the Override Global Allowlist toggle is off, this View can only be accessed and mounted by systems specified by the Global Cohesity cluster Allowlist.
- In case of a multinode HANA setup, specify all nodes individually in the global IP allowlist.

Download the BACKINT Secure Connector Service for Oracle Installer

The BACKINT Secure Connector Service for Oracle is available as a script installer for x86-64 systems.

To download the installer:

1. Navigate to **Data Protection > Sources**, and click **Download Cohesity Agent**.
2. In the **Download Agents** window, click **SAP on Oracle** and then select **Script installer** or **RPM** based on your requirement.
   The installer of BACKINT Secure Connector Service for Oracle is downloaded to your local machine.
3. Copy the installer to the node or nodes of the SAP deployment.

Install the BACKINT Secure Connector Service for Oracle

Cohesity BACKINT Secure Connector Service for Oracle is an SAP-certified service that enables you to establish communication between the Cohesity cluster and the SAP on Oracle nodes.

The BACKINT Secure Connector Service for Oracle can be installed using a `root` user account or using a user account with `sudo` user privileges. Cohesity supports IPv6 for SAP Oracle in x86_64 architectures with gRPC data transfer. If you are installing the BACKINT Secure Connector Service using a `sudo` user account, for example, `oracle`. Ensure that you add the following entries in the `sudoers` file of the `oracle` user account:

```
Defaults:oracle !requiretty
oracle ALL=(ALL) NOPASSWD:ALL
```

Prerequisites

- Before installing the BACKINT Secure Connector service for Oracle you must set the `ORACLE_HOME` and `SAPDATA_HOME` exports on the SAP Oracle host.

  Syntax:

- `export ORACLE_HOME=<oracle database>`

  Example:

  ```
  export ORACLE_HOME=/oracle/SAL/121
  ```

- `export SAPDATA_HOME=<oracle database>`

  Example:

  ```
  export SAPDATA_HOME=/oracle/SAL
  ```

- Run the following command on the SAP Oracle node to synchronize time between the SAP Oracle server and the Cohesity NTP server:

  ```
  timedatectl set-ntp true
  ```

Install the BACKINT Secure Connector Service using the Script Installer

On the SAP node, perform the following to install the BACKINT Secure Connector Service using the script installer, `cohesity_secure_connector_service_sap_oracle_installer_sap_oracle_java_installer`.

To install the BACKINT Secure Connector Service for Oracle:

1. Log on to the SAP node's bash CLI.
2. Navigate to the directory where you have downloaded the script installer.
3. Make the Secure Connector Service installer (`cohesity_secure_connector_service_sap_oracle_installer_sap_oracle_java_installer`) executable using the following command:

   ```
   [oracle@sapnet ~]$ chmod +x cohesity_secure_connector_service_sap_oracle_installer_sap_oracle_java_installer
   ```

4. Install the Secure Connector Service for Oracle using the following command:

   ```
   cohesity_secure_connector_service_sap_oracle_installer_sap_oracle_java_installer -- --install -s <SID> -a <comma_separated_cohesity server_ip_addresses> -v <view_name> -S ora<SID>
   ```

   Where:

   - `-s <SID>` - SAP System Identification (SID) is a unique three-character code that is available for every R/3 installation (SAP system) that consists of a database server and several application servers.
   - `-a <comma_separated_cohesity server_ip_addresses>` - The comma-separated list of virtual IP addresses of the servers in the Cohesity Cluster.

> **Note:** You can also specify the hostname instead of the IP addresses to perform the mount. For successful mount operation, you must specify the proper mapping between the IP addresses and the hostname in the `/etc/hosts` file.

For example:

```
[oracle@sapnet ~]$ nslookup cluster.cohesity.com
Server:         xxx.xx.xx.50
Address:        xxx.xx.xx.50#3
Name:   cluster.cohesity.com
Address: xxx.xx.xx.244
Name:   cluster.cohesity.com
Address: xxx.xx.xx.246
Name:   cluster.cohesity.com
Address: xxx.xx.xx.247
Name:   cluster.cohesity.com
Address: xxx.xx.xx.248
```

- `-v <view_name>` - The Cohesity View that you have created for SAP on Oracle.

The following is a sample output of the successful installation:

```
[oracle@saptest ~]$ ./cohesity_secure_connector_service_6.5.1_sap_oracle_
java_installer -- --install -s SAL -a
10.2.196.217,10.2.196.219,10.2.196.221,10.2.196.223 -v view-oracle
Verifying archive integrity... All good.
Uncompressing Cohesity Backint  100%

=~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
~~=
                  ____ ___  _   _ _____ ____ ___ _______   __

                 / ___/ _ \| | | | ____/ ___|_ _|_   _\ \ / /

                | |  | | | | |_| |  _| \___ \| |  | |  \ V /

                | |__| |_| |  _  | |___ ___) | |  | |   | |

                 \____\___/|_| |_|_____|____/___| |_|   |_|
```

```
=~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~-
~~=
                    ***** NOTE ***** NOTE ***** NOTE *****



      This script based installer is designed for SAP Oracle system like

                       SAP NetWeaver, SAP ECC, and SAP R2.



                    ***** NOTE ***** NOTE ***** NOTE *****

--------------------------------------------------------------------------
--
                 --- S Y S T E M     I N F O R M A T I O N ---

***** DATE AND TIME INFORMATION *****
Wed Jul 22 10:54:28 EDT 2020

***** HOSTNAME INFORMATION *****
   Static hostname: saptest
         Icon name: computer-vm
           Chassis: vm
        Machine ID: 466bc8302ebc4873a6c8541d0854e0fd
           Boot ID: 6e3f19bee42f499d87f65f29b64c91ac
    Virtualization: vmware
  Operating System: Oracle Linux Server 7.7
       CPE OS Name: cpe:/o:oracle:linux:7:7:server
            Kernel: Linux 4.1.12-124.34.1.el7uek.x86_64
      Architecture: x86-64

***** ORACLE DATABASE INFORMATION *****
Oracle Database version : 12.1.0.2.0

***** FILE SYSTEM DISK SPACE USAGE *****
Filesystem           Size  Used Avail Use% Mounted on
devtmpfs             7.7G     0  7.7G   0% /dev
tmpfs                7.8G     0  7.8G   0% /dev/shm
tmpfs                7.8G   45M  7.7G   1% /run
tmpfs                7.8G     0  7.8G   0% /sys/fs/cgroup
/dev/sda1             97G   46G   52G  47% /
/dev/mapper/ol-u01    70G   11G   60G  16% /u01
```

```
/dev/mapper/ol-home   51G   12G   40G  22% /home
tmpfs                1.6G  4.0K  1.6G   1% /run/user/42
tmpfs                1.6G   36K  1.6G   1% /run/user/54321
tmpfs                1.6G     0  1.6G   0% /run/user/0

***** FREE AND USED MEMORY *****
              total        used        free      shared  buff/cache   avail-
able
Mem:       16150368     1559004     4417696     4533196    10173668
9984636
Swap:      17825788           0    17825788

***** SYSTEM UPTIME AND LOAD *****
 10:54:30 up 11:23,  5 users,  load average: 0.31, 0.14, 0.07

***** CURRENTLY LOGGED-IN USERS *****
oracle   :0           2020-07-21 23:31 (:0)
oracle   pts/0        2020-07-21 23:31 (:0)
root     pts/1        2020-07-21 23:58 (sv16-pm-dell6-c1.pm.cohesity.com)
root     pts/2        2020-07-21 23:59 (10.2.54.192)
oracle   pts/3        2020-07-22 10:53 (10.2.252.50)

***** TOP 5 MEMORY-CONSUMING PROCESSES *****
%MEM %CPU COMMAND
 1.7  0.0 ora_mman_sal
 1.6  0.2 gnome-shell
 1.0  0.0 ora_mmon_sal
 0.8  0.0 gnome-software
 0.7  0.0 ora_dbw1_sal
--------------------------------------------------------------------------------
--
==> Installing Cohesity java backint secure connector service for SAP Oracle
--------------------------------------------------------------------------------
--
Are you sure the provided user 'saladm' is correct? Enter y[es]/n[o]: y
Continuing installation...
--------------------------------------------------------------------------------
--
Java Runtime Environment check ..................................... [passed]
NFS client check ................................................... [passed]
Environment variables check ........................................ [passed]
SAP profile files check ............................................ [passed]
Service user OS roles privileges check ............................. [passed]
--------------------------------------------------------------------------------
```

```
--
Architecture Name     :
Operating System Name :
Service User          : saladm
Install Directory     : /usr/sap/SAL/cohesity
Install Log Directory : /tmp/cohesity_oracle_java_backint_installer_
1595429668_51780.log
SAP Database SID      : SAL
Cluster IPs           : 1.1.1.1
View Name             : view1
Sudo Command Prefix   : sudo
-------------------------------------------------------------------------
--
==> Creating the required directories.
Creating install directory '/usr/sap/SAL/cohesity'
Creating scripts directory '/usr/sap/SAL/cohesity/scripts'
Creating profile directory '/usr/sap/SAL/cohesity/profile'
-------------------------------------------------------------------------
--
==> Copying the required files.
Copying 'backint.sh' to install directory '/usr/sap/SAL/cohesity'
Copying 'backint.jar' to install directory '/usr/sap/SAL/cohesity'
Copying 'purge_job.sh' to scripts directory '/usr/sap/SAL/cohesity/scripts'
Copying 'initSID.utl' to profile directory '/usr/sap/SAL/cohesity/profile'
-------------------------------------------------------------------------
--
==> Customizing the required files.
Creating symbolic link from '/usr/sap/SAL/cohesity/backint.sh' to '/us-
r/sap/SAL/SYS/exe/run/backint'
Renaming '/usr/sap/SAL/cohesity/profile/initSID.utl' to
'/usr/sap/SAL/cohesity/profile/initSAL.utl'
Creating symbolic link from '/usr/sap/SAL/cohesity/profile/initSAL.utl' to
'/oracle/SAL/sapprof/initSAL.utl'
-------------------------------------------------------------------------
--
Do you want to continue adding the user 'saladm' in /etc/sudoers file? Enter
y[es]/n[o]: y
Continuing installation...
-------------------------------------------------------------------------
--
==> Adding 'saladm' to sudoers list
Defaults:saladm !requiretty
Removing following line from sudoers
==> Whitelisting following command
```

```
/us-
r/bin/-
java,/us-
r/bin/-
moun-
t,/us-
r/bin/u-
moun-
t,/us-
r/bin/ch-
mod-
,/us-
r/bin/chow-
n,/us-
r/bin/mk-
dir,/usr/bin/rm,/usr/bin/cp,/usr/bin/tee,/usr/bin/stat,/usr/bin/hostname
-------------------------------------------------------------------------
--
==> Checking required permissions for user=saladm
==> User=saladm exists in /etc/passwd
==> User=saladm has all the required permissions to run Cohesity java back-
int secure connector service.
-------------------------------------------------------------------------
--
==> Finding group of service user: saladm
==> Group name from /etc/group : sapsys
==> Group of service user: sapsys
-------------------------------------------------------------------------
--
==> Cohesity java backint secure connector service for SAP Oracle is
installed successfully.
-------------------------------------------------------------------------
--
```

Install the BACKINT Secure Connector Service using the RPM-based Installer

On the SAP HANA node, perform the following to install the BACKINT Secure Connector Service using an RPM-based installer, `cohesity-oracle-backint-6.6-1.x86_64.rpm`. You can install the BACKINT Secure Connector Service using a `root` user account or a user account with `sudo` privileges.

To install the BACKINT secure connector service using a user account with `sudo` privileges:

1. Log in to the SAP node's bash CLI.
2. Navigate to the directory where you have copied the RPM-based installer.
3. Run the following command to install the BACKINT secure connector service:

```
sudo -E rpm -ivh  cohesity-oracle-backint-6.6-1.x86_64.rpm [--
prefix=<install_dir>]
```

Where, `prefix` is an optional parameter. It enables you to define a directory on the SAP node on which you want to install the BACKINT secure connector service. The default installation directory is **/opt**.

The following is a sample output of the successful installation, where `SERVICE_USER` is `saladm` and `SID` is `SAL`:

```
[oracle@saptest ~]$ sudo -E rpm -ivh cohesity-oracle-backint-6.6-1.x86_
64.rpm
Preparing... ################################# [100%]
Pre-install/Upgrade section
==> initialized values
==>
==> User does not exist in /etc/passwd
==> Checking if User exists using getent command.
==> User found with getent command
==> Found java executable in path
==>
java version "1.8.0_261"
Java(TM) SE Runtime Environment (build 1.8.0_261-b12)
Java HotSpot(TM) 64-Bit Server VM (build 25.261-b12, mixed mode)
==> Java Runtime Environment check .................................
[passed]
==> NFS client check ..............................................
[passed]
==> Environment variables check ...................................
[passed]
==> SAP profile files check .......................................
[passed]
==> Service user OS roles privileges check ........................
[passed]
==> Architecture Name : x86_64
==> Operating System Name : Oracle Linux Server 7.7
==> Service User : saladm
==> Package Install Directory : /opt/cohesity/oracle_backint_plugin
==> SAP Database SID : SAL
==> Cluster IPs :
10.2.198.213,10.2.198.217,10.2.198.219,10.2.198.221,10.2.198.223,10.2.
198.225,10.2.198.227,10.2.198.215
==> View Name : ra_view
==> Package installation is happening at /opt/cohesity/oracle_backint_
plugin
```

```
Updating / installing...
1:cohesity-oracle-backint-0.0.0-1 ################################
[100%]
==> Post-install/Upgrade section

------------------------------------------------------------------------
-------
==> Customizing the required files.
Creating symbolic link from '/opt/cohesity/oracle_backint_
plugin/bin/backint.sh' to '/usr/sap/SAL/SYS/exe/run/backint'
Creating symbolic link from '/opt/cohesity/oracle_backint_
plugin/lib/libsbt_6_and_7_linux-x86_64.so' to
'/oracle/SAL/121/lib/libobk.so'
Copying '/opt/cohesity/oracle_backint_plugin/profile/initSID.utl' to
'/opt/cohesity/oracle_backint_plugin/profile/initSAL.utl'
Creating symbolic link from '/opt/cohesity/oracle_backint_
plugin/profile/initSAL.utl' to '/oracle/SAL/sapprof/initSAL.utl'
Backup sap profile file from location /oracle/SAL/sapprof/initSAL.sap
------------------------------------------------------------------------
-------
------------------------------------------------------------------------
-------
==> Adding 'saladm' to sudoers list
==> Whitelisting following command
/usr/bin/java,/usr/bin/mount,/usr/bin/umount,/usr/bin/chmod,/usr/bin/c
hown,/usr/bin/mkdir,/usr/bin/rm,/usr
/bin/cp,/usr/bin/tee,/usr/bin/stat,/usr/bin/hostname
------------------------------------------------------------------------
-------
==> Checking required permissions for user saladm
==> User saladm has all the required permissions to run Cohesity java
backint secure connector service.
------------------------------------------------------------------------
-------

==> Finding group of service user: saladm
==> Group name from /etc/group : sapsys
==> Group of service user: sapsys
Executing backint -V
Cohesity DataProtect 6.5 for SAP Oracle Linux
backint.jar, Build Version: qa-ENG_172760-125cbd48
Successful Installation of Cohesity secure connector service
```

To install the BACKINT secure connector service using the `root` user account:

1. Log in to the SAP node's bash CLI as a root user.
2. Navigate to the directory where you have copied the RPM-based installer.
3. Run the following command to install the BACKINT secure connector service:

```
rpm -ivh cohesity-oracle-backint-6.6-1.x86_64.rpm [--prefix=<install_
dir>]
```

Where, `prefix` is an optional parameter. It enables you to define a directory on the SAP node on which you want to install the BACKINT secure connector service. The default installation directory is **/opt**.

The following is a sample output of the successful installation, where `SERVICE_USER` is `saladm` and `SID` is `SAL`:

```
[root@saptest ~]$ rpm -ivh cohesity-oracle-backint-6.6-1.x86_64.rpm
Preparing... ################################# [100%]
Pre-install/Upgrade section
==> initialized values
==>
==> User does not exist in /etc/passwd
==> Checking if User exists using getent command.
==> User found with getent command
==> Found java executable in path
==>
java version "1.8.0_261"
Java(TM) SE Runtime Environment (build 1.8.0_261-b12)
Java HotSpot(TM) 64-Bit Server VM (build 25.261-b12, mixed mode)
==> Java Runtime Environment check ..............................
[passed]
==> NFS client check ............................................
[passed]
==> Environment variables check .................................
[passed]
==> SAP profile files check .....................................
[passed]
==> Service user OS roles privileges check ......................
[passed]
==> Architecture Name : x86_64
==> Operating System Name : Oracle Linux Server 7.7
==> Service User : saladm
==> Package Install Directory : /opt/cohesity/oracle_backint_plugin
==> SAP Database SID : SAL
==> Cluster IPs :
10.2.198.213,10.2.198.217,10.2.198.219,10.2.198.221,10.2.198.223,10.2.
198.225,10.2.198.227,10.2.198.215
```

```
==> View Name : ra_view
==> Package installation is happening at /opt/cohesity/oracle_backint_
plugin
Updating / installing...
1:cohesity-oracle-backint-0.0.0-1 ################################
[100%]
==> Post-install/Upgrade section

--------------------------------------------------------------------------
-------
==> Customizing the required files.
Creating symbolic link from '/opt/cohesity/oracle_backint_
plugin/bin/backint.sh' to '/usr/sap/SAL/SYS/exe/run/backint'
Creating symbolic link from '/opt/cohesity/oracle_backint_
plugin/lib/libsbt_6_and_7_linux-x86_64.so' to
'/oracle/SAL/121/lib/libobk.so'
Copying '/opt/cohesity/oracle_backint_plugin/profile/initSID.utl' to
'/opt/cohesity/oracle_backint_plugin/profile/initSAL.utl'
Creating symbolic link from '/opt/cohesity/oracle_backint_
plugin/profile/initSAL.utl' to '/oracle/SAL/sapprof/initSAL.utl'
Backup sap profile file from location /oracle/SAL/sapprof/initSAL.sap
--------------------------------------------------------------------------
-------
--------------------------------------------------------------------------
-------
==> Adding 'saladm' to sudoers list
==> Whitelisting following command
/usr/bin/java,/usr/bin/mount,/usr/bin/umount,/usr/bin/chmod,/usr/bin/c
hown,/usr/bin/mkdir,/usr/bin/rm,/usr
/bin/cp,/usr/bin/tee,/usr/bin/stat,/usr/bin/hostname
--------------------------------------------------------------------------
-------
==> Checking required permissions for user saladm
==> User saladm has all the required permissions to run Cohesity java
backint secure connector service.
--------------------------------------------------------------------------
-------

==> Finding group of service user: saladm
==> Group name from /etc/group : sapsys
==> Group of service user: sapsys
Executing backint -V
Cohesity DataProtect 6.5 for SAP Oracle Linux
```

```
backint.jar, Build Version: qa-ENG_172760-125cbd48
Successful Installation of Cohesity secure connector service
```

Update the Cohesity Parameter File

As part of the BACKINT Secure Connector Service for Oracle installation, Cohesity provides a parameter file (`initSID.utl`) that enables the SAP on Oracle node to use the BACKINT Secure Connector Service for a backup job. The parameter file is available at: `<SAPDATA_HOME>/sapprof/initSID.utl`.

The following is a sample `initSID.utl` file:

```
##-------------------------------------------------------------------------
----
## List of IPs' of Cohesity cluster. Only VIPs can be used and not physical
IPs.
##-------------------------------------------------------------------------
----
cluster_ips : "${IP_ADDRESS}"

##-------------------------------------------------------------------------
----
## If service_type is 'REMOTESNAPFS' then port number on Cohesity cluster
nodes
## to which BACKINT executable connects.
## Port 11117 for secure connection (certificate config path should be a
valid
## path) otherwise port 11111 for insecure connection (certificate config
path
## should be an empty string).
## If service_type is 'NFS' then this field is not applicable.
##-------------------------------------------------------------------------
----
cluster_port : 11113

##-------------------------------------------------------------------------
----
## Name of the view on which backup data resides (Note: View Name should not
## contain the View Path or Cluster Name/IP).
##-------------------------------------------------------------------------
----
view_name : "${VIEW_NAME}"

##-------------------------------------------------------------------------
----
## Mount path for nfs approach.
##-------------------------------------------------------------------------
```

```
----
sap_nfs_mount_path : "/tmp/cohesity_nfs"

##------------------------------------------------------------------------
----
## Mount options for nfs approach.
##------------------------------------------------------------------------
----
# rsize, wsize: The  maximum  number of bytes in each network READ/WRITE
request that the NFS client
# can receive/send when reading/writing data from/to a file on an NFS
server.  The actual data
# payload size of each NFS READ/NFS WRITE request is equal to or smaller
than the rsize/wsize
# setting. The largest read/wsize payload supported by the Linux NFS client
is
# 1,048,576 bytes (one megabyte).
# The rsize/wsize mount option as specified on the mount(8) command line
appears
# in the /etc/mtab file. However, the effective rsize/wsize value negotiated
by
# the client and server is reported in the /proc/mounts file.
sap_nfs_mount_options : "rsize=1048576,wsize=1048576"

##------------------------------------------------------------------------
----
## Maximum number of retries during reads/writes.
##------------------------------------------------------------------------
----
sap_max_read_write_retries : 3

##------------------------------------------------------------------------
----
## Maximum size of each IO in bytes. The range is 1048576-4194304 bytes.
##------------------------------------------------------------------------
----
sap_max_io_size_in_bytes : 1048576

##------------------------------------------------------------------------
----
## Maximum number of IO cached/prefetched. The range is 1-128.
##------------------------------------------------------------------------
----
sap_max_data_queue_size : 4
```

```
##-------------------------------------------------------------------------
----
## Maximum number of inflight read-writes at a time.The range is 1-128.
##-------------------------------------------------------------------------
----
sap_max_num_inflight_io : 4

##-------------------------------------------------------------------------
----
## Maximum number of asynchronous callback threads for NFS read-write for
## SAP Hana and for local as well for SAP Oracle.
##-------------------------------------------------------------------------
----
sap_max_io_asynchronous_callback_concurrency : 2

##-------------------------------------------------------------------------
----
## Maximum number of pipes that will be backed up parallelly at any time by
## Cohesity Secure Connector Service during the backup.
##-------------------------------------------------------------------------
----
sap_max_backup_streams : 4

##-------------------------------------------------------------------------
----
## Maximum number of pipes that will be restored parallelly at any time by
## Cohesity Secure Connector Service during the restore.
##-------------------------------------------------------------------------
----
sap_max_restore_streams : 4

##-------------------------------------------------------------------------
----
## To tune, metadata read performance for inquiry and restore.
## The default value is 4.
##-------------------------------------------------------------------------
----
sap_max_metafiles_read_concurrency : 10

##-------------------------------------------------------------------------
----
## Number of days for which inquire information is returned.
##-------------------------------------------------------------------------
```

```
----
sap_max_inquire_period_days : 180

##-------------------------------------------------------------------------
----
## Location on SAP nodes where the Cohesity Secure Connector Service
## will create debug logs.
##-------------------------------------------------------------------------
----
backint_log_dir : "/tmp/cohesity_logs"

##-------------------------------------------------------------------------
----
## Java log level. It can possibly take these values:
## all/trace/debug/info/warn/error/fatal
##-------------------------------------------------------------------------
----
java_log_level : "info"

##-------------------------------------------------------------------------
----
## Service Type
##-------------------------------------------------------------------------
----
service_type : "REMOTESNAPFS"

##-------------------------------------------------------------------------
----
## Database Type
##-------------------------------------------------------------------------
----
database_type : "ORACLE"
```

Edit the parameter file and enter the values for the following parameters:

Parameter	Description
`cluster_ips`	List of IP addresses of the Cohesity Cluster. **Note:** Ensure that you provide only the virtual IP addresses.

Parameter	Description
view_name	The name of the Cohesity view used to store the backup data. **Note:** Ensure that the Cohesity view name does not contain the path of the Cohesity view or Cluster Name or IP address.
sap_nfs_mount_path	The NFS mount path.
sap_nfs_mount_options	The NFS mount options.
sap_max_read_write_retries	The Maximum number of retries during reads or writes.
sap_max_io_size_in_bytes	Maximum size of each IO in bytes. The range is 1048576-4194304 bytes.
sap_max_data_queue_size	Maximum number of IO cached or prefetched. The range is 1-128.
sap_max_num_inflight_io	Maximum number of inflight reads or writes at a given time. The range is 1-128.
sap_max_io_asynchronous_callback_concurrency	Maximum number of asynchronous callback threads for NFS read-write.
sap_max_backup_streams	The number of pipes that are backed up parallelly at any time by Cohesity BACKINT Secure Connector Service for Oracle during the backup process.
sap_max_restore_streams	The number of pipes that are backed up parallelly at any time by Cohesity BACKINT Secure Connector Service for Oracle during the restore process.
sap_max_metafiles_read_concurrency	Tune the metadata read performance for inquiry and restore. The default value is 4.
sap_max_inquire_period_days	The number of days for which the inquire information is returned.
backint_log_dir	Location on the SAP nodes where the BACKINT Secure Connector Service creates the debug logs.

Parameter	Description
java_log_level	The java log level. The supported values are: • all • trace • debug • info • warn • error • fatal
service_type	The type of service. By default, the value is set to REMOTESNAPFS. The supported values are REMOTESNAPFS or NFS.
database_type	The type of database. The supported value is ORACLE.

Uninstall the Cohesity BACKINT Secure Connector Service for Oracle

To uninstall the Cohesity BACKINT Secure Connector Service for Oracle, run the following command with appropriate parameters on SAP bash CLI:

```
cohesity_secure_connector_service_sap_oracle_installer_sap_oracle_java_
installer -- --uninstall -s <SID>
```

The following is a sample of the output:

```
[oracle@saptest ~]$ ./cohesity_secure_connector_service_6.5.1_sap_oracle_
java_installer -- --uninstall -s SAL
Verifying archive integrity... All good.
Uncompressing Cohesity Backint  100%

=~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~-
~~=
                  ____ ___  _   _ _____ ____ ___ _______   __

                 / ___/ _ \| | | | ____/ ___|_ _|_   _\ \ / /

                | |  | | | | |_| |  _| \___ \| |  | |  \ V /

                | |__| |_| |  _  | |___ ___) | |  | |   | |

                 \____\___/|_| |_|_____|____/___| |_|   |_|
```

```
=~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~-
~~=
                    ***** NOTE ***** NOTE ***** NOTE *****


      This script based installer is designed for SAP Oracle system like

                      SAP NetWeaver, SAP ECC, and SAP R2.


                    ***** NOTE ***** NOTE ***** NOTE *****

-------------------------------------------------------------------------------
--
                --- S Y S T E M     I N F O R M A T I O N ---

***** DATE AND TIME INFORMATION *****
Wed Jul 22 11:38:19 EDT 2020

***** HOSTNAME INFORMATION *****
   Static hostname: saptest
         Icon name: computer-vm
           Chassis: vm
        Machine ID: 466bc8302ebc4873a6c8541d0854e0fd
           Boot ID: 6e3f19bee42f499d87f65f29b64c91ac
    Virtualization: vmware
  Operating System: Oracle Linux Server 7.7
       CPE OS Name: cpe:/o:oracle:linux:7:7:server
            Kernel: Linux 4.1.12-124.34.1.el7uek.x86_64
      Architecture: x86-64

***** ORACLE DATABASE INFORMATION *****
Oracle Database version : 12.1.0.2.0

***** FILE SYSTEM DISK SPACE USAGE *****
Filesystem           Size  Used Avail Use% Mounted on
devtmpfs             7.7G     0  7.7G   0% /dev
tmpfs                7.8G     0  7.8G   0% /dev/shm
tmpfs                7.8G   45M  7.7G   1% /run
tmpfs                7.8G     0  7.8G   0% /sys/fs/cgroup
```

```
/dev/sda1              97G   46G   52G  47% /
/dev/mapper/ol-u01     70G   11G   60G  16% /u01
/dev/mapper/ol-home    51G   12G   40G  22% /home
tmpfs                 1.6G  4.0K  1.6G   1% /run/user/42
tmpfs                 1.6G   36K  1.6G   1% /run/user/54321
tmpfs                 1.6G     0  1.6G   0% /run/user/0

***** FREE AND USED MEMORY *****
              total        used        free      shared  buff/cache   avail-
able
Mem:       16150368     1450324     2066340     4533204    12633704
10075340
Swap:      17825788           0    17825788

***** SYSTEM UPTIME AND LOAD *****
 11:38:20 up 12:07,  6 users,  load average: 0.26, 0.15, 0.22

***** CURRENTLY LOGGED-IN USERS *****
oracle   :0           2020-07-21 23:31 (:0)
oracle   pts/0        2020-07-21 23:31 (:0)
root     pts/1        2020-07-21 23:58 (sv16-pm-dell6-c1.pm.cohesity.com)
root     pts/2        2020-07-21 23:59 (10.2.54.192)
oracle   pts/3        2020-07-22 10:53 (10.2.252.50)
oracle   pts/4        2020-07-22 11:24 (10.2.252.50)

***** TOP 5 MEMORY-CONSUMING PROCESSES *****
%MEM %CPU COMMAND
 1.7  0.0 ora_mman_sal
 1.6  0.2 gnome-shell
 1.0  0.0 ora_mmon_sal
 0.8  0.0 gnome-software
 0.5  0.0 ora_cjq0_sal
-------------------------------------------------------------------------------
--
==> Uninstalling Cohesity java backint secure connector service for SAP
Oracle
-------------------------------------------------------------------------------
--
Are you sure you want to uninstall? Enter y[es]/n[o]: y
Continuing uninstallation...
-------------------------------------------------------------------------------
--
Removed symbolic link for '/usr/sap/SAL/SYS/exe/run/backint'
Removed symbolic link for '/oracle/SAL/sapprof/initSAL.utl'
```

```
Removed install directory '/usr/sap/SAL/cohesity' and its contents.
-------------------------------------------------------------------------------
--
==> Cohesity java backint secure connector service for SAP Oracle is unin-
stalled successfully.
```

If you have installed Cohesity BACKINT Secure Connector Service for Oracle using the RPM-based installer, you can uninstall the service using the following command:

```
sudo -E rpm -evh cohesity-oracle-backint-6.6-1.x86_66
```

The following is a sample of the output:

```
[oracle@saptest ~]$ sudo -E rpm -evh cohesity-oracle-backint-0.0.0-1.x86_64
Preparing... ################################# [100%]
==> Pre-uninstall section
backint link path is /usr/sap//SYS/exe/run//backint
-------------------------------------------------------------------------------
--
==> Uninstalling Cohesity java backint secure connector service for SAP
Oracle
-------------------------------------------------------------------------------
--
==> Restoring sap profile file to location /sapprof/init.sap
Cleaning up / removing...
1:cohesity-oracle-backint-0.0.0-1 ################################# [100%]
==> Post-uninstall section
==> Deleting RPM install files at /opt/cohesity/oracle_backint_plugin
Oracle backint plugin is successfully uninstalled
```

Upgrade the Cohesity BACKINT Secure Connector Service for Oracle

To upgrade the SAP HANA secure connector service, run the following command on all SAP HANA nodes:

```
cohesity_secure_connector_service_sap_oracle_installer_sap_oracle_java_
installer -- --upgrade -s <SID>
```

The following is a sample of the output:

```
[oracle@saptest ~]$ ./cohesity_secure_connector_service_6.5.1_sap_oracle_
java_installer -- --upgrade -s SAL
Verifying archive integrity... All good.
Uncompressing Cohesity Backint  100%

=~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~-
~~=
                  ____ ___  _   _ _____ ____ ___ _______   __
```

```
                / ___/ _ \| | | | ____/ ___|_ _|_   _\ \ / /

               | |  | | | | |_| |  _| \___ \| |  | |  \ V /

               | |__| |_| |  _  | |___ ___) | |  | |   | |

                \____\___/|_| |_|_____|____/___| |_|   |_|



=~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~-
~~=
                    ***** NOTE ***** NOTE ***** NOTE *****


      This script based installer is designed for SAP Oracle system like

                      SAP NetWeaver, SAP ECC, and SAP R2.


                    ***** NOTE ***** NOTE ***** NOTE *****

--------------------------------------------------------------------------------
--
                ---  S Y S T E M     I N F O R M A T I O N  ---

***** DATE AND TIME INFORMATION *****
Wed Jul 22 11:37:32 EDT 2020

***** HOSTNAME INFORMATION *****
   Static hostname: saptest
         Icon name: computer-vm
           Chassis: vm
        Machine ID: 466bc8302ebc4873a6c8541d0854e0fd
           Boot ID: 6e3f19bee42f499d87f65f29b64c91ac
    Virtualization: vmware
  Operating System: Oracle Linux Server 7.7
       CPE OS Name: cpe:/o:oracle:linux:7:7:server
            Kernel: Linux 4.1.12-124.34.1.el7uek.x86_64
      Architecture: x86-64
```

```
***** ORACLE DATABASE INFORMATION *****
Oracle Database version : 12.1.0.2.0

***** FILE SYSTEM DISK SPACE USAGE *****
Filesystem           Size  Used Avail Use% Mounted on
devtmpfs             7.7G     0  7.7G   0% /dev
tmpfs                7.8G     0  7.8G   0% /dev/shm
tmpfs                7.8G   45M  7.7G   1% /run
tmpfs                7.8G     0  7.8G   0% /sys/fs/cgroup
/dev/sda1             97G   46G   52G  47% /
/dev/mapper/ol-u01    70G   11G   60G  16% /u01
/dev/mapper/ol-home   51G   12G   40G  22% /home
tmpfs                1.6G  4.0K  1.6G   1% /run/user/42
tmpfs                1.6G   36K  1.6G   1% /run/user/54321
tmpfs                1.6G     0  1.6G   0% /run/user/0

***** FREE AND USED MEMORY *****
              total        used        free      shared  buff/cache   avail-
able
Mem:       16150368     1429088     2105952     4533204    12615328
10078164
Swap:      17825788           0    17825788

***** SYSTEM UPTIME AND LOAD *****
 11:37:34 up 12:06,  6 users,  load average: 0.05, 0.08, 0.20

***** CURRENTLY LOGGED-IN USERS *****
oracle   :0           2020-07-21 23:31 (:0)
oracle   pts/0        2020-07-21 23:31 (:0)
root     pts/1        2020-07-21 23:58 (sv16-pm-dell6-c1.pm.cohesity.com)
root     pts/2        2020-07-21 23:59 (10.2.54.192)
oracle   pts/3        2020-07-22 10:53 (10.2.252.50)
oracle   pts/4        2020-07-22 11:24 (10.2.252.50)

***** TOP 5 MEMORY-CONSUMING PROCESSES *****
%MEM %CPU COMMAND
 1.7  0.0 ora_mman_sal
 1.6  0.2 gnome-shell
 1.0  0.0 ora_mmon_sal
 0.8  0.0 gnome-software
 0.5  0.0 ora_cjq0_sal
--------------------------------------------------------------------------
--
==> Upgrading Cohesity java backint secure connector service for SAP Oracle
```

```
------------------------------------------------------------------------------
--
==> Copying the required files.
Copying 'backint.sh' to install directory '/usr/sap/SAL/cohesity'
Copying 'backint.jar' to install directory '/usr/sap/SAL/cohesity'
Copying 'purge_job.sh' to scripts directory '/usr/sap/SAL/cohesity/scripts'
Copying 'initSID.utl' to profile directory '/usr/sap/SAL/cohesity/profile'
------------------------------------------------------------------------------
--
==> Customizing the required files.
Creating symbolic link from '/usr/sap/SAL/cohesity/backint.sh' to '/us-
r/sap/SAL/SYS/exe/run/backint'
Renaming '/usr/sap/SAL/cohesity/profile/initSID.utl' to
'/usr/sap/SAL/cohesity/profile/initSAL.utl'
Creating symbolic link from '/usr/sap/SAL/cohesity/profile/initSAL.utl' to
'/oracle/SAL/sapprof/initSAL.utl'
------------------------------------------------------------------------------
--
Do you want to continue adding the user 'saladm' in /etc/sudoers file? Enter
y[es]/n[o]: y
Continuing installation...
------------------------------------------------------------------------------
--
==> Adding 'saladm' to sudoers list
Defaults:saladm !requiretty
saladm ALL=(ALL) NOPASSWD:SETENV:/us-
r/bin/-
java,/us-
r/bin/-
moun-
t,/us-
r/bin/u-
moun-
t,/us-
r/bin/ch-
mod-
,/us-
r/bin/chow-
n,/us-
r/bin/mk-
dir,/usr/bin/rm,/usr/bin/cp,/usr/bin/tee,/usr/bin/stat,/usr/bin/hostname
Removing following line from sudoers
==> Whitelisting following command /us-
```

```
r/bin/-
java,/us-
r/bin/-
moun-
t,/us-
r/bin/u-
moun-
t,/us-
r/bin/ch-
mod-
,/us-
r/bin/chow-
n,/us-
r/bin/mk-
dir,/usr/bin/rm,/usr/bin/cp,/usr/bin/tee,/usr/bin/stat,/usr/bin/hostname
-------------------------------------------------------------------------
--
==> Checking required permissions for user=saladm
==> User=saladm exists in /etc/passwd
==> User=saladm has all the required permissions to run Cohesity java back-
int secure connector service.
-------------------------------------------------------------------------
--
==> Finding group of service user: saladm
==> Group name from /etc/group : sapsys
==> Group of service user: sapsys
-------------------------------------------------------------------------
--
==> Cohesity java backint secure connector service for SAP Oracle is
upgraded successfully.
-------------------------------------------------------------------------
--
```

Backup SAP on Oracle Databases

You can back up the Oracle database running on the SAP system using the SAP BR*TOOLS or using a Remote Adaptor Protection Group provided by Cohesity. See the following topics for details:

Backup Using SAP BR*TOOLS

SAP provides BR*TOOLS to back up Oracle databases. The Cohesity BACKINT Secure Connector Service leverages BR*TOOLs for protecting Oracle databases.

The SAP BR*TOOLS includes a BRBACKUP tool that enables you to back up data files, control files, and online redo log files of the database. When you back up the files on the SAP

system using the BRBACKUP tool, it calls the Cohesity BACKINT Secure Connector Service to transfer the backup files to the Cohesity cluster. For more information on SAP BR*TOOLS, refer to SAP on Oracle Documentation.

Sample

The following is a sample of the BRBACKUP command and the output:

```
[saladm@user@123 ~]$ brbackup -d util_file -t offline -m all -p initSAL.sap
-u system/fr9shst8rt -c
BR0051I BRBACKUP 7.40 (40)
BR0055I Start of database backup: bfeftkmh.aff 2020-07-14 08:38:07
BR0484I BRBACKUP log file: /oracle/SAL/sapbackup/bfeftkmh.aff
BR0477I Oracle pfile /oracle/SAL/sapprof/initSAL.ora created from spfile /or-
acle/SAL/121/dbs/spfileSAL.ora
BR1805I Oracle spfile /oracle/SAL/sapprof/spfileSAL.ora created from pfile
/oracle/SAL/sapprof/initSAL.ora

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:12
BR0057I Backup of database: SAL
BR0058I BRBACKUP action ID: bfeftkmh
BR0059I BRBACKUP function ID: aff
BR0110I Backup mode: ALL
BR0077I Database files to backup:
/oracle/SAL/origlogA/log_g11m1.dbf
/oracle/SAL/origlogB/log_g12m1.dbf
/oracle/SAL/origlogA/log_g13m1.dbf
/oracle/SAL/origlogB/log_g14m1.dbf
/oracle/SAL/origlogA/cntrl/cntrlSAL.dbf
BR0061I 20 files found to backup, total size 23939.885 MB
BR0143I Backup type: offline
BR0130I Backup device type: util_file
BR0109I Files will be saved by backup utility at file level
BR0126I Unattended mode active - no operator confirmation required

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:12
BR0307I Shutting down database instance SAL ...

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:30
BR0308I Shutdown of database instance SAL successful

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:30
BR0229I Calling backup utility with function 'backup'...

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:54
#FILE..... /oracle/SAL/sapdata2/sr3_3/sr3.data3
```

```
#SAVED.... 1594730312_cmhn

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:54
#FILE..... /oracle/SAL/sapdata2/sr3_4/sr3.data4
#SAVED.... 1594730312_6v8k

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:55
#FILE..... /oracle/SAL/sapdata2/sr3_1/sr3.data1
#SAVED.... 1594730312_8mts

BR0280I BRBACKUP time stamp: 2020-07-14 08:38:55
#FILE..... /oracle/SAL/sapdata2/sr3_2/sr3.data2
#SAVED.... 1594730312_ulk1

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:08
#FILE..... /oracle/SAL/sapdata2/sr3_6/sr3.data6
#SAVED.... 1594730313_d5cv

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:11
#FILE..... /oracle/SAL/sapdata2/sr3_5/sr3.data5
#SAVED.... 1594730312_b76t

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:11
#FILE..... /oracle/SAL/sapdata3/sr3750_1/sr3750.data1
#SAVED.... 1594730313_hva2

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:13
#FILE..... /oracle/SAL/sapdata3/sr3750_2/sr3750.data2
#SAVED.... 1594730313_oxwn

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:14
#FILE..... /oracle/SAL/sapdata4/sr3usr_1/sr3usr.data1
#SAVED.... 1594730313_gqii

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:18
#FILE..... /oracle/SAL/sapdata3/sr3750_5/sr3750.data5
#SAVED.... 1594730313_x0pe

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:22
#FILE..... /oracle/SAL/sapdata3/sr3750_3/sr3750.data3
#SAVED.... 1594730313_0zsy

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:25
#FILE..... /oracle/SAL/sapdata1/sysaux_1/sysaux.data1
```

```
#SAVED.... 1594730314_mna4

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:25
#FILE..... /oracle/SAL/sapdata1/undo_1/undo.data1
#SAVED.... 1594730313_4od5

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:26
#FILE..... /oracle/SAL/sapdata3/sr3750_4/sr3750.data4
#SAVED.... 1594730313_ze2w

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:27
#FILE..... /oracle/SAL/origlogA/log_g11m1.dbf
#SAVED.... 1594730314_ts4k

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:28
#FILE..... /oracle/SAL/origlogB/log_g12m1.dbf
#SAVED.... 1594730314_fql7

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:28
#FILE..... /oracle/SAL/sapdata1/system_1/system.data1
#SAVED.... 1594730314_9tld

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:28
#FILE..... /oracle/SAL/origlogA/cntrl/cntrlSAL.dbf
#SAVED.... 1594730314_6cyi

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:28
#FILE..... /oracle/SAL/origlogA/log_g13m1.dbf
#SAVED.... 1594730314_e54i

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:30
#FILE..... /oracle/SAL/origlogB/log_g14m1.dbf
#SAVED.... 1594730314_ju2w
BR0278I Command output of '/usr/sap/SAL/SYS/exe/run/brtools -f newgroup /us-
r/sap/SAL/SYS/exe/run/backint -u SAL -f backup -i /or-
acle/SAL/sapbackup/.bfeftkmh.lst -t file -p /oracle/SAL/sapprof/initSAL.utl
-c':
Backup operation finished successfully.

BR0232I 20 of 20 files saved by backup utility

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:31
BR0230I Backup utility called successfully
```

```
BR0280I BRBACKUP time stamp: 2020-07-14 08:39:32
BR0304I Starting and opening database instance SAL ...

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:41
BR0305I Start and open of database instance SAL successful

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:41
BR0229I Calling backup utility with function 'backup'...

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
#PFLOG.... /oracle/SAL/sapprof/initSAL.ora
#SAVED.... 1594730383_q3js

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
#PFLOG.... /oracle/SAL/sapprof/spfileSAL.ora
#SAVED.... 1594730383_60lj

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
#PFLOG.... /oracle/SAL/sapprof/initSAL.sap
#SAVED.... 1594730383_6q25

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
#PFLOG.... /oracle/SAL/sapprof/initSAL.utl
#SAVED.... 1594730383_3h93

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
#PFLOG.... /oracle/SAL/sapreorg/spaceSAL.log
#SAVED.... 1594730383_pxpc

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
#PFLOG.... /oracle/SAL/sapbackup/bfeftkmh.aff
#SAVED.... 1594730384_nqdd

BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
#PFLOG.... /oracle/SAL/sapbackup/backSAL.log
#SAVED.... 1594730384_73h9
BR0278I Command output of '/usr/sap/SAL/SYS/exe/run/brtools -f newgroup /us-
r/sap/SAL/SYS/exe/run/backint -u SAL -f backup -i /or-
acle/SAL/sapbackup/.bfeftkmh.lst -t file -p /oracle/SAL/sapprof/initSAL.utl
-c':
Backup operation finished successfully.

BR0232I 7 of 7 files saved by backup utility
```

```
BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
BR0230I Backup utility called successfully

BR0056I End of database backup: bfeftkmh.aff 2020-07-14 08:39:44
BR0280I BRBACKUP time stamp: 2020-07-14 08:39:44
BR0052I BRBACKUP completed successfully
```

Backup Using BACKINT Secure Connector Service

SAP provides BR*TOOLS to back up Oracle databases. The Cohesity BACKINT Secure Connector Service leverages BR*TOOLs for protecting Oracle databases. When the BACKINT Secure Connector Service for Oracle is installed and configured on the SAP system, it manages the communication between the SAP system and the Cohesity cluster.

The SAP BR*TOOLS includes a BRBACKUP tool that enables you to back up data files, control files, and online redo log files of the database. When you back up the files on the SAP system using the BRBACKUP tool, it calls the Cohesity BACKINT Secure Connector Service to transfer the backup files to the Cohesity cluster. For more information on SAP BR*TOOLS, refer to SAP on Oracle Documentation.

Considerations

Review and understand the following considerations before you protect your SAP on Oracle databases using Cohesity BACKINT Secure Connector Service:

- Incremental backups of the Oracle databases are not supported. For incremental backups, use the custom Oracle SBT along with the BACKINT Secure Connector Service. For more information, see Backup Using Oracle SBT and BACKINT Secure Connector Service.
- Source-side deduplication is not supported if you use only BACKINT Secure Connector to backup SAP on Oracle databases.

Backup Using a Remote Adaptor Protection Group

As part of the BACKINT Secure Connector Service for Oracle installation, Cohesity provides a backup script for the backup of SAP on the Oracle database. The backup script leverages the SAP BR*TOOLS capabilities provided by SAP. You can create a Remote Adaptor Protection Group and define the script parameters that enable you to back up the SAP on the Oracle database.

Backup Script Parameters

In the Remote Adaptor Protection Group, you can define the following script argument for your SAP on Oracle backup configuration:

Backup Script Parameters	Descriptions
`-m, --mode`	The mode of backup. You can specify the following: • **full** for full backup. • **incr** for incremental backup • **arch** to archive offline redo log files
`-t, --type`	The type of backup. You can specify offline or online backups based on your requirements. **Note:** In case of offline backup, the database is shut down before performing the backup.
`-D, --device`	The device type for performing the backup. The following are the device types, you can specify: • **util_file**: To backup all the files using the BACKINT Secure Connector Service for Oracle. • **rman_util**: To backup the data files using the custom SBT extension and other files using the BACKINT Secure Connector Service for Oracle.
`-l, --log-file [file]`	The filename of the log file where you want to log the backup script logs. By default the logs are stored as follows: `/tmp/cohesity_backup_script_logs/cohesity_oracle_backup_script_<time_in_seconds>_<pid>.log`
`-c, --channels`	Specify the number of channels configured for backups using RMAN. By default, the number of channels configured in the `rman_channels` parameter in the **init<SID>.sap** file is used.
`-r, --retention-days`	The retention period in days. The backups older than the specified retention period are deleted from both the catalog and the Cohesity View.
`-d, --debug`	Allows you to enable shell script debugging with set -x. By default, shell script debugging is disabled.

Create a Remote Adapter Protection Group for SAP on Oracle

To create a Remote Adapter Protection Group:

1. Navigate to **Data Protection** > **Protection** and do one of the following:
    a. To create a Protection Group—Click the plus icon on the top-right of the page and then select the type of Protection Group.
    b. To edit a Protection Group—Click the action icon next to the Protection Group and select **Edit**.
2. Here the *Job Type* is **Remote Adapter**. **Name**: Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes and periods. This field is required.
3. **Linux Hostname or IP**: Specify the Hostname or the IP Address of the Linux system where the scripts are located and will be run.
4. **Username**: Specify the username on the Linux system that will be running the scripts.
5. **Cluster SSH Public Key**: Set up permissions on the Linux system to allow the Cohesity cluster to run the scripts remotely on the system.
    a. Click **Copy Key to Clipboard** to copy the SSH public key into the clipboard.
    b. Log in to the Linux system with the username specified earlier.
    c. In the home directory of the user, determine if there is a .ssh directory that contains a file called authorized_keys:

    ```
    ls -lr ~/.ssh
    ```

    d. Add the SSH key to the `authorized_keys` file located in the **~/.ssh** directory of the specified User:
        1. If the directory and file do not exist, create the .ssh directory in the home directory, change directories into the .ssh directory, create the authorized_keys file and paste the SSH key into a new line of the file.
        2. If the **.ssh** directory exists but the file does not, change directories into the .ssh directory, create the authorized_keys file and paste the SSH key into a new line of the file.
        3. If the authorized_keys file exists in the .ssh directory, edit the authorized_keys file and paste the SSH key into a new line of the file.
    e. Set the **.ssh** directory and authorized_keys file permissions:

    ```
    chmod 700 .ssh
    chmod 640 .ssh/authorized_keys
    ```

6. **Policy**: Select an existing protection policy or create a new one by selecting **New Policy**.
7. **Script (with Full Path)**: Enter the full path to the script, for example:

```
 rhel-sap-oracle-003:devadm> /usr/sap/SAL/cohesity/scripts/backup_
job.sh
```

8. **Parameters**: Enter the parameters for the backup script based on your requirements. For more information, see Backup Script Parameters.

   For example:

Use Cases	Sample Script Arguments
Create a full online backup using the BACKINT Secure Connector Service.	`-m full -t online -D util_file`
Create a full online backup using the BACKINT Secure Connector Service with 30 days retention period.	`-m full -t online -D util_file -r 30`
Archive offline redo log files using BACKINT Secure Connector Service.	`-m arch -D util_file -r 30`

Backup Using Oracle SBT and BACKINT Secure Connector Service

Oracle database provides a custom third-party backup extension called SBT (Secure Backup to Tapes) for Oracle database backups. Cohesity provides a custom SBT extension that provides full backup and restore capabilities for Oracle databases.

SAP BR*TOOL supports extensions to back up the Oracle databases using the custom SBT extension libraries. You can use the custom SBT extension along with the BACKINT Secure Connector Service to perform:

1. Incremental Oracle database backups.
2. Multistream the multi-node backups using Oracle RMAN channels, thus enabling you to split the data between multiple nodes for optimal performance.
3. Source side deduplication.

> **Note:** Backing up SAP on Oracle databases using custom Oracle SBT and BACKINT Secure Connector Service is supported only on SAP on Oracle databases running on Linux (x86-64) system.

Backup Process

The Oracle database is backed up using the custom SBT extension in the following phases:

1. The BRBACKUP or BRARCHIVE commands call RMAN and start the Oracle server process. The Oracle server process reads the Oracle database files to be backed up

from the Oracle database and transfers it to the external backup library. The external backup library acts as a backup client, and it passes the data to the backup server. The server backs up the data to the Cohesity cluster.

2. The BRBACKUP uses the BACKINT Secure Connector Service to transfer the control file, the initialization profile, and the log files to the backup client. The backup client then transfers the data to the backup server, and the backup server backs up the data to the Cohesity cluster.

> **Note:** The control file contains the information related to the files' backup from the first phase.

For more information, see the values `rman_disk` or `rman_stage` for the backup_dev_type parameter.

Backup Using a Remote Adaptor Protection Group

As part of the BACKINT Secure Connector Service for Oracle installation, Cohesity provides a backup script for the backup of SAP on the Oracle database. The backup script leverages the SAP BR*TOOLS capabilities provided by SAP. You can create a Remote Adaptor Protection Group and define the script parameters that enable you to back up the SAP on the Oracle database.

Backup Script Parameters

In the Remote Adaptor Protection Group, you can define the following script argument for your SAP on Oracle backup configuration:

Backup Script Parameters	Descriptions
`-m, --mode`	The mode of backup. You can specify the following: • **full** for full backup. • **incr** for incremental backup • **arch** to archive offline redo log files
`-t, --type`	The type of backup. You can specify offline or online backups based on your requirements. **Note:** In case of offline backup, the database is shut down before performing the backup.

Backup Script Parameters	Descriptions
`-D, --device`	The device type for performing the backup. The following are the device types, you can specify: • **util_file**: To backup all the files using the BACKINT Secure Connector Service for Oracle. • **rman_util**: To backup the data files using the custom SBT extension and other files using the BACKINT Secure Connector Service for Oracle.
`-l, --log-file [file]`	The filename of the log file where you want to log the backup script logs. By default the logs are stored as follows: `/tmp/cohesity_backup_script_logs/cohesity_oracle_backup_script_<time_in_seconds>_<pid>.log`
`-c, --channels`	Specify the number of channels configured for backups using RMAN. By default, the number of channels configured in the `rman_channels` parameter in the **init<SID>.sap** file is used.
`-r, --retention-days`	The retention period in days. The backups older than the specified retention period are deleted from both the catalog and the Cohesity View.
`-d, --debug`	Allows you to enable shell script debugging with set -x. By default, shell script debugging is disabled.

Create a Remote Adapter Protection Group for SAP on Oracle

To create a Remote Adapter Protection Group:

1. Navigate to **Data Protection** > **Protection** and do one of the following:
   a. To create a Protection Group—Click the plus icon on the top-right of the page and then select the type of Protection Group.
   b. To edit a Protection Group—Click the action icon next to the Protection Group and select **Edit**.
2. Here the *Job Type* is **Remote Adapter**. **Name**: Enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes and periods. This field is required.
3. **Linux Hostname or IP**: Specify the Hostname or the IP Address of the Linux system where the scripts are located and will be run.
4. **Username**: Specify the username on the Linux system that will be running the scripts.

5. **Cluster SSH Public Key**: Set up permissions on the Linux system to allow the Cohesity cluster to run the scripts remotely on the system.
   a. Click **Copy Key to Clipboard** to copy the SSH public key into the clipboard.
   b. Log in to the Linux system with the username specified earlier.
   c. In the home directory of the user, determine if there is a .ssh directory that contains a file called authorized_keys:

      ```
      ls -lr ~/.ssh
      ```

   d. Add the SSH key to the `authorized_keys` file located in the **~/.ssh** directory of the specified User:
      1. If the directory and file do not exist, create the .ssh directory in the home directory, change directories into the .ssh directory, create the authorized_keys file and paste the SSH key into a new line of the file.
      2. If the **.ssh** directory exists but the file does not, change directories into the .ssh directory, create the authorized_keys file and paste the SSH key into a new line of the file.
      3. If the authorized_keys file exists in the .ssh directory, edit the authorized_keys file and paste the SSH key into a new line of the file.
   e. Set the **.ssh** directory and authorized_keys file permissions:

      ```
      chmod 700 .ssh
      chmod 640 .ssh/authorized_keys
      ```

6. **Policy**: Select an existing protection policy or create a new one by selecting **New Policy**.
7. **Script (with Full Path)**: Enter the full path to the script, for example:

   ```
    rhel-sap-oracle-003:devadm> /usr/sap/SAL/cohesity/scripts/backup_
   job.sh
   ```

8. **Parameters**: Enter the parameters for the backup script based on your requirements. For more information, see Backup Script Parameters.

   For example:

Use Cases	Sample Script Arguments
Create a full online backup using the custom SBT extension and BACKINT Secure Connector Service.	`-m full -t online -D rman_util`

Use Cases	Sample Script Arguments
Create a full online backup using the custom SBT extension and BACKINT Secure Connector Service with 30 days retention period and the number of channels set to 8.	`-m full -t online -D rman_util -c 8 -r 30`
Create an online incremental backup using the custom SBT extension and BACKINT Secure Connector Service.	`-m incr -t online -D rman_util -c 8 -r 30`
Archive offline redo log files using the custom SBT extension and BACKINT Secure Connector Service.	`-m arch -D rman_util -r 30`

Archive Offline Redo Log Files

The SAP BR*TOOLS includes a BRARCHIVE tool that enables you to archive offline redo log files. When you archive the offline redo log files on the SAP system using the BRARCHIVE tool, it calls the Cohesity BACKINT Secure Connector Service to transfer the archived files to the Cohesity cluster. For more information on SAP BR*TOOL, refer to SAP on Oracle Documentation.

Sample

The following is a sample of the BRARCHIVE command and the output:

```
[saladm@user@123 ~]$ brarchive -d util_file -cds -p initSAL.sap -u sys-
tem/fr8shst8rt -c
BR0002I BRARCHIVE 7.40 (40)
BR0006I Start of offline redolog processing: afeftlky.cds 2020-07-14
08:48:48
BR0484I BRARCHIVE log file: /oracle/SAL/saparch/afeftlky.cds
BR0477I Oracle pfile /oracle/SAL/sapprof/initSAL.ora created from spfile /or-
acle/SAL/121/dbs/spfileSAL.ora
BR1805I Oracle spfile /oracle/SAL/sapprof/spfileSAL.ora created from pfile
/oracle/SAL/sapprof/initSAL.ora
BR0030I BRARCHIVE summary log '/oracle/SAL/saparch/archSAL.log' not found

BR0280I BRARCHIVE time stamp: 2020-07-14 08:48:49
BR0008I Offline redolog processing for database instance: SAL
BR0009I BRARCHIVE action ID: afeftlky
BR0010I BRARCHIVE function ID: cds
BR0048I Archive function: copy_delete_save
BR0011I 36 offline redolog files found to process, total size 5993.521 MB
BR0130I Backup device type: util_file
BR0109I Files will be saved by backup utility at file level
```

```
BR0126I Unattended mode active - no operator confirmation required

BR0280I BRARCHIVE time stamp: 2020-07-14 08:48:49
BR0229I Calling backup utility with function 'backup'...

BR0280I BRARCHIVE time stamp: 2020-07-14 08:48:57
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_30_1042718659.dbf
#SAVED.... 1594730931_dydc

BR0280I BRARCHIVE time stamp: 2020-07-14 08:48:57
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_31_1042718659.dbf
#SAVED.... 1594730931_7ijb

BR0280I BRARCHIVE time stamp: 2020-07-14 08:48:57
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_33_1042718659.dbf
#SAVED.... 1594730932_zxbm

BR0280I BRARCHIVE time stamp: 2020-07-14 08:48:58
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_32_1042718659.dbf
#SAVED.... 1594730932_m4oe

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:00
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_34_1042718659.dbf
#SAVED.... 1594730932_yi6o

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:00
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_36_1042718659.dbf
#SAVED.... 1594730932_l5au

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:00
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_37_1042718659.dbf
#SAVED.... 1594730932_926c

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:00
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_35_1042718659.dbf
#SAVED.... 1594730932_56um

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:02
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_38_1042718659.dbf
#SAVED.... 1594730932_kufx

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:02
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_40_1042718659.dbf
#SAVED.... 1594730933_y89y
```

```
BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:02
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_41_1042718659.dbf
#SAVED.... 1594730933_3rum

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:02
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_39_1042718659.dbf
#SAVED.... 1594730932_333s

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:04
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_42_1042718659.dbf
#SAVED.... 1594730933_lxrv

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:04
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_45_1042718659.dbf
#SAVED.... 1594730933_06jk

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:05
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_44_1042718659.dbf
#SAVED.... 1594730933_p7uf

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:05
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_43_1042718659.dbf
#SAVED.... 1594730933_nrb9

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:06
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_46_1042718659.dbf
#SAVED.... 1594730933_9t7n

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:06
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_48_1042718659.dbf
#SAVED.... 1594730933_k9jn

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:06
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_49_1042718659.dbf
#SAVED.... 1594730934_7csc

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:07
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_47_1042718659.dbf
#SAVED.... 1594730933_nac7

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:07
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_53_1042718659.dbf
#SAVED.... 1594730934_4wsv
```

```
BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:07
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_52_1042718659.dbf
#SAVED.... 1594730934_89kw

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:08
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_51_1042718659.dbf
#SAVED.... 1594730934_8gbf

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:08
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_50_1042718659.dbf
#SAVED.... 1594730934_yfaa

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:10
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_55_1042718659.dbf
#SAVED.... 1594730934_78it

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:10
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_54_1042718659.dbf
#SAVED.... 1594730934_1ru5

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:11
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_57_1042718659.dbf
#SAVED.... 1594730934_1wml

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:12
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_56_1042718659.dbf
#SAVED.... 1594730934_6l7p

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:13
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_58_1042718659.dbf
#SAVED.... 1594730935_yv36

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:14
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_59_1042718659.dbf
#SAVED.... 1594730935_flne

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:15
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_60_1042718659.dbf
#SAVED.... 1594730935_ihbd

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:16
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_61_1042718659.dbf
#SAVED.... 1594730935_x5gf
```

```
BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:16
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_63_1042718659.dbf
#SAVED.... 1594730935_ftc5

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:17
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_62_1042718659.dbf
#SAVED.... 1594730935_9acq

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:17
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_64_1042718659.dbf
#SAVED.... 1594730935_me46

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:19
#ARCHIVE.. /oracle/SAL/oraarch/SALarch1_65_1042718659.dbf
#SAVED.... 1594730935_z1pa
BR0278I Command output of '/usr/sap/SAL/SYS/exe/run/brtools -f newgroup /us-
r/sap/SAL/SYS/exe/run/backint -u SAL -f backup -i /or-
acle/SAL/saparch/.afeftlky.lst -t file -p /oracle/SAL/sapprof/initSAL.utl -
c':
Backup operation finished successfully.

BR0232I 36 of 36 files saved by backup utility

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:20
BR0230I Backup utility called successfully

BR0016I 36 offline redolog files processed, total size 5993.521 MB

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:20
BR0229I Calling backup utility with function 'backup'...

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
#PFLOG.... /oracle/SAL/sapprof/spfileSAL.ora
#SAVED.... 1594730962_c539

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
#PFLOG.... /oracle/SAL/sapprof/initSAL.ora
#SAVED.... 1594730962_l7c1

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
#PFLOG.... /oracle/SAL/sapreorg/spaceSAL.log
#SAVED.... 1594730962_ucyr
```

```
BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
#PFLOG.... /oracle/SAL/sapprof/initSAL.sap
#SAVED.... 1594730962_53cy

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
#PFLOG.... /oracle/SAL/sapprof/initSAL.utl
#SAVED.... 1594730962_upql

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
#PFLOG.... /oracle/SAL/saparch/afeftlky.cds
#SAVED.... 1594730962_zss8

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
#PFLOG.... /oracle/SAL/saparch/archSAL.log
#SAVED.... 1594730962_8bya
BR0278I Command output of '/usr/sap/SAL/SYS/exe/run/brtools -f newgroup /us-
r/sap/SAL/SYS/exe/run/backint -u SAL -f backup -i /or-
acle/SAL/saparch/.afeftlky.lst -t file -p /oracle/SAL/sapprof/initSAL.utl -
c':
Backup operation finished successfully.

BR0232I 7 of 7 files saved by backup utility

BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
BR0230I Backup utility called successfully

BR0007I End of offline redolog processing: afeftlky.cds 2020-07-14 08:49:23
BR0280I BRARCHIVE time stamp: 2020-07-14 08:49:23
BR0003I BRARCHIVE completed successfully
```

Restore SAP on Oracle Databases

The SAP BR*TOOLS includes a BRRESTORE tool that enables you to restore an entire database backup or parts of it when the backup was performed using the BRBACKUP tool. When you restore the file to the SAP system using the BRRESTORE tool, it calls the Cohesity BACKINT Secure Connector Service to restore the files from the Cohesity cluster. For more information on SAP BR*TOOL, refer to SAP on Oracle Documentation.

Sample

The following is a sample of the BRRESTORE command and the output:

```
[saladm@saptest ~]$ brrestore -d util_file -b last -m full -p initSAL.sap -u
system/fr8shst8rt
BR0401I BRRESTORE 7.40 (40)
```

```
BR0405I Start of file restore: rfefulbb.rsb 2020-07-14 13:37:27
BR0484I BRRESTORE log file: /oracle/SAL/sapbackup/rfefulbb.rsb

BR0428W File /oracle/SAL/sapdata2/sr3_1/sr3.data1 will be overwritten
BR0428W File /oracle/SAL/sapdata2/sr3_2/sr3.data2 will be overwritten
BR0428W File /oracle/SAL/sapdata2/sr3_3/sr3.data3 will be overwritten
BR0428W File /oracle/SAL/sapdata2/sr3_4/sr3.data4 will be overwritten
BR0428W File /oracle/SAL/sapdata2/sr3_5/sr3.data5 will be overwritten
BR0428W File /oracle/SAL/sapdata2/sr3_6/sr3.data6 will be overwritten
BR0428W File /oracle/SAL/sapdata3/sr3750_1/sr3750.data1 will be overwritten
BR0428W File /oracle/SAL/sapdata3/sr3750_2/sr3750.data2 will be overwritten
BR0428W File /oracle/SAL/sapdata3/sr3750_3/sr3750.data3 will be overwritten
BR0428W File /oracle/SAL/sapdata3/sr3750_4/sr3750.data4 will be overwritten
BR0428W File /oracle/SAL/sapdata3/sr3750_5/sr3750.data5 will be overwritten
BR0428W File /oracle/SAL/sapdata4/sr3usr_1/sr3usr.data1 will be overwritten
BR0428W File /oracle/SAL/sapdata1/undo_1/undo.data1 will be overwritten
BR0428W File /oracle/SAL/sapdata1/sysaux_1/sysaux.data1 will be overwritten
BR0428W File /oracle/SAL/sapdata1/system_1/system.data1 will be overwritten
BR0428W File /oracle/SAL/origlogA/log_g11m1.dbf will be overwritten
BR0428W File /oracle/SAL/mirrlogA/log_g11m2.dbf will be overwritten
BR0428W File /oracle/SAL/origlogB/log_g12m1.dbf will be overwritten
BR0428W File /oracle/SAL/mirrlogB/log_g12m2.dbf will be overwritten
BR0428W File /oracle/SAL/origlogA/log_g13m1.dbf will be overwritten
BR0428W File /oracle/SAL/mirrlogA/log_g13m2.dbf will be overwritten
BR0428W File /oracle/SAL/origlogB/log_g14m1.dbf will be overwritten
BR0428W File /oracle/SAL/mirrlogB/log_g14m2.dbf will be overwritten
BR0428W File /oracle/SAL/origlogA/cntrl/cntrlSAL.dbf will be overwritten
BR0428W File /oracle/SAL/origlogB/cntrl/cntrlSAL.dbf will be overwritten
BR0428W File /oracle/SAL/sapdata1/cntrl/cntrlSAL.dbf will be overwritten

BR0280I BRRESTORE time stamp: 2020-07-14 13:37:27
BR0256I Enter 'c[ont]' to continue, 's[top]' to cancel BRRESTORE:
BR0280I BRRESTORE time stamp: 2020-07-14 13:37:29
BR0257I Your reply: 'c'
BR0259I Program execution will be continued...

BR0280I BRRESTORE time stamp: 2020-07-14 13:37:29
BR0407I Restore of database: SAL
BR0408I BRRESTORE action ID: rfefulbb
BR0409I BRRESTORE function ID: rsb
BR0449I Restore mode: FULL
BR0411I Database files to restore:
/oracle/SAL/origlogA/log_g11m1.dbf
/oracle/SAL/mirrlogA/log_g11m2.dbf
```

```
/oracle/SAL/origlogB/log_g12m1.dbf
/oracle/SAL/mirrlogB/log_g12m2.dbf
/oracle/SAL/origlogA/log_g13m1.dbf
/oracle/SAL/mirrlogA/log_g13m2.dbf
/oracle/SAL/origlogB/log_g14m1.dbf
/oracle/SAL/mirrlogB/log_g14m2.dbf
/oracle/SAL/origlogA/cntrl/cntrlSAL.dbf
/oracle/SAL/origlogB/cntrl/cntrlSAL.dbf
/oracle/SAL/sapdata1/cntrl/cntrlSAL.dbf
BR0419I Files will be restored from backup: bfeftkmh.aff 2020-07-14 08:38:07
BR0416I 20 files found to restore, total size 23939.885 MB
BR0421I Restore device type: util_file

BR0280I BRRESTORE time stamp: 2020-07-14 13:37:29
BR0256I Enter 'c[ont]' to continue, 's[top]' to cancel BRRESTORE:
BR0280I BRRESTORE time stamp: 2020-07-14 13:37:30
BR0257I Your reply: 'c'
BR0259I Program execution will be continued...

BR0280I BRRESTORE time stamp: 2020-07-14 13:37:31
BR0229I Calling backup utility with function 'restore'...
BR0278I Command output of '/usr/sap/SAL/SYS/exe/run/backint -u SAL -f
restore -i /oracle/SAL/sapbackup/.rfefulbb.lst -t file -p /or-
acle/SAL/sapprof/initSAL.utl':
Restore operation finished successfully.

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata2/sr3_1/sr3.data1
#RESTORED. 1594730312_8mts

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata2/sr3_2/sr3.data2
#RESTORED. 1594730312_ulk1

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata2/sr3_3/sr3.data3
#RESTORED. 1594730312_cmhn

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata2/sr3_4/sr3.data4
#RESTORED. 1594730312_6v8k

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata2/sr3_5/sr3.data5
```

```
#RESTORED. 1594730312_b76t

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata2/sr3_6/sr3.data6
#RESTORED. 1594730313_d5cv

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata3/sr3750_1/sr3750.data1
#RESTORED. 1594730313_hva2

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata3/sr3750_2/sr3750.data2
#RESTORED. 1594730313_oxwn

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata3/sr3750_3/sr3750.data3
#RESTORED. 1594730313_0zsy

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata3/sr3750_4/sr3750.data4
#RESTORED. 1594730313_ze2w

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata3/sr3750_5/sr3750.data5
#RESTORED. 1594730313_x0pe

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata4/sr3usr_1/sr3usr.data1
#RESTORED. 1594730313_gqii

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata1/undo_1/undo.data1
#RESTORED. 1594730313_4od5

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata1/sysaux_1/sysaux.data1
#RESTORED. 1594730314_mna4

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/sapdata1/system_1/system.data1
#RESTORED. 1594730314_9tld

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/origlogA/log_g11m1.dbf
```

```
#RESTORED. 1594730314_ts4k

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/origlogB/log_g12m1.dbf
#RESTORED. 1594730314_fql7

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/origlogA/log_g13m1.dbf
#RESTORED. 1594730314_e54i

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/origlogB/log_g14m1.dbf
#RESTORED. 1594730314_ju2w

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
#FILE..... /oracle/SAL/origlogA/cntrl/cntrlSAL.dbf
#RESTORED. 1594730314_6cyi

BR0374I 20 of 20 files restored by backup utility

BR0280I BRRESTORE time stamp: 2020-07-14 13:38:15
BR0230I Backup utility called successfully

BR0351I Restoring /oracle/SAL/mirrlogA/log_g11m2.dbf
BR0355I from /oracle/SAL/origlogA/log_g11m1.dbf ...

BR0351I Restoring /oracle/SAL/mirrlogB/log_g12m2.dbf
BR0355I from /oracle/SAL/origlogB/log_g12m1.dbf ...

BR0351I Restoring /oracle/SAL/mirrlogA/log_g13m2.dbf
BR0355I from /oracle/SAL/origlogA/log_g13m1.dbf ...

BR0351I Restoring /oracle/SAL/mirrlogB/log_g14m2.dbf
BR0355I from /oracle/SAL/origlogB/log_g14m1.dbf ...

BR0351I Restoring /oracle/SAL/origlogB/cntrl/cntrlSAL.dbf
BR0355I from /oracle/SAL/origlogA/cntrl/cntrlSAL.dbf ...

BR0351I Restoring /oracle/SAL/sapdata1/cntrl/cntrlSAL.dbf
BR0355I from /oracle/SAL/origlogA/cntrl/cntrlSAL.dbf ...

BR0406I End of file restore: rfefulbb.rsb 2020-07-14 13:38:16
BR0280I BRRESTORE time stamp: 2020-07-14 13:38:16
BR0403I BRRESTORE completed successfully with warnings
```

Chapter 21: Kubernetes

Kubernetes is an open-source software platform that automates the deployment, management, and scaling of containerized applications. Using Cohesity Data Cloud, you can backup (file-based) and recover the Kubernetes cluster resources, PersistentVolumesClaims (PVCs), and metadata of the Kubernetes namespaces. Cohesity uses the Cohesity's Data Mover pod that will be deployed in the Kubernetes cluster for backing up the PVCs of the Kubernetes namespaces. Cohesity uses the third-party software, Velero, to back up the cluster and namespace scoped Kubernetes resources.

Plan and Prepare for Protecting Kubernetes Clusters

Before you register your Kubernetes cluster with the Cohesity cluster and protect the namespaces of the Kubernetes container, ensure that:

- Cohesity cluster supports the Kubernetes cluster version you plan to protect. For more information, see Supported Versions.
- The firewall ports required for the communication between the Cohesity cluster and Kubernetes cluster are open. For more information, see Network Firewall Considerations.
- You have the minimum permission required to register and manage your Kubernetes cluster. For more information, see Minimum Permission.
- You meet the prerequisites to register and manage your Kubernetes cluster. For more information, see Prerequisites.
- You have read the considerations for the backup and recovery of the Kubernetes cluster. For more information, see Considerations.

Supported Versions

Cohesity cluster supports data protection and recovery of the Kubernetes clusters. For supported versions of the Kubernetes clusters, see Supported Versions.

Network Firewall Considerations

If there is a firewall between the Cohesity cluster and the Kubernetes cluster, the following ports need to be opened.

Cohesity cluster	Type:Port	Direction	Kubernetes cluster	Type:Port	Usage Notes
Cohesity cluster nodes	TCP:	Outgoing	Kubernetes cluster api-address	TCP: <Kubernetes API server port>	Kubernetes cluster and Velero API Access. For more details on Velero support, see the Considerations section.
Cohesity cluster nodes	TCP:	Outgoing	Kubernetes cluster worker nodes	TCP:30000-32767[1] TCP:32000-(32000+#<worker nodes>)[2]	Cohesity cluster to Cohesity Datamover communication
Cohesity cluster VIPs	TCP:3000	Incoming	Kubernetes cluster	TCP	Velero S3 communication for metadata backups and restores

1 TCP:30000-32767 are supported only on releases 6.6.0a to 6.6.0d_u1.

2 TCP:32000-(32000+#<worker nodes>) are supported only on releases 6.6.0d_u2 and later.

Minimum Permission

To register the Kubernetes cluster with the Cohesity cluster, you must provide the Bearer token ID of the Kubernetes service account. Ensure that the service account has the cluster-admin role on the Kubernetes cluster.

> **Note:** This Bearer token ID will be used for all communications with the Kubernetes API server.

Prerequisites

Ensure that you meet the following prerequisites to register and manage your Kubernetes cluster:

- **Valid Certificate with Subject Alternative Name (SAN) field** —On the Cohesity cluster, install a valid certificate with the SAN field. The SAN field is used to specify additional host names (such as sites, IP addresses, and common names) to be protected by a single SSL certificate, such as a Multi-Domain or Extend Validation Multi-Domain Certificate. For more information, see Update SSL Certificates Using Cohesity CA or Customer CA certificate for Agent to Cluster Communication.
- **Create a Cohesity service account on Kubernetes to register Kubernetes on Cohesity cluster**—For registering a Kubernetes cluster with the Cohesity cluster, you can either use a default service account with the cluster-admin role or a dedicated service account with the cluster-admin role account. Cohesity recommends you to create a dedicated account in the default namespace.

  For example, to create a 'cohesity' service account on Kubernetes, use the following command:

  ```
  kubectl create serviceaccount cohesity -n default
  ```

  Assign the cluster-admin role to the 'cohesity' service account using the following command:

  ```
  kubectl create clusterrolebinding cohesity-admin --clusterrole=cluster-admin --serviceaccount=default:cohesity
  ```

  The above example can be run from any client or workstation from which `kubectl` is typically run for a given Kubernetes cluster that is registered and protected by the Cohesity cluster.

  > **Note:** Custom roles can be used as an Early Access feature after contacting Cohesity support.

- If you use a Kubernetes cluster 1.24 or later, run the following command to generate the Bearer Token ID or secret token for your service account.

  ```
  kubectl apply -f - <<EOF
  apiVersion: v1
  kind: Secret
  metadata:
    name: cohesity-token
    annotations:
      kubernetes.io/service-account.name: cohesity
  type: kubernetes.io/service-account-token
  EOF
  ```

  Sample Ouput

  ```
  secret/cohesity-token created
  ```

- **Extract the Bearer token ID**—To register the Kubernetes cluster with Cohesity cluster, you must have the Bearer token ID of the Kubernetes service account. Extract the Bearer token of the service account that has admin privileges on the Kubernetes cluster.

  Extract the Bearer token using the following command:

  ```
  kubectl describe secrets -n default cohesity-token | grep token: | awk
  '{print $2}' | head -1

  eyJhbGciOiJSUzI1NiIsImtpZCI6IjRtZzBkX0VRLXMwV05veFdvcjR3eGw1N2c4MUk0RX
  lhWFBPdXBHNUVoTlkifQ.eyJpc3MiOiJrdWJqcm5ldGVzL3NlcnZpY2VhY2NvdW50Iiwia
  3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9uYW1lc3BhY2UiOiJkZWZhdWx0Iiwia3V
  iZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9zZWNyZXQubmFtZSI6ImNvaGVzaXR5LXRva
  2VuL2Y0ZG02Iiwia3ViZXJuZXRlcy5pbF9zZXJ2aWNlYWNjb3VudC9zZXJ2aWNlLWFjY29
  1bnQubmFtZSI6ImNvaGVzaXR5Iiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9zZ
  XJ2aWNlLWFjY291bnQudWlkIjoiZDg4NmE5ZWMtMzQxYi00ZTM3LTg1MjktMzYxNGJiY2U
  0OTAxIiwic3ViIjoic3lzdGVtOnNlcnZpY2VhY2NvdW50OmRlZmF1bHQ6Y29oZXNpdHeif
  Q.Of4nzdrTDsbxzHWbzskwk6_
  mdwSvduSLPhSivW3MdHJNcGMa6nGmn5iKUQqDG4oe1tu71w0r9rT1Cp8sGHOLxpjGEhaKS
  f7gIGX5fUbw2jY6m7gL1U-4zj5FrglJ9XdE1yE70qNE3nFL1QZexYjAW80_
  rjMBv0lW4unq1o2IO21e5PC31SIGMfs1rlSvToy4_
  vJEVFzX9E0e4ha1RSsEIYFjdwt3sgG6y-
  PGoR4UBgnNnBaMpmy6NXYToaChzteJw93HQyyRVGBFNYgNTinqc5edu_
  zyMdAl6Chba6q10ffjC_zjhU_Lm7wp31-BTMelg-1W_5b3uni4pCGD8EPGPw
  ```

- **Obtain the Cohesity Datamover Docker Image**—To back up the PVCs of the Kubernetes namespaces, you must push the Cohesity Datamover image to an image registry accessible by your Kubernetes cluster.

  To install the Cohesity Datamover image:

  1. Obtain the Datamover image that matches closely with your Cohesity cluster version from the Cohesity Download page.
  2. To push the Datamover image to your Docker registry:
     1. Run the following command to log into your Docker image registry:

        ```
        docker login <your image registry server>.fqdn
        ```

        If prompted, enter the credentials that have the permissions to push the Cohesity Datamover image to your image registry. Skip this login command if your image registry does not require authentication to push images.
     2. Load the Datamover image:

```
docker load --input cohesity-datamover-<your Datamover version>.tar
```

Replace <your Datamover version> in the preceding command with the version in your Datamover image name.

Example:

```
docker load --input cohesity-datamover-6.6.0d_u6.tar
```

```
docker load --input cohesity-datamover-6.8.1_u1.tar
```

3. Run the following command, and from the command result, make a note of the Image ID of the Datamover image:

```
docker images
```

4. Tag the Datamover image with the Image ID you made a note of in step c:

```
docker image tag <ID of your image> <your image registry server>.fqdn/path/cohesity-datamover:<your Datamover version>
```

5. Push the Datamover image to your Docker registry:

```
docker push <your image registry server>.fqdn/path/cohesity-datamover:<your Datamover version>
```

> **Note:** The preceding example assumes you use Docker to load, tag, push, and pull the Docker images. If you use Podman, replace `docker` with `podman`.

Considerations

- Cohesity uses the third-party software Velero to facilitate the backup and restore of namespace data.

  If your Kubernetes clusters are of versions earlier than 1.16, use Velero 1.4. If your Kubernetes clusters are of versions 1.16 or later, use Velero 1.16.2. For more information, see Upgrade Velero and Datamover.
- If Kubernetes is not accessed through the internet, you must provide the path to the registry from which Velero should be pulled.
- Cohesity cluster only supports the backup of the user-created application namespaces of the Kubernetes cluster and does not support the backup of default or infrastructure namespaces such as kube-node-lease, kube-public, kube-system, and so on.
- Cohesity supports restoring namespaces backed up from one Kubernetes cluster to a different Kubernetes cluster if the storage class of the Kubernetes clusters is the same.

- Cohesity does not support Tanzu Kubernetes Grid Integrated Edition (TKGI).
- Cohesity does not support data protection of Kubernetes clusters in private networks behind NAT gateways (Dynamic/Static NAT, Port Address Translation PAT).
- For new protection jobs, S3 embedded credentials are used by default.
- For existing protection jobs, incremental backups continue using the Cohesity S3 account ID. S3 embedded credentials are used only after a full backup is performed.
- For recovery, the credential type depends on the snapshot selected:
  - If the snapshot was created using the Cohesity S3 account, recovery uses the same.
  - If the snapshot was created using S3 embedded credentials, recovery uses embedded credentials.
- For replication, the target cluster must support S3 embedded credentials (version 7.3 or later). Otherwise, the replication job will fail.

> **Note:** Do not delete the Cohesity user if any protection job is still using the old approach (Cohesity S3 account ID). Deleting the user will cause incremental backups to fail, and existing snapshots will become unrecoverable.

- When restoring to an alternate cluster, Cohesity (through Velero) does not automatically restore Custom Resource Definitions (CRDs). This is an intentional design choice meant to avoid potential conflicts or compatibility issues that could occur if the target cluster already has CRDs with different versions or schemas.
- In Kubernetes and OpenShift environments, restore operations might fail partially with the **PartiallyFailed** status due to resource restore order dependencies. This might happen in namespaces with complex inter-resource dependencies. The following error appears:

  "[KubernetesError]: Restore job failed with reason PartiallyFailed."
  Workaround: If the restore operation fails partially due to resource restore order dependencies, retry the restore operation. The initial restore attempt might have already recovered prerequisite resources, and therefore the dependency that caused the failure might be satisfied during the subsequent restore attempt.

Register and Manage Kubernetes Cluster

To protect the namespaces of the Kubernetes cluster, you must first register your Kubernetes cluster as a source on the Cohesity cluster. After registering your source, you can also perform the following actions:

- Update the Kubernetes cluster source configuration.
- Refresh the source details.

- Unregister the Kubernetes cluster from the Cohesity cluster.
- Configure User-defined Priority Class, Labels, and Annotations for Cohesity Deployed Resource

Register Kubernetes Cluster as a Source

Before you register the Kubernetes cluster as a source on the Cohesity cluster, ensure that you:

- Review the requirements, considerations, and best practices.
- Perform the prerequisite tasks.

For more information, see Plan and Prepare for Protecting Kubernetes Clusters.

To register the Kubernetes cluster as a source on the Cohesity cluster:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** and then select **Kubernetes Cluster**.

   The **Register Kubernetes Source** page appears.

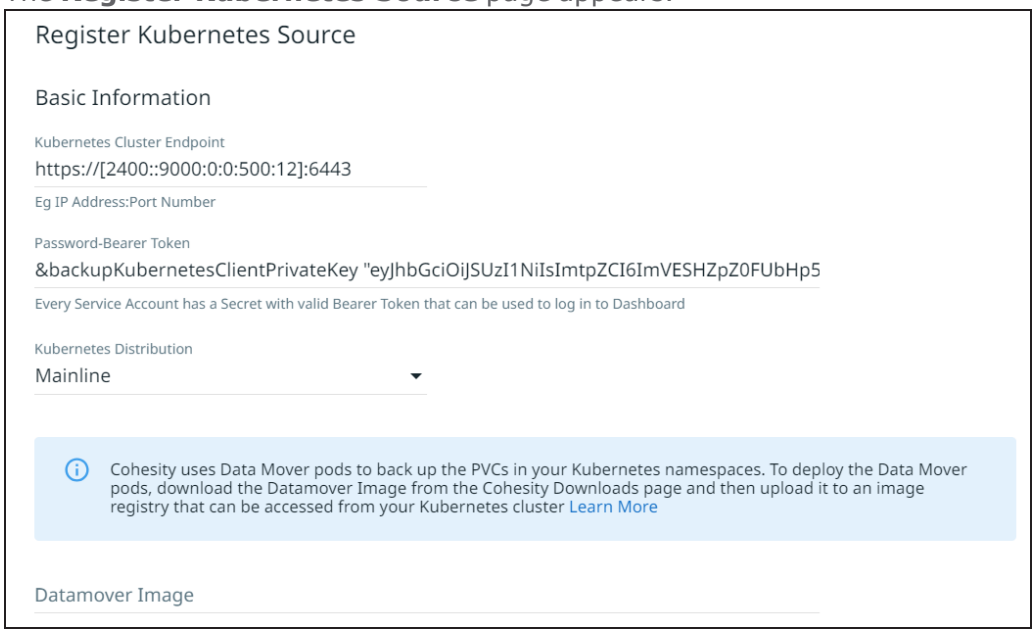

3. In the **Kubernetes Cluster Endpoint** field, based on your requirements, enter the Kubernetes API Server endpoints in one of the following formats:

   https://<FQDN>:<port_number>

   IPv4

   Format: https://<IPv4 address>:<port_number>

   Example: https://10.14.58.2:6443

   IPv6

Format: https://[<IPv6 address>]:<port_number>

Example: https://[2400::9000:0:0:500:12]:6443

4. In the **Password-Bearer Token** field, enter the Bearer token ID of the Kubernetes service account. For information on obtaining the Bearer token ID, see Plan and Prepare for Protecting Kubernetes Clusters.
5. From the **Kubernetes Distribution** drop-down, select one of the following Kubernetes distributions you want to backup and recover:
   - Mainline (Vanilla Kubernetes)
   - Redhat Openshift
   - VMware Tanzu

   > **Note:** Rancher (RKE2), Azure Kubernetes Service (AKS) (supported only on Cohesity Cloud Edition on Azure), and Nutanix Kubernetes Platform (NKP) managed Kubernetes clusters are not listed as separate options during registration. These clusters are supported as upstream Kubernetes distributions. Select **Mainline** to proceed with the registration.

6. *(Optional)* **Velero OpenShift Plugin Image**: This field is displayed if you have selected Redhat OpenShift as the Kubernetes Distribution.

   Download the latest plugin from the Red Hat website. In the **Velero OpenShift Plugin Image** field, enter the location where the Velero Openshift image is saved in your Docker registry. If you leave this field blank, the Cohesity cluster will automatically attempt to pull the latest image from the Red Hat website.

   You can download the Cohesity Velero OpenShift Plugin Image from the Cohesity Downloads page.

   > **Note:** This option is available only if you have selected Redhat OpenShift as the Kubernetes Distribution.

7. **Datamover Image**: Enter the location where you have the Datamover image saved in your Docker registry. The image must be entered in the format: `<Image_name>:<TAG>`.

   > **Note:** To protect the namespace level data in Vanilla Kubernetes 1.22 clusters, use the DataMover version that matches closely with your Cohesity cluster version. For more information, see Upgrade Velero and Datamover.

8. *(Optional)* **Service Type**: Enable this field to route external traffic to your Kubernetes cluster node via dedicated VLANs.

   From the drop-down, select one of the following service types:

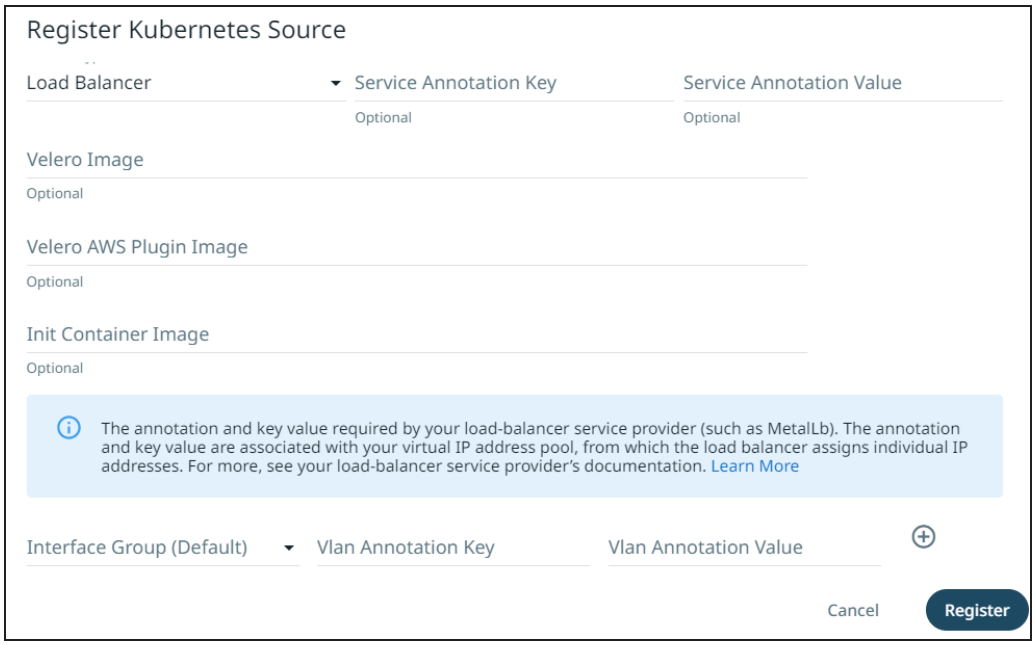

   - **NodePort** service type allows you to route external traffic through a fixed port (which you can set within the range of 30000-32767) on every Kubernetes node. In a NodePort setup, external traffic has direct access to the Kubernetes node.
   - **Load Balancer** service type allows you to route external traffic via an external Load Balancer functionality. Load Balancers monitor the availability of your Kubernetes pods and route incoming service requests to the relevant pods. In a Load Balancer setup, external traffic has access to the Kubernetes node through the Load Balancer only.

     Service Annotation Key and **Service Annotation Value**

     Enter the Service annotation key-value pairs. These key-value pairs are utilized to store metadata that offers information about the selected Kubernetes service that Cohesity deploys. This metadata does not affect the behavior of the Kubernetes cluster.

     > **Note:** The Service Type field is an Early Access feature. Contact your Cohesity account team to enable it.

9. *(Optional)* **Velero Image**: Enter the location where you have the Velero image saved in your Docker registry. If you leave this field blank, then the Velero image will be

pulled from a public image registry on the internet. You can download the Cohesity Velero Image from the Cohesity Downloads page.

> **Note:** You need Velero 1.8.1 to protect the namespace level data in Vanilla Kubernetes 1.22 clusters. For more information, see Upgrade Velero and Datamover.

10. *(Optional)* **Velero AWS Plugin Image**: Enter the location where you have the Velero plugin image for AWS saved in your Docker registry. If you leave this field blank, then the Velero plugin image for AWS will be pulled from a public image registry on the internet. You can download the Cohesity Velero AWS Plugin Image from the Cohesity Downloads page.
11. *(Optional)* **Init Container Image**: Enter the location where you have the Init Container image saved in your Docker registry. If you leave this field blank, then Cohesity will use the Datamover image as the Init Container image.
12. **Interface Group (Default)**: Displays all the Interface Groups set in your cluster. Click the **+** icon to add as many Interface Groups as needed. Once the Interface Groups are created, the traffic of the Service Type (you selected in Step 8) will only flow through the Interface Group(s), creating a separate channel for the traffic.

    Vlan Annotation Key and **Vlan Annotation Value**: Displayed only if the **Load Balancer Service Type** is selected. Enter the key-value pair to specify a VLAN that the Load Balancer service deploys to.

    Load balancer services allow you to configure a specific kind of address pool (in this case VLAN). For example, if you use the MetalLB LoadBalancer service, you need to set the **Vlan Annotation Key** to "metallb.universe.tf/address-pool" and the **Vlan Annotation Value** should be the name of the address pool. For more information, see the documentation of your LoadBalancer service.

    > **Note:** The Interface Group field is an Early Access feature. Contact your Cohesity account team to enable it.

13. Click **Register**.

Post-Registration Task

Once you register the Kubernetes cluster with the Cohesity cluster, perform the following actions:

- Run the following command to view the ports associated with each node of the Kubernetes cluster.

```
kubectl get services -n cohesity-<cluster-id>
```

Ensure the ports associated with the nodes are open to enable communication between the Kubernetes cluster and the Cohesity cluster.

For information on other required ports, see Network Firewall Considerations.

- Verify all of the pods in the Kubernetes cluster are in running state by running the following command:

```
kubectl get pods -n cohesity-<cluster-id>
```

The namespaces are displayed on the **Sources** page of the Cohesity Dashboard only after all the pods are either in the running state or error state.

- Refresh the Kubernetes source if there is an error on the source page. For more information, see Refresh the Kubernetes Source.

Update the Kubernetes Cluster Source Configuration

You can update the Kubernetes cluster details that you had provided during the registration process with the latest Kubernetes cluster configuration.

To edit the Kubernetes cluster source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Kubernetes cluster source and select **Edit**.
3. In the Register Kubernetes Source page, update the respective configurations.
4. Click **Register**.

> **Note:** When a Kubernetes source is edited to update the Velero image, Velero AWS plugin image, OpenShift plugin image, or DataProtect plugin image, the changes are not applied to the source, even though the edit appears to be successful.
> As a workaround, manually delete the Velero deployment from the Cohesity namespace in the source Kubernetes cluster, and then perform the source edit.

Refresh the Kubernetes Source

You can refresh the Kubernetes and fetch the latest changes on the Kubernetes cluster. To refresh the Kubernetes cluster source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Kubernetes cluster source and select **Refresh**.

Unregister Kubernetes Cluster

If you plan to stop taking the backup of your Kubernetes cluster namespace, you can unregister the Kubernetes cluster source from the Cohesity cluster.

Before unregistering, ensure to delete all the Protection Groups associated with the Kubernetes source.

To unregister the Kubernetes cluster:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Kubernetes cluster source and select **Unregister**

Configure User-defined PriorityClass, Labels, and Annotations for Cohesity Deployed Resources

You can provide the priority class names for pods deployed by it. You also have the option to select labels and annotations that will be applied to the pods and services deployed by Cohesity.

The Kubernetes source must include a PriorityClass resource with the specified name. The PriorityClass is assigned to each pod in its configuration. Labels track pods, while annotations on pods and services manage external network access rules.

You can specify the priority class for pods and also apply user-defined labels and annotations to the pods and services deployed by Cohesity at the Kubernetes source for data protection.

To provide Priority Class, Labels, and Annotations:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** and then select **Kubernetes Cluster**.

   The **Register Kubernetes Source** page appears.
3. Click the actions menu (⋮) next to the Kubernetes cluster source and select **Edit**.
4. Select the **Optional** field.
5. (*Optional*) **Priority Class Name**: In the Priority Class Name field, enter the name of the PriorityClass object available at the source. Pods deployed by Cohesity will specify this priority class name in their specifications. If a priority class for the pods is

specified during the registration or modification of the source, the priority class will first be validated to ensure it exists at the source. If it does not exist, the registration or editing process will fail.

During registration, the Velero and Datamover pods use the priority class to ensure they receive the appropriate priority assignment.

The Velero and Datamover pods will be updated with the new value during modification. Consequently, the pods will restart automatically.

If a priority class was specified during registration and you later deleted it from the Kubernetes source, the existing pods will remain unchanged. Upon the next refresh, it will be detected that the priority class no longer exists in the source, and an error message will be displayed on the sources page.

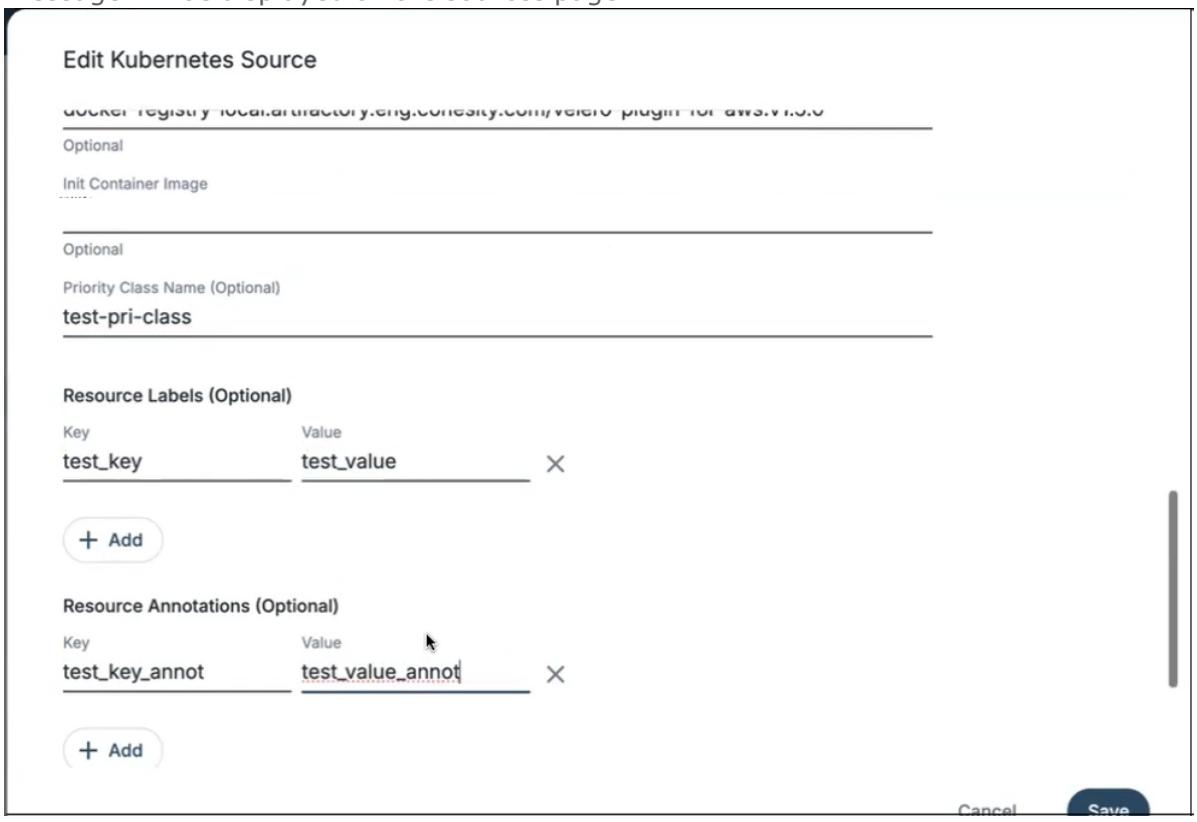

> **Note:** Before adding the priority class name, make sure that the priority class is already existing in the Kubernetes source.

6. (*Optional*) **Resource Labels Key** and **Resource Labels Value**: Enter the Resource Labels key-value pairs.

   (*Optional*): **Resource Annotations Key** and **Resource Annotations Value**: Enter the Resource Annotations key-value pairs.

If labels or annotation maps are provided during source registration, Velero and Datamover pods and services will be labeled and annotated accordingly.

When source modifications occur, such as changes to labels or annotation maps, both the Velero and Datamover pods are updated with the new labels and annotations. This update triggers a restart of the pods, which can disrupt ongoing backup and recovery processes. Any subsequent attempts within the same run will be successful. While the Datamover services will be updated with new labels and annotations, they will not be restarted.

If a label or annotation with a specific key is deleted, the associated pods and services will be updated to reflect an empty string as the value for those keys. Additionally, if the Velero deployment or the data mover DaemonSet and services are deleted, and a manual refresh of the source is initiated, any labels and annotations with empty string values will be removed.

When you perform a search, you can filter the labels or annotations you provide and list the resources accordingly.

This is supported for both Load Balancer and NodePort service types.

Diagnostics Collection

The Kubernetes cluster log collection feature provides a streamlined way to gather comprehensive diagnostic data required for troubleshooting backup and restore issues.

This capability collects:

- Cluster-scoped resources
- Namespace-scoped resources
- Velero pod logs
- Datamover pod logs
- Cohesity namespace objects
- Additional user-selected namespaces

This feature captures the state of the cluster at the time of issue, reducing manual effort and improving troubleshooting efficiency.

To collect the diagnostics:

1. Navigate to **Data Protection > Sources > Register**. Register a Kubernetes source.
2. Select the registered Kubernetes source.
3. Click **More options** (⋮).
4. Select **Collect Diagnostics**.

The Collect Diagnostics dialog appears.

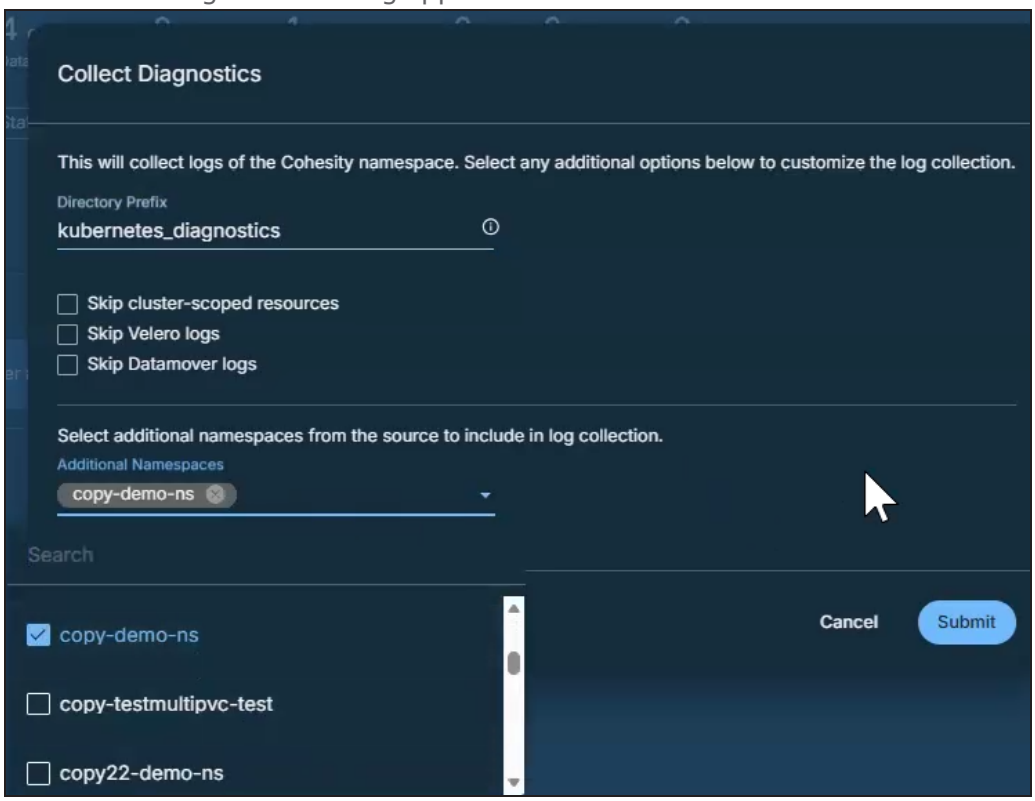

5. Enter a folder or directory name to store the diagnostic bundle.
6. You can choose to skip specific data:
    - Skip cluster-scoped resources
    - Skip Velero logs
    - Skip Datamover logs

    By default, the Cohesity namespace is included.
7. Optionally select additional namespaces for extended diagnostics.
8. Click **Submit** to initiate log collection.

Diagnostic data is collected as a bundle. Status is recorded in the Audit Log. Data can be downloaded and shared with support teams

Upgrade Velero and Datamover

Cohesity uses third-party software Velero 1.4 to back up and recover data at the namespace level.

If Kubernetes is not accessed through the internet, you must provide the path to the registry from which Velero should be pulled. For Cohesity Cluster version 7.2 and later, if

Kubernetes is accessible from the internet, the Cohesity Kubernetes code will pull Velero version 1.16.2.

Configure Velero while Registering a Kubernetes Cluster

When you are registering the Kubernetes cluster as a source or re-registering a deleted Kubernetes cluster, on the Register Kubernetes Source page:

- In the **Velero Image** field, enter the location of the Velero image saved in your Docker registry.
- In the **DataMover Image** field, enter the location of the supported DataMover image saved in your Docker registry.

Upgrade Velero in Existing Kubernetes Cluster Sources

If you have existing Kubernetes Clusters registered with the Cohesity cluster, verify whether the Cohesity namespace is present or absent in the registered source. The procedure to upgrade Velero is different when the Cohesity namespace is present and when the Cohesity namespace is absent.

Verify if the Cohesity Namespace is Present

Run the following command on the master node of the Kubernetes cluster:

```
- kubectl get ns
```

Sample Output:

NAME	STATUS	AGE
<auto-rec-jenkins-del	Active	23h
cohesity-21234567891234	Active	7d8h
default	Active	24d
Jenkins	Active	21d
kube-node-lease	Active	24d
Kube-public	Active	24d
Kube-system	Active	24d
longhorn-system	Active	20d
mysql	Active	12d
wordpress	Active	21d

Check if there are namespaces with Cohesity cluster IDs in the command output. For example, `cohesity-21234567891234` is the Cohesity namespace in the preceding sample.

Upgrade When Cohesity Namespace is Absent

If the Cohesity namespace is absent, no warnings or alerts are displayed in the Cohesity cluster regarding the requirement to upgrade Velero and Data Mover.

To upgrade Velero:

1. Run the following command on any node of the Kubernetes cluster (or from a node in another Kubernetes cluster that has access to the Kubernetes command-line tool) to delete the Velero Custom Resource Definition (CRD):

   ```
   - kubectl delete crds -l component=velero
   ```

2. To edit the Kubernetes cluster source configuration:
   a. Navigate to **Data Protection** > **Sources**.
   b. Click the Actions menu (⋮) next to the Kubernetes cluster source and select **Edit**.
   c. In the Register Kubernetes Source page, enter the locations of the new Velero and Data Mover images in the **Datamover Image** and **Velero Image** fields, respectively.
   d. Click **Register**.

   When registration is complete, a new entry appears on the Sources page.

Upgrade When the Cohesity Namespace is Present

If the Cohesity namespace is present in a registered Kubernetes cluster, a message is displayed in the Cohesity cluster indicating that you must upgrade the Cohesity data protection resources in the Kubernetes cluster. The message will also help you understand which Velero and Data Mover versions you must upgrade to.

To upgrade Velero:

1. Stop all the Kubernetes protection jobs that are currently running and ensure there are no active protection runs on the Kubernetes source.
2. Run the following command to delete the Cohesity namespace in the Kubernetes cluster:

   ```
   - kubectl delete ns <name of the Cohesity namespace>
   ```

   Example

   ```
   - kubectl delete ns cohesity-21234567891234
   ```

3. Run the following command on any node of that Kubernetes cluster (or from a node in another Kubernetes cluster that has access to the Kubernetes command-line tool) to delete the Velero CRD:

   ```
   - kubectl delete crds -l component=velero
   ```

4. To edit the Kubernetes cluster source configuration:
   a. Navigate to **Data Protection** > **Sources**.
   b. Click the Actions menu (⋮) next to the Kubernetes cluster source and select **Edit**.

c. In the Register Kubernetes Source page, enter the locations of the new Velero and Data Mover images in the **Datamover Image** and **Velero Image** fields, respectively.

d. Click **Register**.

When registration is complete, a new entry appears on the Sources page.

Backup Kubernetes Cluster Resources, PVCs, and Namespaces

Cohesity supports the backup of PVCs, Kubernetes cluster resources and metadata of the Kubernetes namespaces. Cohesity uses the Data Mover pod that will be deployed in the Kubernetes cluster for backing up the PVCs of the Kubernetes namespaces.

Cohesity supports the backup of PVCs with the volumeMode set to Filesystem.

The supported file systems include ext3, ext4, xfs, vxfs, ReiserFS, autofs, cvfs, zfs, fuse.mfs, nfs, nfs3, cifs, gpfs, and mmfs.

Cohesity also supports the backup of PVCs with CSI drivers using CSI volume snapshotting capability.

> **Note:** The backup of PVCs with CSI drivers using CSI volume is an Early Access feature. Contact Cohesity Support support for enabling this feature.

After registering the Kubernetes cluster as a source on the Cohesity cluster, you can start configuring your backups. Cohesity recommends the first full and an incremental forever approach for the backup. However, you can do periodic full backups, based on your requirements.

The temporary pods created during backups by Cohesity will have their priority class set to the specified value for the source. Additionally, the labels and annotations will be applied to the temporary pods.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Protect Kubernetes Namespaces

In the Cohesity cluster, Kubernetes namespaces are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group.

To protect the namespaces of the Kubernetes containers:

1. In the Cohesity cluster UI, perform one of the following:
   - Select **Data Protection > Protection**. Click **Protect** displayed on the top-right of the page and then select Kubernetes Cluster.
   - Select **Data Protection > Sources**. Click the Actions menu (⋮) next to the Kubernetes source and then select Protect.
2. From the **Registered Source** drop-down list, select the registered Kubernetes source or click **Register Source** to register an Kubernetes source. For more information, see Register Kubernetes Cluster as a Source.
3. From the **Objects** list, select the namespaces that you plan to backup and click **Save Selection**.

   > **Note:** You can use the advanced filters that are available to search and find the namespaces you plan to protect. Also, you can click the shield icon to enable Auto Protect at the Kubernetes cluster level or based on the namespace label(s). For more information, see Enable or Disable Auto Protect.

   Exclude or Include PVCs

   Cohesity enables you to choose which PVC to back up or exclude in a Kubernetes Namespace or Protection Group. (For details on how to include or exclude a PVC from a Protection Group, see the Advanced Settings section.)

   To include or exclude PVCs for Kubernetes namespace:

   a. In the **New Protection** dialog, click the **Edit** (✎) icon next to that selected namespace.

      > **Note:** To turn ON the **PVC Inclusion/Exclusion** option to an existing Protection Group, click the Actions menu (⋮) for that Protection Group and select **Edit**. In the **Edit Protection** dialog, under **Objects**, select the namespace and click the **Edit** icon.

   b. In the **Options** dialog, turn ON the **Enable PVC Inclusion/Exclusion** toggle.

c. You can choose to either include or exclude the PVCs. Select **Include** or **Exclude**, and then select the PVCs from the drop-down.

   After selecting an option, the **Advanced: Configure Metadata per PVC** setting becomes available.

   - **Include**: If you included PVCs, then you can additionally use the advanced dialog to configure metadata-only backup for such PVCs.
     - Full data and metadata, or
     - Metadata only by enabling the **Only Metadata** toggle.

     The **Only Metadata** toggle determines whether a PVC is backed up with metadata only or with both data and metadata.

     Example

     - PVC1: Only Metadata enabled - Only metadata is backed up.
     - PVC2: Only Metadata disabled - Full data and metadata are backed up.

     If there are any other non selected PVCs, those will get fully excluded.
   - **Exclude**: If you excluded PVCs, then you can additionally use the advanced dialog to configure metadata-only backup for such PVCs.
     - If Only Metadata is disabled - The PVC is completely excluded (no data or metadata is backed up).
     - If Only Metadata is enabled - Only the metadata is backed up, and the data is excluded.

     Example:

     - PVC3: Only Metadata enabled - Only metadata is backed up.
     - PVC4: Only Metadata disabled - PVC is fully excluded from backup.

     Any PVCs which were not selected will get fully included.

d. Click **Save**.

Once saved, the **Customized** link appears next to the **Edit** icon. Click the link if you want to modify the **PVC inclusion/exclusion** settings.

> **Note:** PVC Inclusion or Exclusion settings of a Kubernetes namespace take precedence over the settings of its Protection Group.

Exclude or Include Resources

Cohesity enables you to choose which Custom Resources to back up or exclude in a Kubernetes namespace or Protection Group. (For details on how to include or exclude a resource from a Protection Group, see the Advanced Settings section.)

To include or exclude Resource for Kubernetes namespace:

a. In the **New Protection** dialog, click the **Edit** ( ) icon next to that selected namespace.

b. In the **Options** dialog, turn ON the **Enable Resource Inclusion/Exclusion** toggle.

c. You can choose to either include or exclude the resource. Select **Include** or **Exclude**, and then select the resource from the drop-down.

   You can add other custom resources by clicking **Add Custom Resources**.

   - When you select a resource from the drop-down menu, only that specific resource will be backed up.
   - When you select the **Backup only PVC and related resources** option, the **Enable Resource Inclusion/Exclusion** option will be disabled. As a result, you can back up only PVCs and dependent resources, and skip other resources (for example Secrets).
   - To back up a PVC and include or exclude other resources, you have two options:
     - **Specify Inclusion/Exclusion for Individual PVCs**: Navigate to the PVC Inclusion/Exclusion panel to configure the inclusion or exclusion settings for individual PVCs. Instead of selecting the **Backup Only PVC** option, configure resource-level inclusion/exclusion through the **Enable Resource Inclusion/Exclusion** dropdown. In this scenario, the PVC-related resources will be automatically identified, since the individual PVC inclusion/exclusion settings have been specified.
     - **Configure Resource-Level Inclusion/Exclusion Directly**: Instead of setting inclusion/exclusion at the individual PVC level, you can directly configure resource-level inclusion/exclusion through the **Enable Resource Inclusion/Exclusion** dropdown. With this method, you must explicitly select the PVC-related resources you want to include or exclude, along with any other resource categories you wish to manage.

d. Click **Save**.

Once saved, the **Customized** link appears next to the **Edit** icon. Click the link if you want to modify the **Resource inclusion/exclusion** settings.

4. In the Protection Group section, perform one of the following:

- Select **New Group** to create a new Protection Group. In the **Name** field, enter a name from the Protection Group and go to step 5.
- Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.

5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, Create or Use an Existing Protection Policy .
6. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
7. In the **Start Time** field, select the time and time zone to start the protection run. This option is available only if the selected policy is set to Backup Daily. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
8. Cohesity supports protecting PVCs with the snapshot and restore feature of Container Storage Interface (CSI) drivers on CSI volumes. This option is disabled by default in your Cohesity protection job. To enable this feature, turn ON the **Leverage CSI snapshot** toggle. Once enabled, Cohesity clusters will begin utilizing snapshots captured by CSI drivers to protect the PVCs. These snapshots are crash-consistent, meaning that the PVC is in the same state as it was during the crash.
9. If you need to change any of the advanced settings on the Protection Group page, click the down arrow icon next to **Advance Settings** and click **Edit** on the right. For more information, Advanced Settings.
10. Click **Protect**.

    Your new Protection Group is now active and running.

Advanced Settings

Advance Settings	Description
Pause Future Runs	Enable **Pause Future Runs** to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Enable **End Date** and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
Priority	Select a priority for the Protection Group eProtection Groupxecution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs

Advance Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	Enter a brief description about the Protection Group.
Include or Exclude PVC	Cohesity enables you to choose which PVCs to include or exclude in the Protection Group. To select the PVCs: a. Turn ON the **Enable PVC Inclusion/Exclusion** toggle. b. In the **Logical Rule** drop-down, select “Match Any of the following labels” or “All of the following labels”. c. Select **Include** or **Exclude**. d. Click **+ Add**. e. Enter the PVC’s label. The label is the unique identifier assigned to that PVC in the Kubernetes cluster. You can add more PVC until you have all the PVCs you want to be included or excluded for that Protection Group. > **Note:** PVC Inclusion or Exclusion settings of the Namespace take precedence over the settings of its Protection Group.
Include or Exclude Resource	Cohesity enables you to choose which resource types to include or exclude in the Protection Group. To select the Resource: a. Turn ON the **Enable Resource Inclusion/Exclusion** toggle. b. Select **Include** or **Exclude**. c. Click **+ Add**.

Create or Use an Existing Protection Policy

A Protection Policy is a reusable set of settings. It enables you to define how and when Kubernetes namespaces are protected, replicated, and archived. A standard policy provides various settings that you can configure for a Protection Group.

Create a Protection Policy with the settings listed in the Policy Settings for Kubernetes. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Kubernetes

The following table lists the applicable policy settings for Kubernetes:

Field	Description
DataLock	If applicable, add a DataLock for compliance and regulatory requirements, to ensure that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users, including those who have the Data Security role in Cohesity Data Cloud, from modifying or deleting any Snapshots that were generated by the Protection Groups that use this Policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Backup	Defines when backups are captured, how long they are retained and the type of backup to capture.
Periodic Full Backup	Use this option to choose when to run a periodic full backup.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
Quiet Times	Quiet times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet times are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet times.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, the Cohesity cluster attempts to capture Snapshots three times before the Protection Group fails. The default time between retries is five minutes. You can customize the number of retries and how long to wait between each attempt.

Field	Description
Replication	To replicate snapshots to another, click the **Replication** icon from the floating menu. The copied snapshots are stored in the storage domain of the remote cluster.
Archival	To add an archive schedule, click the **Archival** icon from the floating menu.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with applicable policy settings for Kubernetes cluster. In this case, all the policy settings apart from the settings applicable to Kubernetes cluster are ignored.

To use and existing Protection Policy:

1. In the Cohesity cluster, navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for Kubernetes cluster. For information on the applicable policy settings, see Policy Settings for Kubernetes.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage the Kubernetes Backup Configuration

In Cohesity Data Cloud, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered Kubernetes cluster and the assigned Storage Domain in the Protection Group settings.

Add Namespaces to the Protection Group

You can add new namespaces to the Protection Group for backup. The newly added namespaces are backed up from the next protection run, and it will be a full backup.

To add namespaces to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for the Kubernetes cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.

4. On the Kubernetes page, click the edit icon in the **Objects** section.
5. From the namespaces list, select the namespaces that you plan to add to the backup and click **Save Selection**.
6. Click **Save**.

Remove Namespaces from the Protection Group

You can remove the Kubernetes namespaces from the Protection Group that you do not want to back up in the next protection run. When you remove a namespace from the Protection Group, the incremental backup of this namespace is not performed from the future protection run. The existing backup for the namespace is preserved until the snapshot expires.

To remove namespaces from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for the Kubernetes cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the Kubernetes page, click the edit icon in the **Objects** section.
5. From the namespaces list, unselect the namespaces that you plan to remove to the backup and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

You can enable the Auto Protect option at the:

- **Kubernetes cluster level**: Automatically protect the new namespaces that are added to the Kubernetes container.
- **Label level of the namespaces**: Automatically protect the new namespaces that are tagged with the same label you selected while enabling the auto protect.

When the Auto Protect option is enabled:

- All the namespaces that are added to the Kubernetes cluster or tagged with the same label in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the namespaces that are removed from the Kubernetes cluster or untagged from the specified label in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the namespace is preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly added namespaces to the Protection Group.

- From the Protection Group, remove existing namespaces that are deleted on Kubernetes. The existing backups of the namespaces are preserved until the snapshot expires.

Enable or Disable Auto Protect at Kubernetes cluster level

To enable or disable Auto Protect at the Kubernetes cluster level:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for the Kubernetes cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the Kubernetes page, click the edit icon in the **Objects** section.
5. Click the icon in the cluster level to enable or disable Auto Protect and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect at Label level

To enable or disable Auto Protect at the Kubernetes label level:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for the Kubernetes cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Kubernetes** page, click the edit icon in the **Objects** section.
5. Select the label icon ( ) available on the top-right of the **Objects** section.
6. From the **label** drop-down list, select the label(s) for which you want to enable **Auto Protect**.
7. Click the icon in the enable or disable Auto Protect and click **Save Selection**.
8. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Kubernetes.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up the namespaces of the Kubernetes cluster. It provides you with options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.
- Add replication.
- Add Archival.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Kubernetes.
3. Click the Actions menu next (⋮) to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    a. In the **Backup** drop-down list, select any one of the following:
        - **Backup all Objects in the Protection Group**
        - **Backup only selected Objects in the Protection Group**
    b. In the **Backup Type** drop-down list, select the backup type. By default, an incremental backup will be performed.
    c. Optional. In the **Replicate** section, select the replication target and the retention duration.
    d. Optional. In the **Archive** section, select the Archival target and the retention duration.
5. Click **Run Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Kubernetes.
3. Click the Actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had paused.
3. Click the Actions menu (⋮) next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- Delete Protection Group Only: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- Delete Protection Group and Snapshots: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

  After deleting the Protection Group and Snapshots, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Kubernetes.
3. Click the Actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - **Delete Protection Group Only**
   - **Delete Protection Group and Snapshots**

Quiesce Kubernetes Workloads during Backups

Quiescing involves pausing an application to achieve a consistent state in preparation for a backup.

Cohesity supports quiescing stateful Kubernetes workloads before capturing Persistent Volume Claims (PVCs) volume snapshots during backup in a namespace. After the snapshots are taken, the Kubernetes workloads are automatically unquiesced. This ensures consistent and reliable snapshots for stateful applications running in Kubernetes.

When running a protection job for a namespace, there are pre-snapshot and post-snapshot hooks. Before taking a snapshot of any PVC attached to a Kubernetes pod within an application, the application should be in a consistent state. After taking the snapshot, a post-snapshot hook is executed called "unquiescing," which allows applications to resume their operations.

You can define actions to be executed before and after taking a snapshot. Hooks allow specific scripts to run within a group of pods identified by key-value pairs known as label selectors (each pod will have its labels). You can select pods constituting an application (which needs to be quiesced) by specifying labels. Velero will execute the designated script within the pod's container and provide the exit status of the script.

Each script has a timeout feature, with a default of 60 seconds. You can extend this duration if necessary. The scripts are executed at the container level, running on each container within a pod. You can specify a container name to execute a script on specific containers within the same pod.

To quiesce Kubernetes workloads during backups:

1. Register Kubernetes source.
2. Edit the Protection Group that contains Kubernetes source.
3. In the **Add Objects** section, click the **Edit** icon.
4. In the list of protected objects, click the **Edit** icon next to the Kubernetes namespace for which you want to create or modify quiescing rules.
5. Switch to the **Scripts** tab and select **Quiesce Rules** to view the different quiesce modes.
6. Choose a Quiesce Mode to add different rules.. The following quiesce modes are available:
    - Apply following rules together
    - Apply following rules independently
    - Apply following rules sequentially

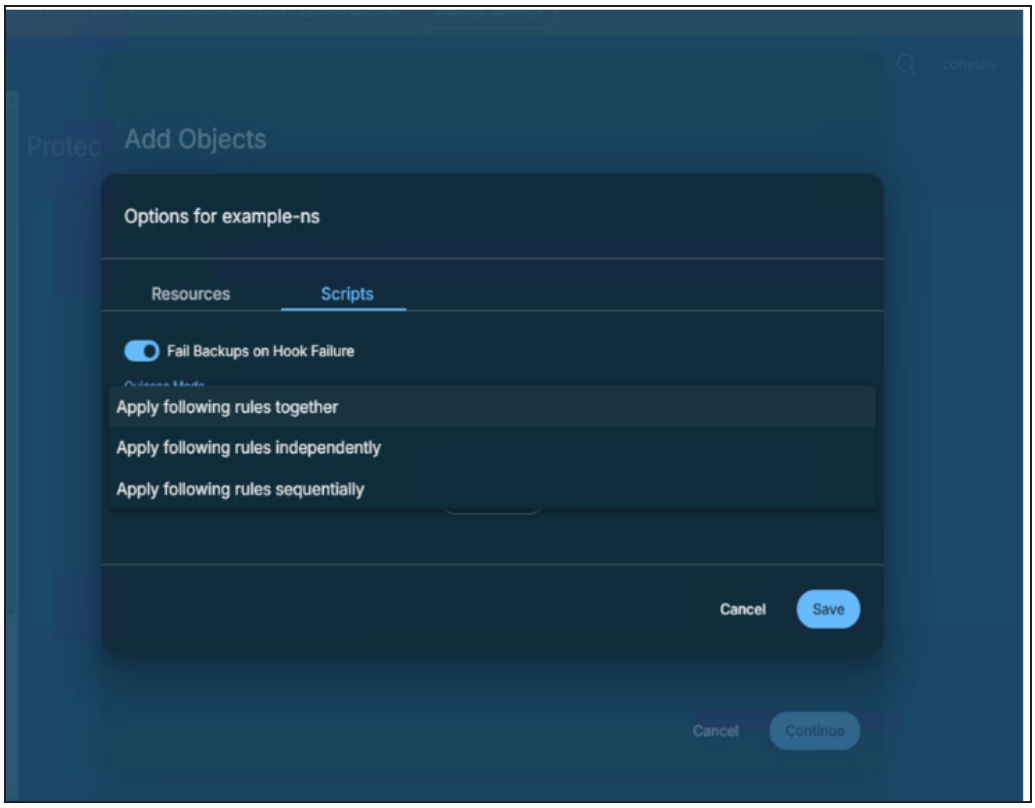

Application quiescing involves defining one or more rules (for example, Rule 1, Rule 2, and Rule 3) to pause or quiesce an application , allowing snapshots to be taken. Once the snapshots have been captured, the application resumes and the backup process continues. The execution of these rules is managed by the quiescing mode, which offers two options: parallel mode, in which rules are executed simultaneously, and sequential mode, where rules are executed one after another.

- **Together mode**: This mode operates using a single volume group, allowing all processes to run independently and in parallel. Quiescing occurs simultaneously, with all rules executed before a snapshot is taken. After applying the rules, snapshots are created, followed by unquiescing, and then the backup is performed. This approach results in faster backup speeds compared to sequential mode.
- **Independent mode**: This mode operates with multiple volume groups, allowing for parallel and independent processes. Quiescing happens simultaneously; for example, while volume group 1 is quiescing for Rules 1 and 2, volume group 2 may have already completed the quiescing and begun the snapshot process. As a result, this mode provides the fastest backup speed.
- **Sequential mode**: This mode operates with a single volume group, where all processes are interdependent and executed in a specific sequence. Quiescing occurs in stages, following defined rules. For example, Rule 1 is executed first,

then Rule 2, followed by Rule 3, and so on. After applying these rules, snapshots are created, the system is unquiesced, and the backup process is performed. This sequential approach applies to both the quiescing and unquiescing processes. Consequently, the backup speed in this mode is the slowest compared to the other two options.

In sequential mode, rules are executed one after another. If Rule 2 encounters an issue, you can either continue or stop all backups at the namespace level. This means that any error will cause the backup to fail for all rules. This mechanism manages backup processes in case of failures related to quiescing rules.

7. Click **Add Rule** to add a rule. You can add multiple rules by clicking **Add Rule**.
8. Add the Key value under the **Pod Selector Labels** for the rules to run on.

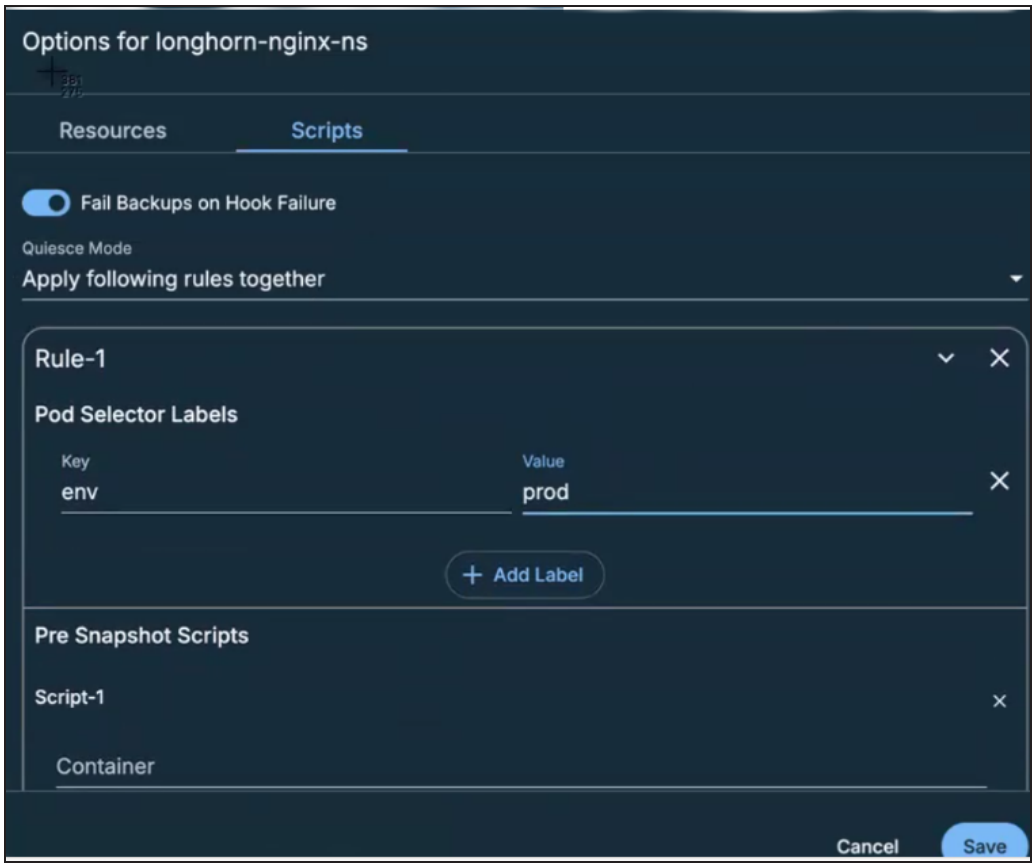

9. Select the **Pre Snapshot Scripts**.
10. Enter the container name, the command that needs to be run, and the timeout value in seconds. You can add multiple scripts by clicking Add Scripts for a particular rule selected under Pre Snapshots Scripts.
Enter the command `sh -c pre_script.sh` in the command prompt.

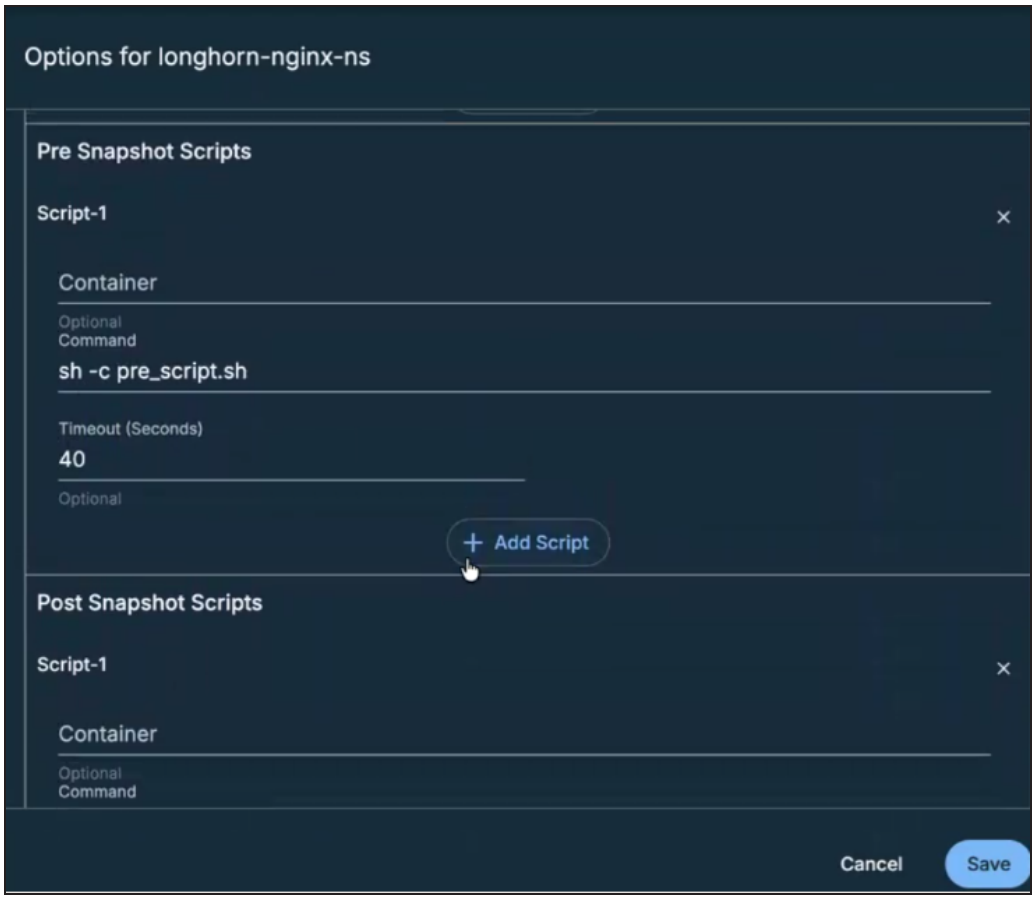

Example script on the source:

```
#!/bin/sh
LOG_FILE="/tmp/pre_script.log"
echo "$(date) : Starting WordPress maintenance mode ENABLE" >>
"$LOG_FILE"
if wp maintenance-mode activate >> "$LOG_FILE" 2>&1; then
echo "$(date) : Maintenance mode ENABLE completed successfully" >>
"$LOG_FILE"
else
echo "$(date) : ERROR - Maintenance mode ENABLE failed" >>
"$LOG_FILE"
exit 1
Fi
```

11. Select the **Post Snapshot Scripts**.
12. Enter the container name, the command that needs to be run, and the timeout value in seconds. You can add multiple scripts by clicking **Add Scripts** for a particular rule selected under **Post Snapshots Scripts**.

    Enter the command `sh -c post_script.sh` in the command prompt.

    Example script on the source:

```
#!/bin/sh
LOG_FILE="/tmp/post_script.log"
echo "$(date) : Starting WordPress maintenance mode DISABLE" >>
"$LOG_FILE"
if wp maintenance-mode deactivate >> "$LOG_FILE" 2>&1; then
echo "$(date) : Maintenance mode DISABLE completed successfully" >>
"$LOG_FILE"
else
echo "$(date) : ERROR - Maintenance mode DISABLE failed" >>
"$LOG_FILE"
exit 1
Fi
```

13. You can choose between "True" or "False" to indicate whether the backup should continue or fail when the script fails by selecting **Continue Backup when Script fails**. If the script fails, the protection job will still back up the PVC. However, if the script fails, the protection job will also fail.
14. Run the scripts. Ensure that the specified script is located at the correct path within the pod before you execute it.

    The pre-snapshot hook is triggered prior to taking snapshots of the PVCs. After the backup of the PVCs is completed, the post-snapshot hook is triggered to finalize the process. The Backup Task Activity for quiescing-ins, as shown in the pulse logs, will provide details about the execution of both the pre and post-snapshot hooks.

Recover Kubernetes

PVCs and the metadata of namespaces backed up to the Cohesity cluster can be restored. You can also perform granular recovery of files and folders within the PVCs.

- Recover Kubernetes Cluster
- Granular Recovery of Files and Folders

Recover Kubernetes Cluster

You can restore the PVCs and the metadata of the namespaces that are backed up onto the Cohesity cluster. You can restore the namespaces to the same location or to an alternate location.

Note:

- Cohesity does not support the granular-level recovery of the namespaces.
- When recovering a namespace in Kubernetes, you may encounter a "Timed out" error on the Cohesity UI if the namespace is being recovered to an existing namespace where one or more pods are already awaiting a volume/PVC resource being recovered.

  These pods might have been introduced in previous recovery attempts with certain excluded PVCs or through manual deployment.

Recovering Namespaces to the Same Location

When you are recovering the namespaces to the same or original location, then only the resources, PVCs, or the metadata of the Kubernetes namespaces that are missing or deleted from the original location are restored from the namespace snapshot you selected for recovery. Cohesity does not overwrite or restore the resources, PVCs, or the metadata of the Kubernetes namespace that already exists in the original location.

For example, a Kubernetes namespace you want to restore to the same location contains a deployment resource and a service account. When Cohesity initiates the restore to the original location and if the service account is missing from the original location, then Cohesity will restore only the service account. The restore of the deployment resource is skipped.

To restore the namespaces of a Kubernetes cluster to the same location:

1. In the Cohesity cluster UI, navigate to **Data Protection > Recoveries**.
2. Click **Recoveries** on the top-right of the page and then select **Kubernetes Cluster** as the recovery source.

   The Kubernetes Namespaces page appears.
3. Search for a namespace or Protection Group that contains the snapshots to recover from by entering characters of the namespace name, Protection Group name, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, Storage Domain, or date.
4. From the list of search results, select one or more objects (namespaces) or Protection Groups that contain the snapshots to recover.

> **Note:**
> - If you have selected a Protection Group, then all the namespaces backed up in the latest protection run of this Protection Group will be selected for recovery. The option to select these namespaces individually will be greyed out in the search result list.
> - You can only select Protection Groups and namespaces with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a namespace or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Protection Group	a. From the selected lists, mouse over the snapshot, and then click the **edit** icon. b. From the **Protection Run** drop-down menu, select the snapshot from which you want to recover the namespaces. By default the snapshot is recovered from the local cluster. c. Optionally, you can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a..	Then..
Namespace	a. From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <namespace_name> page appears. By default, the page displays the snapshots available in the last one month. b. Optionally, you can also select a snapshot from a different date range or by selecting the date range from the **Date** drop-down list. Once you select the date range and then click **Apply**, the Edit recovery point page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. By default, the snapshot is recovered from the local cluster. d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the namespace to the original location.
8. From the **PVC Selection** option, select the PVCs of the namespace you want to recover.
9. Optional. Edit the default settings of the following **Recovery Options**:
    - **Rename**: Add Prefix and/or Suffix strings to the names of the new namespaces created by this task. For example, you could specify Prefix Test_ and Suffix _QA to the namespace's name WinOS and when the namespace is recovered, the new namespace that is created is called Test_WinOS_QA.
        - **Rename Namespaces**: Allows you to rename individual objects being recovered. You can specify a new name for each object by entering the updated name next to the corresponding object in the list. If a new name is not specified for an object, the name is determined by the configured

prefix and suffix applied to the original object name.

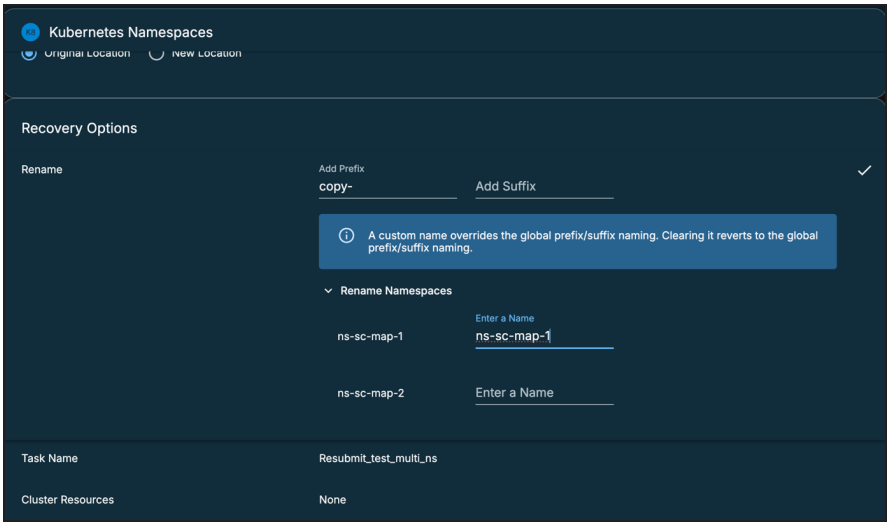

> **Note:** If you recover a namespace using the Rename Namespace and delete overwrites option, the original protection job becomes invalid. You must re-protect the namespace by creating a new protection job or adding it back to an existing protection group.

Task Name: Change the default name of the recovery task.

- **Namespace Resources**: Allows you to select the namespace resources you want to recover. Enable the following options to customize these settings: Click the edit icon on **Namespace Resources**.
  - **Resources**: In the Resources you can select the following options:
    a. **Persistent Volume Claim (PVC) Inclusion/Exclusion**: Enable this option to choose which PVCs to include or exclude in the recovery task.

    The PVCs that were part of the backup can be restored.

    You also have an option to recover only PVCs and their dependent resources. The dependent resources include: ConfigMaps and PersistentVolumes.

    - **Include**: Use this option to include PVCs during restore. For each selected PVC, you can restore:
      - Full data and metadata, or
      - Metadata only by enabling the **Only Metadata** toggle

The **Only Metadata** toggle determines how a PVC is restored.

> **Note:** If a PVC was backed up with metadata only, it can be restored only as metadata. In this case, the **Only Metadata** toggle is locked and cannot be changed.

Example

PVC1: Backed up with metadata only - Only metadata is restored (toggle locked).

PVC2: Backed up with data and metadata – You can choose to recover only metadata or both data and metadata.

If you do not enable the PVC Inclusion/Exclusion option or do not select specific PVCs:

- Each PVC is restored in the same mode in which it was backed up.

Example

- PVC1: Restored with metadata only.
- PVC2: Restored with data and metadata.

- **Exclude**: Use this option to exclude selected PVCs from the restore.

  For each selected PVC:

  - **Only Metadata** enabled: Only metadata is restored (data is excluded).
  - **Only Metadata** disabled: PVC is not restored.

  Example

  - PVC1: Only Metadata enabled - Only metadata is restored.
  - PVC3: Only Metadata disabled - PVC is not restored.

  > **Note:** Each PVC is restored in the same mode in which it was backed up, unless explicitly modified using the inclusion or exclusion settings.

PVCs that are not selected in the exclude list are restored based on how they were backed up.

- If a PVC was backed up with metadata only, it is restored with metadata only.
- If a PVC was backed up with data and metadata, it is restored with data and metadata.

b. **Resource Inclusion/Exclusion**: Enable this toggle to include or exclude resource classes, and then select the protected resource from the drop-down.

- If **Only Metadata** is disabled: The PVC is completely excluded (no data or metadata is backed up).
- If **Only Metadata** is enabled: Only the metadata is backed up, and the data is excluded.

- **Storage Class**: Enable the following options to customize these settings:

a. **Unbind PVCs from original PV mapping**: Before using this recovery option, you must specify the location of the Cohesity DataProtect Plugin image in the **Cohesity DataProtect Plugin Image Path** field during registration of a source.

When enabled, the DataProtect Plugin Image clears certain fields from all PVCs within the namespace during the recovery process. This action removes references to the original PersistentVolume (PV), allowing the recovered PVCs to dynamically provision a new PV or bind to any available PV.

This option is useful if the snapshot contains PVCs with a Retain reclaimPolicy.

b. To map recovered PVCs to a different storage class:

i. Click **+ Add** to create a new mapping.
Each mapping includes two drop-down lists:

- Left drop-down: Select a storage class from the source (backed-up) Kubernetes cluster.
- Right drop-down: Select a storage class from the target Kubernetes cluster.

This mapping ensures that any PVCs using the selected source storage class are recreated using the selected target storage class during recovery.

ii. To create additional mappings, click **+ Add**.

10. Click **Recovery**.

If a namespace with the same name already exists in the target cluster, the **Overwrite options** dialog appears. You can select the namespaces to be deleted and recreated. Any namespaces that are not selected are overwritten.

To prevent accidental data loss, you must explicitly confirm this action by entering **YES**. If the operation proceeds, the selected action is recorded in the audit logs.

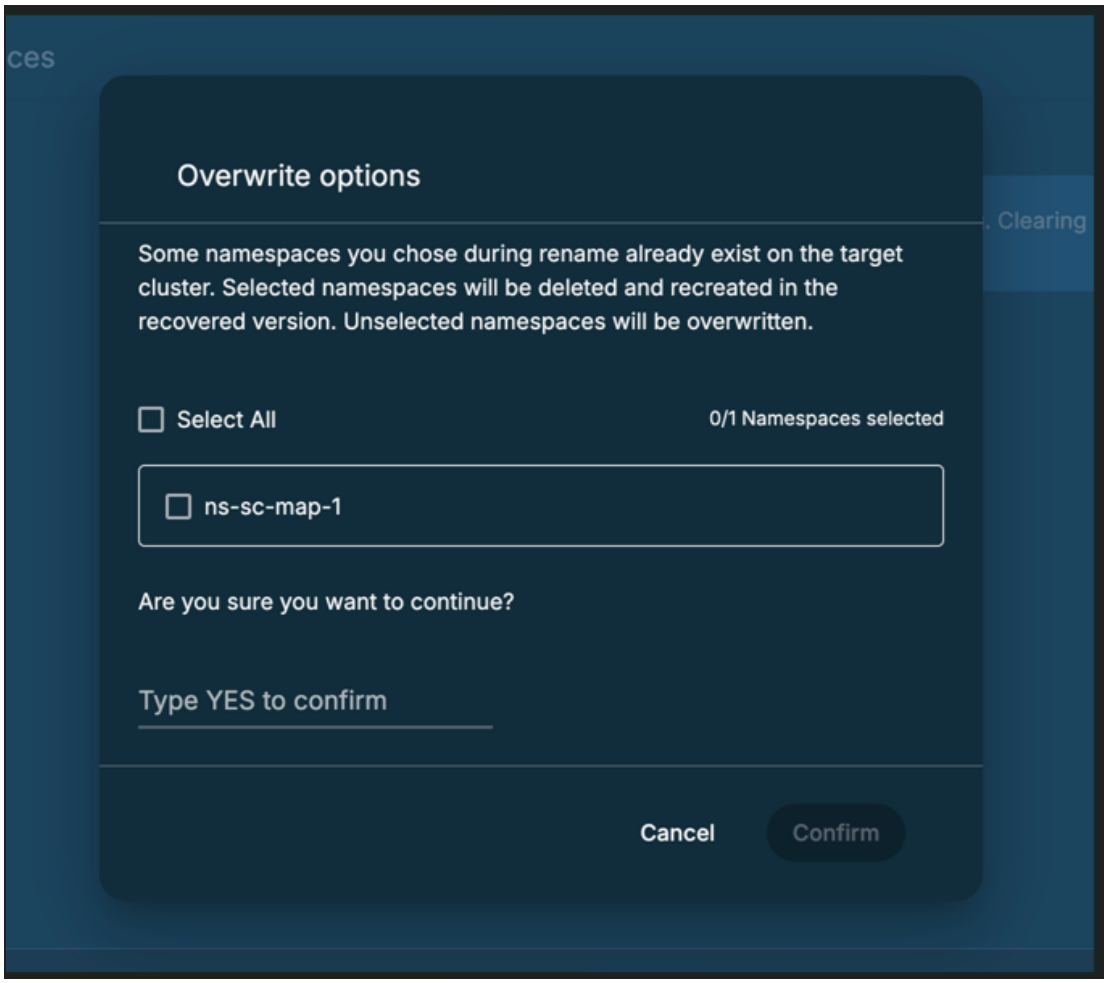

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recovering Namespaces to an Alternate Location

You can recover the namespaces of the Kubernetes to an alternate Kubernetes cluster that is registered with the Cohesity cluster. If you are recovering the Kubernetes namespace to an alternate location, then the entire namespace is restored. However, if you are restoring to this location again, then Cohesity will not restore the resources, PVCs, or the metadata of the Kubernetes namespace that already exist in the location.

To restore the namespaces of a Kuberneters cluster to the alternate location:

1. In the Cohesity cluster UI, navigate to **Data Protection > Recoveries**.
2. Click **Recoveries** on the top-right of the page and then select **Kubernetes Cluster** as the recovery source.

   The Kubernetes Namespaces page appears.
3. Search for a namespace or Protection Group that contains the snapshots to recover from by entering characters of the namespace name, Protection Group name, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, Storage Domain, or date.
4. From the list of search results, select one or more objects (namespaces) or Protection Groups that contain the snapshots to recover.

   > **Note:**
   > - If you have selected a Protection Group, then all the namespaces backed up in the latest protection run of this Protection Group will be selected for recovery. The option to select these namespaces individually will be greyed out in the search result list.
   > - You can only select Protection Groups and namespaces with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a namespace or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, perform the following steps depending on the object selected:

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Protection Group	a. From the selected lists, mouse over the snapshot, and then click the **edit** icon. b. From the **Protection Run** drop-down menu, select the snapshot from which you want to recover the namespaces. By default the snapshot is recovered from the local cluster. c. Optionally, you can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6.
Namespace	a. From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <namespace_name> page appears. By default, the page displays the snapshots available in the last one month. b. Optionally, you can also select a snapshot from a different date range or by selecting the date range from the **Date** drop-down list. Once you select the date range and then click **Apply**, the Edit recovery point page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. By default, the snapshot is recovered from the local cluster. d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** to recover the namespace to the new location.
8. From the **PVC Selection** option, select the PVCs of the namespace you want to recover.
9. Optional. Edit the default settings of the following **Recovery Options**:
    - **Rename**: Add Prefix and/or Suffix strings to the names of the new namespaces created by this task. For example, you could specify Prefix Test_ and Suffix _QA

to the namespace's name WinOS and when the namespace is recovered, the new namespace that is created is called Test_WinOS_QA.

- **Rename Namespaces**: Allows you to rename individual objects being recovered. You can specify a new name for each object by entering the updated name next to the corresponding object in the list. If a new name is not specified for an object, the name is determined by the configured prefix and suffix applied to the original object name.

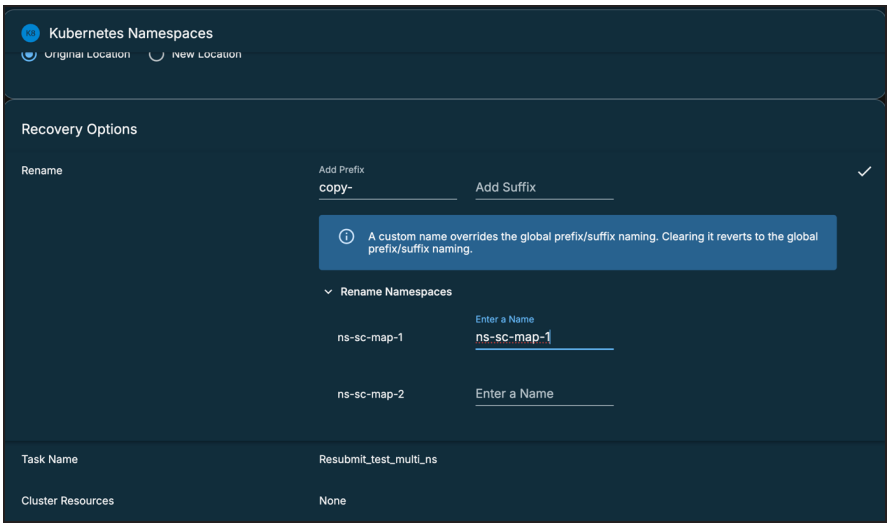

> **Note:** If you recover a namespace using the Rename Namespace and delete overwrites option, the original protection job becomes invalid. You must re-protect the namespace by creating a new protection job or adding it back to an existing protection group.

Task Name: Change the default name of the recovery task.

- **Namespace Resources**: Allows you to select the namespace resources you want to recover. Enable the following options to customize these settings: Click the edit icon on **Namespace Resources**.
  - **Resources**: In the Resources you can select the following options:
    a. **Persistent Volume Claim (PVC) Inclusion/Exclusion**: Enable this option to choose which PVCs to include or exclude in the recovery task.

       The PVCs that were part of the backup can be restored.

       You also have an option to recover only PVCs and their dependent resources. The dependent resources include: ConfigMaps and PersistentVolumes.

- **Include**: Use this option to include PVCs during restore. For each selected PVC, you can restore:
  - Full data and metadata, or
  - Metadata only by enabling the **Only Metadata** toggle

  The **Only Metadata** toggle determines how a PVC is restored.

  > **Note:** If a PVC was backed up with metadata only, it can be restored only as metadata. In this case, the **Only Metadata** toggle is locked and cannot be changed.

  Example

  PVC1: Backed up with metadata only - Only metadata is restored (toggle locked).

  PVC2: Backed up with data and metadata – You can choose to recover only metadata or both data and metadata.

  If you do not enable the PVC Inclusion/Exclusion option or do not select specific PVCs:

  - Each PVC is restored in the same mode in which it was backed up.

  Example

  - PVC1: Restored with metadata only.
  - PVC2: Restored with data and metadata.
- **Exclude**: Use this option to exclude selected PVCs from the restore.

  For each selected PVC:

  - **Only Metadata** enabled: Only metadata is restored (data is excluded).
  - **Only Metadata** disabled: PVC is not restored.

  Example

  - PVC1: Only Metadata enabled - Only metadata is restored.
  - PVC3: Only Metadata disabled - PVC is not restored.

> **Note:** Each PVC is restored in the same mode in which it was backed up, unless explicitly modified using the inclusion or exclusion settings.

PVCs that are not selected in the exclude list are restored based on how they were backed up.

- If a PVC was backed up with metadata only, it is restored with metadata only.
- If a PVC was backed up with data and metadata, it is restored with data and metadata.

b. **Resource Inclusion/Exclusion**: Enable this toggle to include or exclude resource classes, and then select the protected resource from the drop-down.

- If **Only Metadata** is disabled: The PVC is completely excluded (no data or metadata is backed up).
- If **Only Metadata** is enabled: Only the metadata is backed up, and the data is excluded.

- **Storage Class**: Enable the following options to customize these settings:

a. **Unbind PVCs from original PV mapping**: Before using this recovery option, you must specify the location of the Cohesity DataProtect Plugin image in the **Cohesity DataProtect Plugin Image Path** field during registration of a source.

When enabled, the DataProtect Plugin Image clears certain fields from all PVCs within the namespace during the recovery process. This action removes references to the original PersistentVolume (PV), allowing the recovered PVCs to dynamically provision a new PV or bind to any available PV.

This option is useful if the snapshot contains PVCs with a Retain reclaimPolicy.

b. To map recovered PVCs to a different storage class:

i. Click **+ Add** to create a new mapping.
Each mapping includes two drop-down lists:

- Left drop-down: Select a storage class from the source (backed-up) Kubernetes cluster.

- Right drop-down: Select a storage class from the target Kubernetes cluster.

This mapping ensures that any PVCs using the selected source storage class are recreated using the selected target storage class during recovery.

ii. To create additional mappings, click **+ Add**.

10. Click **Recovery**.

If a namespace with the same name already exists in the target cluster, the **Overwrite options** dialog appears. You can select the namespaces to be deleted and recreated. Any namespaces that are not selected are overwritten.

To prevent accidental data loss, you must explicitly confirm this action by entering **YES**. If the operation proceeds, the selected action is recorded in the audit logs.

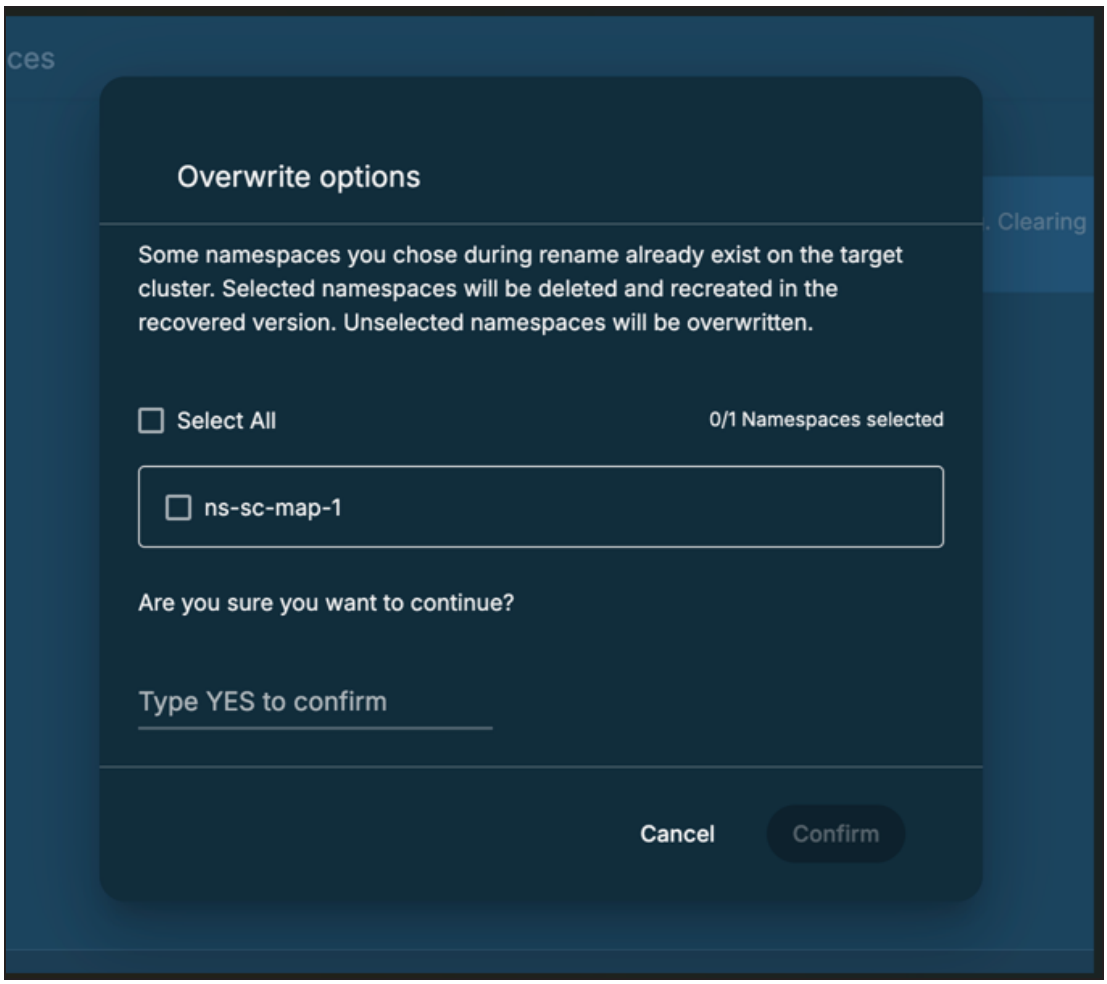

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Kubernetes Namespaces

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

The temporary pods created during recovery by Cohesity will have their priority class set to the specified value for the source. Additionally, the labels and annotations will be applied to the temporary pods.

Granular Recovery of Files and Folders

Cohesity supports Kubernetes clusters for granular recovery of files and folders within Persistent Volume Claims (PVCs). It allows for the selective recovery of files and folders from backed-up data volumes instead of restoring the entire volume in case of data corruption. This capability enables users to recover specific files or related files and folders efficiently.

While you register a Kubernetes source on the Cohesity cluster and create a Protection Group, you can perform the backup of namespaces. This allows you to search and browse the backed-up PVC folders and files you want to recover by providing their names. Each file and folder is indexed within volumes of type PersistentVolumeClaims in Kubernetes namespaces from snapshots taken during Kubernetes backups. Once indexed, the files and folders are available for searching and browsing.

You have the option to either download the files or initiate the recovery of files and folders. You can specify the download path if you choose to download the files. If you choose the restore option, you can select one or more files and folders from a single object. The recovered files and folders can be restored to their original location (such as the source, namespace, PVC, and file path) or an alternate location wherein the target source must be registered.

You can specify the recovery path and choose to preserve the original permissions and attributes of the recovered files and folders at the time of the backup. Additionally, you can overwrite files and folders or download them from selected snapshots created by a Cohesity Protection Group.

> **Note:**
> - For archived items, only downloading is supported.
> - If restoring to a different location, you must turn on the PVC inclusion/exclusion feature and enable the corresponding magneto flag. This is necessary to list the PVCs available for selection during recovery.

Considerations

- Granular recovery from archival is not supported. Only download is supported.
- If recovering to an alternate location, ensure that the file system of the target supports the recovery of files.
- Granular recovery is only supported for volumes of type PersistentVolumeClaims.
- Recovery of files and folders to an alternate location is restricted to already registered Kubernetes sources.
- Granular recovery does not support file and folder recovery involving multiple namespaces.
- Granular file recovery from the archive to the original sources is not supported. Only file downloads are allowed.
- Attempt to resubmit a failed Kubernetes restore job may fail with the following error message:

  "Error while calling FileStat for /cohesity-data/cohesity-data/pv-5tb: [kNotFound]: Bridge error: Directory not present."

  The issue occurs due to a known UI bug, which adds an extra "cohesity-data" to the file path with each resubmission. A fix for this issue is expected in a future release.

For more information about the recovery options:

- Recover Files or Folders to the Original Location
- Recover Files or Folders to a New Location within the original source

Recover Files or Folders to the Original Location

You can recover the files and folders within the PVCs to the original location. To restore the files and folders to the same location:

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** at the top right of the page and then select **Kubernetes Cluster** > **Files or Folders** as the recover source.

   The Files recovery page appears.
3. Select any one of the following options to search for the file or folder:
   a. **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   b. **Browse**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.
4. Search for files and folders that contain the snapshots to recover from by entering characters of the files or folder name or by specifying the wildcard character *.

5. From the search results list, select one or more objects (files or folders or Protection Groups containing the snapshots to recover).
6. You have the option to select either **Download** or **Recover files**. To download, select the recovery point by choosing **Select Recovery Point** and then click on the **Download** option to download files and folders. To recover, proceed to the next step.
7. Click **Next: Recover Options**.
8. Select **Original Server**, to recover the files or folders to the original location.
9. Optional. Edit the default settings of the following **Recovery Options**: You can select any one of the Recovery options:
    a. **Overwrite Existing File/Folder**: Allows to overwrite the existing files and folders with new ones during recovery.
    b. **Preserve File/Folder Attributes**: Allows to preserve the attributes of the existing files and folders during recovery.
    c. **Continue on Error**: Allows to continue recovery even if one of the objects encounters an error
    d. **Cluster Interface**: Is auto-selected,
    e. **Task Name**: Change the default name of the recovery task.
10. Click **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Files or Folders to a New Location

To recover files and folders to a new location:

> **Note:** It is recommended that the file system at the new location be the same as the source.

1. Navigate to **Data Protection** > **Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Kubernetes Cluster** > **Files or Folders**.

   The FIles recovery page appears.
3. Select any one of the following options to search for the file or folder:
    a. **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
    b. **Browse**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

4. Search for files and folders that contain the snapshots to recover from by entering characters of the files or folder name or by specifying the wildcard character *.
5. From the search results list, select one or more objects (files or folders or Protection Groups containing the snapshots to recover).
6. You have the option to select either **Download** or **Recover files**. To download, select the recovery point by choosing **Select Recovery Point** and then click on the **Download** option to download files and folders. To recover, proceed to the next step.
7. Click **Next: Recover Options**.
8. Select **New Server** , to recover the files or folders to the new location.
9. Select any of the registered source, namespace, persistent volume claim, and the path to save the recovered files or folders.
10. Optional. Edit the default settings of the following Recovery Options: You can select any one of the Recovery options:
    a. **Overwrite Existing File/Folder**: Allows to overwrite the existing files and folders with new ones during recovery.
    b. **Preserve File/Folder Attributes**: Allows to preserve the attributes of the files from the snapshot, regardless of their existence.
    c. **Continue on Error**: Allows to continue recovery even if one of the objects encounters an error.
    d. **Cluster Interface**: Is auto-selected,
    e. **Task Name**: Change the default name of the recovery task.
11. Click **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Red Hat OpenShift Virtualization

Red Hat OpenShift Virtualization is a feature of Red Hat OpenShift that allows you to manage and run VMs and container workloads together. It is based on the open-source project KubeVirt and uses a Kernel-based VM (KVM) within Kubernetes.

OpenShift API for Data Protection (OADP)

Cohesity now supports the OpenShift API for Data Protection (OADP) operator, a Red Hat solution that simplifies the registration and maintenance of Velero and Plugin versions. OADP is based on the open-source tool Velero, which allows for effective metadata backup.

Plan and Prepare for Protecting Red Hat OpenShift Cluster

Before you register your Red Hat OpenShift cluster on the Cohesity platform, ensure that:

- Cohesity cluster supports the Red Hat OpenShift cluster version you plan to protect. For more information, see Supported Versions.
- The firewall ports required for the communication between the Cohesity cluster and Red Hat OpenShift cluster are open. For more information, see Network Firewall Considerations.
- You have the minimum permission required to register and manage your Red Hat OpenShift cluster. For more information, see Minimum Permissions.
- You meet the prerequisites to register and manage your Red Hat OpenShift cluster. For more information, see Prerequisites.
- You have read the considerations for the backup and recovery of the Red Hat OpenShift cluster. For more information, see Considerations.

Supported Versions

Cohesity cluster supports data protection and recovery of the Red Hat OpenShift cluster. For supported versions of the Red Hat OpenShift clusters, see Supported Versions.

Network Firewall Considerations

If there is a firewall between the Cohesity cluster and the Red Hat OpenShift cluster, the following ports need to be opened.

Cohesity cluster	Type:Port	Direction	Red Hat OpenShift cluster	Type:Port	Usage Notes
Cohesity cluster nodes	TCP:	Outgoing	Red Hat OpenShift cluster api-address	TCP: <Red Hat OpenShift API server port>	Red Hat OpenShift cluster and Velero API Access. For more details on Velero support, see the Considerations section.
Cohesity cluster nodes	TCP:	Outgoing	Red Hat OpenShift cluster worker nodes	TCP:30000-32767[1] TCP:32000-(32000+#<worker nodes>)[2]	Cohesity cluster to Cohesity Datamover communication
Cohesity cluster VIPs	TCP:3000	Incoming	Red Hat OpenShift cluster	TCP	Velero S3 communication for metadata backups and restores

1 TCP:30000-32767 are supported only on releases 6.6.0a to 6.6.0d_u1.

2 TCP:32000-(32000+#<worker nodes>) are supported only on releases 6.6.0d_u2 and later.

Minimum Permissions

To register the Red Hat OpenShift cluster with the Cohesity cluster, you must provide the Bearer token ID of the service account. Ensure that the service account has the cluster-admin role on the Red Hat OpenShift cluster.

Prerequisites

Ensure that you meet the following prerequisites to register and manage your Red Hat OpenShift cluster:

- **Valid Certificate with Subject Alternative Name (SAN) field** —On the Cohesity cluster, install a valid certificate with the SAN field. The SAN field is used to specify additional host names (such as sites, IP addresses, and common names) to be protected by a single SSL certificate, such as a Multi-Domain or Extend Validation Multi-Domain Certificate. For more information, see Update SSL Certificates Using Cohesity CA or Customer CA certificate for Agent to Cluster Communication.
- **Create a Cohesity service account**—For registering a Red Hat OpenShift cluster with the Cohesity cluster, you can either use a default service account with the cluster-admin role or a dedicated service account with the cluster-admin role account. Cohesity recommends you to create a dedicated account in the default namespace.

  For example, to create a 'cohesity' service account, use the following command:

  ```
  kubectl create serviceaccount cohesity -n default
  ```

  Assign the cluster-admin role to the 'cohesity' service account using the following command:

  ```
  kubectl create clusterrolebinding cohesity-admin --clusterrole=cluster-admin --serviceaccount=default:cohesity
  ```

  The above example can be run from any client or workstation from which `kubectl` is typically run for a given Red Hat OpenShift cluster that is registered and protected by the Cohesity cluster.
- For current version of Red Hat OpenShift, run the following command to generate the Bearer Token ID or secret token for your service account.

  ```
  kubectl apply -f - <<EOF
  apiVersion: v1
  kind: Secret
  metadata:
    name: cohesity-token
    annotations:
      kubernetes.io/service-account.name: cohesity
  ```

```
type: kubernetes.io/service-account-token
EOF
```

Sample Ouput

```
secret/cohesity-token created
```

- **Extract the Bearer token ID**—To register the Red Hat OpenShift cluster with Cohesity cluster, you must have the Bearer token ID of the Kubernetes service account. Extract the Bearer token of the service account that has admin privileges on the Red Hat OpenShift cluster.

  Extract the Bearer token using the following command:

```
kubectl describe secrets -n default cohesity-token | grep token: | awk
'{print $2}' | head -1

eyJhbGciOiJSUzI1NiIsImtpZCI6IjRtZzBkX0VRLXMwV05veFdvcjR3eGw1N2c4MUk0RX
lhWFBPdXBHNUVoTlkifQ.eyJpc3MiOiJrdWJqcm5ldGVzL3NlcnZpY2VhY2NvdW50Iiwia
3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9uYW1lc3BhY2UiOiJkZWZhdWx0Iiwia3V
iZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9zZWNyZXQubmFtZSI6ImNvaGVzaXR5LXRva
2VuL2Y0ZG02Iiwia3ViZXJuZXRlcy5pbF9zZXJ2aWNlYWNjb3VudC9zZXJ2aWNlLWFjY29
1bnQubmFtZSI6ImNvaGVzaXR5Iiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9zZ
XJ2aWNlLWFjY291bnQudWlkIjoiZDg4NmE5ZWMtMzQxYi00ZTM3LTg1MjktMzYxNGJiY2U
0OTAxIiwic3ViIjoic3lzdGVtOnNlcnZpY2VhY2NvdW50OmRlZmF1bHQ6Y29oZXNpdHeif
Q.Of4nzdrTDsbxzHWbzskwk6_
mdwSvduSLPhSivW3MdHJNcGMa6nGmn5iKUQqDG4oe1tu71w0r9rT1Cp8sGHOLxpjGEhaKS
f7gIGX5fUbw2jY6m7gL1U-4zj5FrglJ9XdE1yE70qNE3nFL1QZexYjAW80_
rjMBv0lW4unq1o2IO21e5PC31SIGMfs1rlSvToy4_
vJEVFzX9E0e4ha1RSsEIYFjdwt3sgG6y-
PGoR4UBgnNnBaMpmy6NXYToaChzteJw93HQyyRVGBFNYgNTinqc5edu_
zyMdAl6Chba6q10ffjC_zjhU_Lm7wp31-BTMelg-1W_5b3uni4pCGD8EPGPw
```

- **Obtain the Cohesity Datamover Docker Image**—To back up the PVCs of the Kubernetes namespaces, you must push the Cohesity Datamover image to an image registry accessible by your Red Hat OpenShift cluster.

  To install the Cohesity Datamover image:

  1. Obtain the Datamover image that matches closely with your Cohesity cluster version from the Cohesity Download page.
  2. To push the Datamover image to your Docker registry:
     1. Run the following command to log into your Docker image registry:

```
docker login <your image registry server>.fqdn
```

If prompted, enter the credentials that have the permissions to push the Cohesity Datamover image to your image registry. Skip this login command if your image registry does not require authentication to push images.

You can push the Cohesity Datamover image to your container image registry using either Docker or Skopeo:

Option 1 - Using Docker

2. Load the Datamover image:

```
docker load --input cohesity-datamover-<your Datamover version>.tar
```

Replace <your Datamover version> in the preceding command with the version in your Datamover image name.

Example:

```
docker load --input cohesity-datamover-6.6.0d_u6.tar
```

```
docker load --input cohesity-datamover-6.8.1_u1.tar
```

3. Run the following command, and from the command result, make a note of the Image ID of the Datamover image:

```
docker images
```

4. Tag the Datamover image with the Image ID you made a note of in step c:

```
docker image tag <ID of your image> <your image registry server>.fqdn/path/cohesity-datamover:<your Datamover version>
```

5. Push the Datamover image to your Docker registry:

```
docker push <your image registry server>.fqdn/path/cohesity-datamover:<your Datamover version>
```

> **Note:** The preceding example assumes you use Docker to load, tag, push, and pull the Docker images. If you use Podman, replace `docker` with `podman`.

Option 2 - Using Skopeo

You can push the Datamover image tarball directly to the image registry without loading it locally.

Run the following command:

```
skopeo copy docker-archive:cohesity-datamover-<your Datamover
version>.tar docker://<your image registry
server>.fqdn/path/cohesity-datamover:<your Datamover version>
```

- **Obtain the Cohesity DataProtect Plugin Image**—To upgrade the Cohesity DataProtect Plugin for a Red Hat OpenShift source, you must upload the Cohesity DataProtect Plugin container image to an image registry that is accessible by your Red Hat OpenShift cluster.

  > **Note:** Docker commands must be executed from an amd64 machine.

  To upload the Cohesity DataProtect Plugin image:

  1. **Download the DataProtect Plugin image** - Obtain the DataProtect Plugin tarball that closely matches your Cohesity cluster version from the Cohesity Download page.
  2. **Load the plugin image** - Load the downloaded tarball as a container image to their private image registry using the following command:

     ```
     docker load -i cohesity-dataprotect-plugin-<version>.tar
     ```

  3. **Verify the loaded image** - Run the following command to check the loaded image and note the Image ID:

     ```
     docker images
     ```

  4. **Tag and push the image to your registry** - Tag the image using your private image registry location using the following commands:

     ```
     docker tag <image-id> <image-registry>/cohesity-dataprotect-
     plugin:<tag>
     ```

     ```
     docker push <image-registry>/cohesity-dataprotect-plugin:<tag>
     ```

  To prepare the source Cohesity cluster:

  Before upgrading the DataProtect Plugin, ensure that the Red Hat OpenShift source is in a clean state.

  1. Stop all running Red Hat OpenShift protection jobs and verify that no protection runs are currently active.
  2. Delete the Cohesity namespace from the Red Hat OpenShift cluster using the following command:

     ```
     oc delete ns <cohesity-namespace>
     ```

  To update the Red Hat OpenShift source configuration:

1. Navigate to **Data Protection > Sources**.
2. From the **Actions** (⋮) menu next to the Red Hat OpenShift source, select **Edit**.
3. In the **Register Kubernetes Source** page, specify the location of the new DataProtect Plugin image.
4. Click **Save** and wait for the registration process to complete.

   The registration process redeploys the Cohesity Data Mover with the updated DataProtect Plugin.

Considerations

- Cohesity does not support incremental backup.
- Red Hat OpenShift Virtualization is supported only in on-premise environments.
- While protecting VMs with Block mode, ensure that you are using the latest version of Datamover.
- Cohesity does not support the restoration of VMs with partially selected PVCs. As a result, the VMs will not get provisioned. Also, configuration change is required to remove the excluded disk configurations.
- PVCs not attached to any pods are protected if a CSI snapshot is available for the Volume.
- You must toggle ON the CSI snapshot for crash-consistent backups. If a CSI snapshot is not enabled on the Protection Group or fails during its creation, the backup may still proceed for the VM. However, this backup may be inconsistent.
- Cohesity does not support granular recovery of files and folders within VMs.
- Cohesity does not support quiescing VMs during backups.
- If SELinux is enabled on the VM, correct permissions should be set on qemu-agent to allow it to run guest-fsfreeze commands.
- Cohesity does not support recovery from backups if the storage class is deleted after the backup is made from the source and that same storage class is not present on the alternate cluster.
- If the Machine type is not supported on the alternate cluster then recovery is successful but the VM quits unexpectedly. Ensure that the required Machine type is available on the alternate location cluster.
- In scenarios where the Red Hat OpenShift cluster lacks external IP support, a configuration change may be needed to support backups and recoveries. Contact your Cohesity account team to enable the feature.
- During the backup process, metadata for all resources in the namespace is backed up, but only the volumes of the selected VMs are included in the backup.
- During the recovery process, if all VMs in a namespace are being recovered, the entire metadata for that namespace is also recovered along with the VMs and their

associated volumes.

- If VMs are selectively recovered, only the metadata of the resources connected to those selected or backed up VMs, along with their volumes, will be restored.
- If the version of the Red Hat OpenShift cluster is lower than 4.17, the deployment of the OpenShift API for Data Protection (OADP) operator is not supported.
- During the deployment of the operator, the subscription channel is set to stable-1.4, which automatically manages any minor updates to the operator. If you have deployed another OADP operator with a different stable channel, then the Cohesity operator deployment is not supported.
- If the Velero CRDs are already present, you cannot deploy the OADP operator by default. You must delete the CRDs and click refresh on the OpenShift source in the Cohesity UI to deploy the operator.
- After restoring a VM with a hotplug disk, the disk appears in the VM configuration but is not visible within the VM. As a workaround, either restart the VM or remove the older hotplug pod that is currently running to detect the hotplug disk within the VM.
- Backups fail if the namespace contains a VM that is in a paused state. As a workaround, stop the paused VM and reinitiate the backup. As a workaround, for assured data consistency during backups, VMs should be powered down before initiating the backup.
- Backing up a VM that uses LVM or has a file system spanning multiple disks may result in data corruption. The issue occurs when the snapshots of the underlying volumes or PVCs are not captured at the same point in time, resulting in inconsistencies
- For IPv6 support, the Cohesity cluster must be configured with a preferred IPv6 address.
- Instant recovery feature is supported starting from OpenShift 4.18.
- VM recovery using the override option is not supported.
- When creating target data volumes for VM migration, the volume mode (Filesystem or Block) is not explicitly specified.
- A Velero backup may fail if the PVC associated with a VM is deleted before the backup is triggered. As a result, the Velero fails to backup deleted PVC and marks the metadata backup as failed:

  “Backup job failed by Velero.”
- After upgrading the Cohesity cluster from version 7.2.2_u2 to 7.3.1, if an incremental backup is run using an existing Protection Group to generate a new snapshot, instant recovery of a VM from that snapshot may fail with the following error:

  “Failed to clone the namespace backup view.”

Instant Recovery is supported starting from version 7.3.1. For existing Protection Groups, a forced full backup is required. Run a full backup of the affected Protection Group to ensure that subsequent backups create separate views.

- If a VM's Network Attachment Definitions (NADs) are located in a different namespace than the VM being backed up, those NADs are not included in the backup for that VM.
- NFS-mounted file systems (and other file systems that do not support extended attributes) such as security.selinux are not supported as target locations for file-level recovery, and restores may fail.
- On OpenShift 4.19 for Red Hat OpenShift Virtualization, File-Level Recovery (FLR) may fail with "Hotplug volumes did not reach Ready state within 300 seconds" when a block-mode PVC is already hotplugged, causing subsequent hotplug operations to fail unless you manually delete the existing hotplug pod.
- **For KubeVirt**

  Below are the considerations for KubeVirt. For more information, see limitations:

  - Migration of hot-plug disks is not supported.
  - Migration of shareable disks is not supported, as data consistency cannot be guaranteed with multiple writers.
  - Live migration is not supported when using virtiofs because it does not currently provide live migration capabilities.
  - Original disks may support the SCSI protocol, but the destination PVC storage class might not support it.
  - KubeVirt currently supports live migration only between separate nodes.

> **Note:** After upgrading to version 7.3.1, ensure you run a full backup of all VMs in the existing protection jobs.

Register and Manage Red Hat OpenShift Cluster

To protect the namespaces of the Red Hat OpenShift cluster, you must first register your Red Hat OpenShift cluster as a source on the Cohesity cluster. After registering your source, you can also perform the following actions:

- Update the Red Hat OpenShift cluster source configuration.
- Refresh the source details.
- Unregister the Red Hat OpenShift cluster from the Cohesity cluster.

Register Red Hat OpenShift Cluster as a Source

Before you register the Red Hat OpenShift cluster as a source on the Cohesity cluster, ensure that you:

Review the requirements, considerations, and best practices.

Perform the prerequisite tasks.

For more information, see Plan and Prepare for Protecting Red Hat OpenShift Cluster.

To register the Red Hat OpenShift cluster as a source on the Cohesity cluster:

1. Navigate to **Data Protection > Sources**.
2. Click **Register** and then select **Kubernetes Cluster**.

   The **Register Kubernetes Source** page appears.

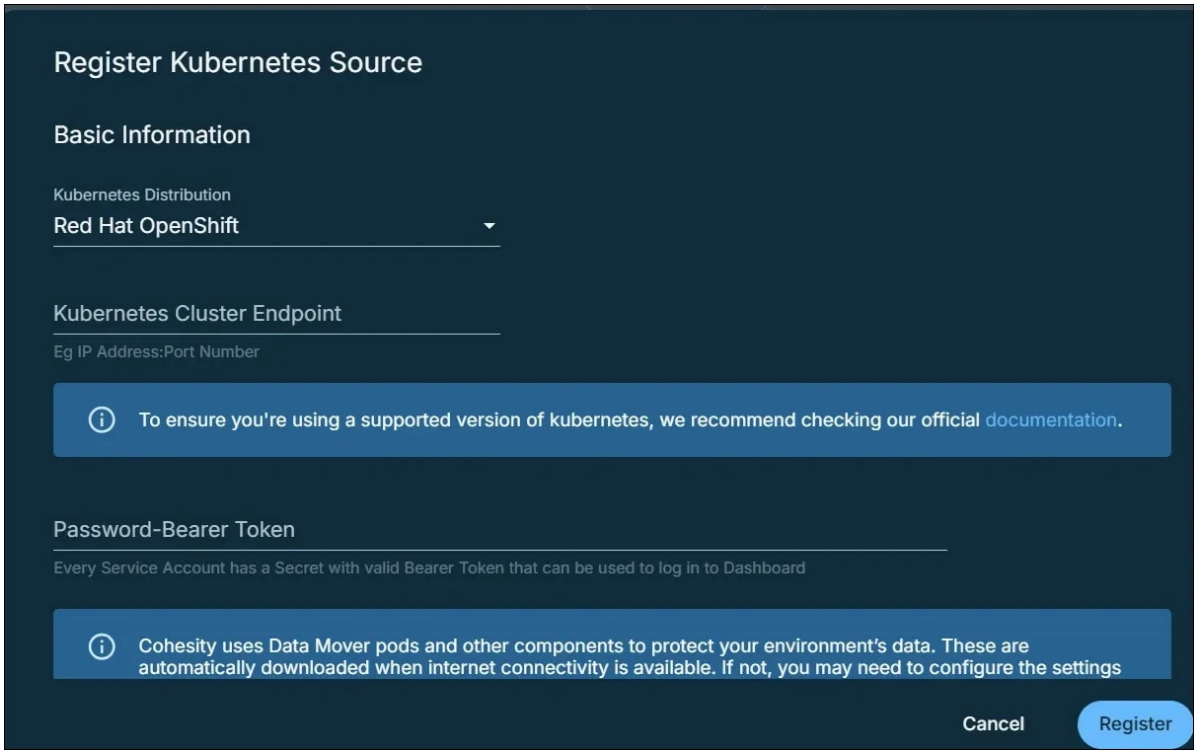

3. From the **Kubernetes Distribution** drop-down, select Red Hat OpenShift Kubernetes distributions to backup and recover.
4. In the **Kubernetes Cluster Endpoint** field, based on your requirements, enter the Kubernetes API Server endpoints in one of the following formats:

   https://<FQDN>:<port_number>

   IPv4

   Format: https://<IPv4 address>:<port_number>

   Example: https://10.14.58.2:6443IPv6

   Format: https://[<IPv6 address>]:<port_number>

   Example: https://[2400::9000:0:0:500:12]:6443

In the **Password-Bearer Token** field, enter the Bearer token ID of the Kubernetes service account. For information on obtaining the Bearer token ID, see Plan and Prepare for Protecting Red Hat OpenShift Cluster.

5. *(Optional)***Service Type**: Enable this field to route external traffic to your Red Hat OpenShift cluster node via dedicated VLANs.

   From the drop-down, select one of the following service types:

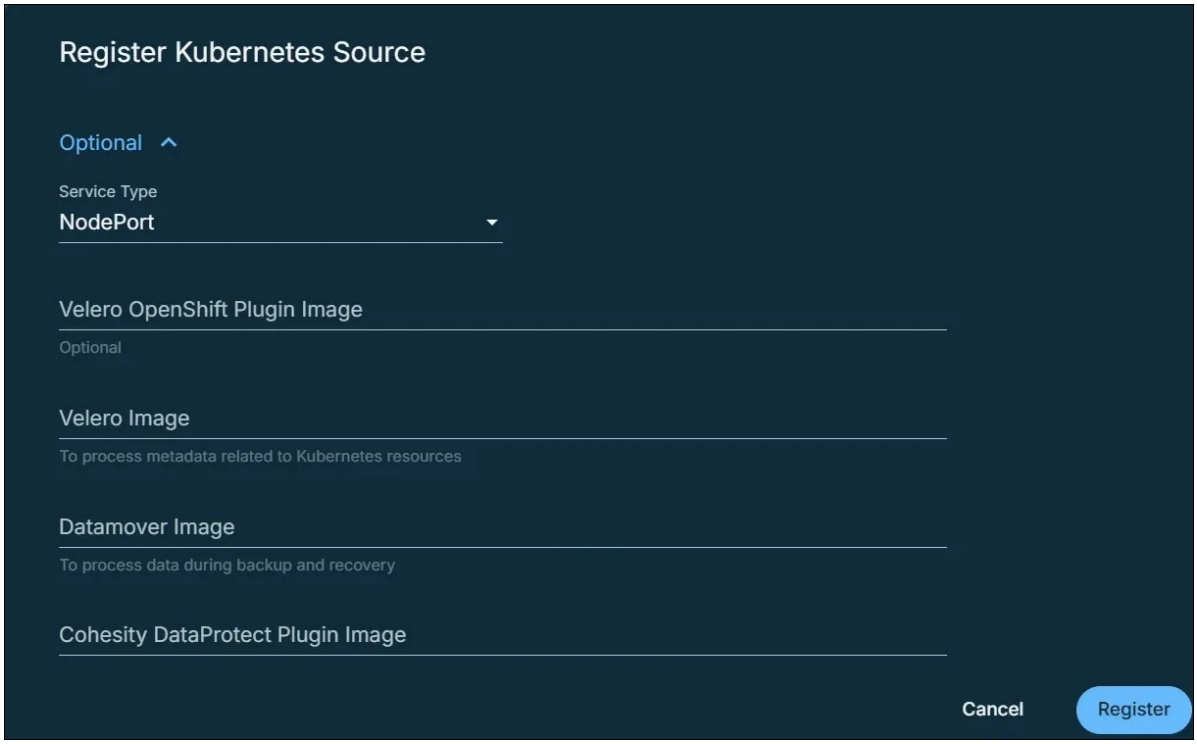

   - **NodePort** service type allows you to route external traffic through a fixed port (which you can set within the range of 30000-32767) on every Kubernetes node. In a NodePort setup, external traffic has direct access to the Kubernetes node.
   - **Load Balancer** service type allows you to route external traffic via an external Load Balancer functionality. Load Balancers monitor the availability of your Kubernetes pods and route incoming service requests to the relevant pods. In a Load Balancer setup, external traffic has access to the Kubernetes node through the Load Balancer only.

6. *(Optional)***Velero OpenShift Plugin Image**: This plugin image will be deployed by the OADP operator for Red Hat OpenShift Virtualization.
7. *(Optional)***Velero Image**: This plugin image will be deployed by the OADP operator for Red Hat OpenShift Virtualization.
8. **Datamover Image**: Enter the location where you have the Datamover image saved in your registry accessible from the OpenShift cluster. The image must be entered in

the format: `<Image_name>:<TAG>`.

9. *(Optional)***Cohesity DataProtect Plugin Image**: Enter the location where you have the Cohesity DataProtect Plugin image saved. This plugin is required when selecting the VM.
10. *(Optional)***Velero Kubevirt Plugin Image**: This plugin image will be deployed by the OADP operator for Red Hat OpenShift Virtualization.
11. *(Optional)***Velero AWS Plugin Image**: This plugin image will be deployed by the OADP operator for Red Hat OpenShift Virtualization.
12. *(Optional)***Init Container Image**: Enter the location where you have the Init Container image saved in your Docker registry. If you leave this field blank, then Cohesity will use the Datamover image as the Init Container image

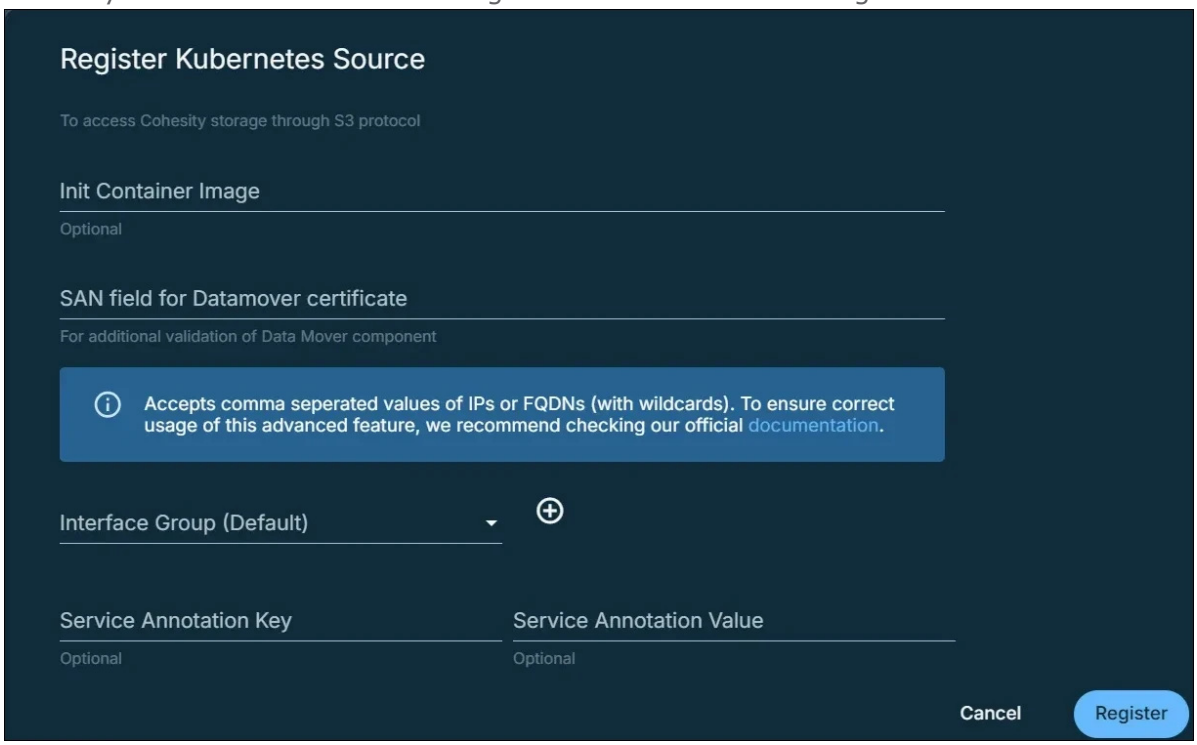

13. **Interface Group (Default)**: Displays all the Interface Groups set in your cluster. Click the **+** icon to add as many Interface Groups as needed. Once the Interface Groups are created, the traffic of the Service Type (you selected in Step 8) will only flow through the Interface Group(s), creating a separate channel for the traffic.

    Vlan Annotation Key and **Vlan Annotation Value**: Displayed only if the **Load Balancer Service Type** is selected. Enter the key-value pair to specify a VLAN that the Load Balancer service deploys to.

    Load balancer services allow you to configure a specific kind of address pool (in this case VLAN). For example, if you use the MetalLB LoadBalancer service, you need to set the **Vlan Annotation Key** to "metallb.universe.tf/address-pool" and the **Vlan**

Annotation Value should be the name of the address pool. For more information, see the documentation of your LoadBalancer service.

> **Note:** The Interface Group field is an Early Access feature. Contact your Cohesity account team to enable it.

14. **Service Annotation Key** and **Service Annotation Value** Enter the Service annotation key-value pairs. These key-value pairs are utilized to store metadata that offers information about the selected Kubernetes service that Cohesity deploys. This metadata does not affect the behavior of the Red Hat OpenShift cluster.

    > **Note:** The Service Type field is an Early Access feature. Contact your Cohesity account team to enable it.

15. Click **Register**.

Post-Registration Task

Once you register the Red Hat OpenShift cluster with the Cohesity cluster, perform the following actions:

- Run the following command to view the ports associated with each node of the Red Hat OpenShift cluster.

  ```
  kubectl get services -n cohesity-<cluster-id>
  ```

  Ensure the ports associated with the nodes are open to enable communication between the Red Hat OpenShift cluster and the Cohesity cluster.

  For information on other required ports, see Network Firewall Considerations.

- Verify all of the pods in the Red Hat OpenShift cluster are in running state by running the following command:

  ```
  kubectl get pods -n cohesity-<cluster-id>
  ```

  The namespaces are displayed on the **Sources** page of the Cohesity Dashboard only after all the pods are either in the running state or error state.

- Refresh the Red Hat OpenShift source if there is an error on the source page. For more information, see Refresh the Red Hat OpenShift Source.

Update the Red Hat OpenShift Cluster Source Configuration

You can update the Red Hat OpenShift cluster details that you had provided during the registration process with the latest Red Hat OpenShift cluster configuration.

To edit the Red Hat OpenShift cluster source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Red Hat OpenShift cluster source and select **Edit**.
3. In the Register Kubernetes Source page, update the respective configurations.
4. Click **Register**.

Refresh the Red Hat OpenShift Source

You can refresh the Red Hat OpenShift and fetch the latest changes on the Red Hat OpenShift cluster. To refresh the Red Hat OpenShift cluster source configuration:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Red Hat OpenShift cluster source and select **Refresh**.

Unregister Red Hat OpenShift Cluster

If you plan to stop taking the backup of your Red Hat OpenShift cluster, you can unregister the Red Hat OpenShift cluster source from the Cohesity cluster.

Before unregistering, ensure to delete all the Protection Groups associated with the Kubernetes source.

To unregister the Red Hat OpenShift cluster:

1. Navigate to **Data Protection > Sources**.
2. Click the actions menu (⋮) next to the Red Hat OpenShift cluster source and select **Unregister**.

Perform a Manual Cleanup

To perform a manual cleanup:

1. Stop all the Red Hat OpenShift protection jobs that are currently running and ensure there are no active protection runs on the Red Hat OpenShift source.
2. Run the following command to delete the Cohesity namespace in the Red Hat OpenShift cluster:

   ```
   - oc delete ns <name of the Cohesity namespace>
   ```

   Example

   ```
   - oc delete ns cohesity-21234567891234
   ```

3. Run the following command to delete the CRDs specific to Cohesity:

   ```
   - oc delete crds -l operators.coreos.com/redhat-oadp-operator.<Cohesity namespace>
   ```

   Example

```
- oc delete crds -l operators.coreos.com/redhat-oadp-operator.cohesity-21234567891234
```

4. Delete the operator from the cluster using the following command:

   ```
   - oc delete operator redhat-oadp-operator.<Cohesity namespace>
   ```

   Example

   ```
   - oc delete operator redhat-oadp-operator.cohesity-21234567891234
   ```

Upgrade to OpenShift API for Data Protection (OADP)

After upgrading the Cohesity cluster, make sure you follow steps to make use of the newly supported features. Additionally, update the Datamover image to the latest version to.

Upgrade When the Cohesity Namespace is Present

To upgrade from Velero deployment to OADP:

1. Stop all the Red Hat OpenShift protection jobs that are currently running and ensure there are no active protection runs on the Red Hat OpenShift source.
2. Run the following command to delete the Cohesity namespace in the Red Hat OpenShift cluster:

   ```
   - oc delete ns <name of the Cohesity namespace>
   ```

   Example

   ```
   - oc delete ns cohesity-21234567891234
   ```

3. Run the following command on any node of that Red Hat OpenShift cluster (or from a node in another Red Hat OpenShift cluster that has access to the OpenShift command-line tool) to delete the Velero CRD:

   ```
   - oc delete crds -l component=velero
   ```

4. To edit the Red Hat OpenShift cluster source configuration:
   a. Navigate to **Data Protection** > **Sources**.
   b. Click the Actions menu (⋮) next to the Red Hat OpenShift cluster source and select **Edit**.
   c. In the Register Kubernetes Source page, enter the locations of the new DataProtect Plugin and Data Mover images in the **Datamover Image** field, respectively.
   d. Click **Register**.

Backup Red Hat OpenShift Virtualization

Cohesity supports the backup of PVCs, VMs, Red Hat OpenShift cluster resources and metadata of the Kubernetes namespaces. Cohesity uses the Data Mover pod that will be deployed in the Red Hat OpenShift cluster for backing up the PVCs of the Kubernetes namespaces.

Cohesity supports the backup of PVCs with the volumeMode set to Filesystem or Block.

Cohesity uses the VM Snapshot API that enables taking a single snapshot of all volumes associated with a VM in one backup operation. PVCs that are not associated with the VM but exist within the same namespace will continue to use individual volume snapshots for backup purposes.

The supported file systems include ext3, ext4, xfs, vxfs, ReiserFS, autofs, cvfs, zfs, fuse, mfs, nfs, nfs3, cifs, gpfs, and mmfs.

Cohesity also supports the backup of PVCs with CSI drivers using CSI volume snapshotting capability.

> **Note:** Cohesity recommends that the entire namespace containing VMs should be protected, without any exclusions. Additionally, it is recommended to use CSI snapshots.

After registering the Red Hat OpenShift cluster as a source on the Cohesity cluster, you can start configuring your backups.

Protect Red Hat OpenShift Cluster

In the Cohesity cluster, Kubernetes namespaces are protected using a Protection Group and a Protection Policy. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated. You can select which standard or custom policy to use when creating a Protection Group.

To protect the namespaces of the Kubernetes containers:

1. In the Cohesity cluster UI, perform one of the following:
    - Select **Data Protection > Protection**. Click **Protect** displayed on the top-right of the page and then select Kubernetes Cluster.
    - Select **Data Protection > Sources**. Click the Actions menu (⋮) next to the Red Hat OpenShift source and then select Protect.
2. From the **Registered Source** drop-down list, select the registered Red Hat OpenShift source or click **Register Source** to register a Red Hat OpenShift source. For more information, see Register Red Hat OpenShift Cluster as a Source.

3. From the **Objects** list, select the namespaces that you plan to backup and click **Save Selection**.

   > **Note:** You can use the advanced filters that are available to search and find the namespaces you plan to protect. Also, you can click the shield icon to enable Auto Protect at the Red Hat OpenShift cluster level or based on the namespace label(s). For more information, see Enable or Disable Auto Protect.

4. In the Protection Group section, perform one of the following:
   - Select **New Group** to create a new Protection Group. In the **Name** field, enter a name from the Protection Group and go to step 5.
   - Select **Existing Group** to use the settings of an existing Protection Group. From the **Group** list, select an existing Protection Group and click **Protect**.
5. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Use an Existing Protection Policy .
6. Select the appropriate **Storage Domain**. For more information, see Manage Storage Domains.
7. In the **Start Time** field, select the time and time zone to start the protection run. This option is available only if the selected policy is set to Backup Daily. The current time is displayed by default but you can change it. Enter the hour and minutes or use the up and down arrows on your keyboard. Verify the AM or PM setting. The default time zone is the browser's time zone. You can change the time zone of the protection run by selecting a different time zone.
8. Cohesity supports protecting PVCs with the snapshot and restore feature of Container Storage Interface (CSI) drivers on CSI volumes. This option is disabled by default in your Cohesity protection job. To enable this feature, turn ON the **Leverage CSI snapshot** toggle. Once enabled, Cohesity clusters will begin utilizing snapshots captured by CSI drivers to protect the PVCs. These snapshots are crash-consistent, meaning that the PVC is in the same state as it was during the crash. If you are protecting entire namespaces, this setting must be enabled manually. When protecting individual VMs, the setting is applied automatically.
9. If you need to change any of the advanced settings on the Protection Group page, click the down arrow icon next to **Advance Settings** and click **Edit** on the right. For more information, see Advanced Settings.
10. Click **Protect**.

    Your new Protection Group is now active and running.

Advanced Settings

Advance Settings	Description
Pause Future Runs	Enable **Pause Future Runs** to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Enable **End Date** and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.

Advance Settings	Description
Priority	Select a priority for the Protection Group eProtection Groupxecution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	Enter a brief description about the Protection Group.

Advance Settings	Description
Include or Exclude PVC	Cohesity enables you to choose which PVCs to include or exclude in the Protection Group. To select the PVCs: a. Turn ON the **Enable PVC Inclusion/Exclusion** toggle. b. In the **Logical Rule** drop-down, select “Match Any of the following labels” or “All of the following labels”. c. Select **Include** or **Exclude**. d. Click **+ Add**. e. Enter the PVC’s label. The label is the unique identifier assigned to that PVC in the Red Hat OpenShift cluster. You can add more PVC until you have all the PVCs you want to be included or excluded for that Protection Group. **Note:** • PVC Inclusion or Exclusion settings of the Namespace take precedence over the settings of its Protection Group. • Do not modify these settings while protecting VMs.

Create or Use an Existing Protection Policy

A Protection Policy is a reusable set of settings. It enables you to define how and when Kubernetes namespaces are protected, replicated, and archived. A standard policy provides various settings that you can configure for a Protection Group.

Create a Protection Policy with the settings listed in the Policy Settings for Kubernetes. For detailed instructions on how to create a Protection Policy, see Create or Edit a Standard Policy.

Policy Settings for Red Hat OpenShift

The following table lists the applicable policy settings for Red Hat OpenShift:

Field	Description
DataLock	If applicable, add a DataLock for compliance and regulatory requirements, to ensure that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration. Once applied, a DataLocked Snapshot will be deleted only after its retention period expires. A DataLock prevents all users, including those who have the Data Security role in Cohesity Data Cloud, from modifying or deleting any Snapshots that were generated by the Protection Groups that use this Policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy.
Backup	Defines when backups are captured, how long they are retained and the type of backup to capture.
Periodic Full Backup	Use this option to choose when to run a periodic full backup.
Extended Retention	Add this to retain a subset of snapshots (backups) for longer than defined by the protection schedule.
Quiet Times	Quiet times define time periods when new protection runs are not started. For example, you might want to configure hourly backups that run during weekdays but not on Saturdays or Sundays. Quiet times only prevent new protection runs from starting during the specified time period. The protection runs that started before the quiet times are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet times.
Retry Options	Use this option to customize the retry settings for capturing Snapshots. By default, the Cohesity cluster attempts to capture Snapshots three times before the Protection Group fails. The default time between retries is five minutes. You can customize the number of retries and how long to wait between each attempt.
Replication	To replicate snapshots to another, click the **Replication** icon from the floating menu. The copied snapshots are stored in the storage domain of the remote cluster.
Archival	To add an archive schedule, click the **Archival** icon from the floating menu.

Use an Existing Protection Policy

You can edit an existing Protection Policy and configure it with applicable policy settings for Red Hat OpenShift cluster. In this case, all the policy settings apart from the settings applicable to Red Hat OpenShift cluster are ignored.

To use an existing Protection Policy:

1. In the Cohesity cluster, navigate to **Data Protection > Policies**.
2. On the **Policies** page, select the policy that you plan to edit.
3. On the **Policy Details** page, click **Edit** and configure the applicable policy settings for Red Hat OpenShift cluster. For information on the applicable policy settings, see Policy Settings for Red Hat OpenShift.
4. Click **Save**.

The updated policy settings are applied to the Protection Group in the next protection run.

Manage the Red Hat OpenShift Backup Configuration

In Cohesity Data Cloud, a backup configuration includes a Protection Group and a Protection Policy. You can modify the existing Protection Group and the Protection Policy based on your requirements.

> **Note:** You cannot update the registered Red Hat OpenShift cluster and the assigned Storage Domain in the Protection Group settings.

Add Namespaces to the Protection Group

You can add new namespaces to the Protection Group for backup. The newly added namespaces are backed up from the next protection run, and it will be a full backup.

To add namespaces to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for the Red Hat OpenShift cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the Kubernetes page, click the edit icon in the **Objects** section.
5. From the namespaces list, select the namespaces that you plan to add to the backup and click **Save Selection**.
6. Click **Save**.

Remove Namespaces from the Protection Group

You can remove the Kubernetes namespaces from the Protection Group that you do not want to back up in the next protection run. When you remove a namespace from the Protection Group, the incremental backup of this namespace is not performed from the future protection run. The existing backup for the namespace is preserved until the snapshot expires.

To remove namespaces from the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you had created for the Red Hat OpenShift cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the Kubernetes page, click the edit icon in the **Objects** section.
5. From the namespaces list, unselect the namespaces that you plan to remove to the backup and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect

You can enable the Auto Protect option at the:

- **Red Hat OpenShift cluster level**: Automatically protect the new namespaces that are added to the Kubernetes container.
- **Label level of the namespaces**: Automatically protect the new namespaces that are tagged with the same label you selected while enabling the auto protect.

When the Auto Protect option is enabled:

- All the namespaces that are added to the Red Hat OpenShift cluster or tagged with the same label in the future are automatically added to the Protection Group and are protected from the next protection run.
- All the namespaces that are removed from the Red Hat OpenShift cluster or untagged from the specified label in the future are automatically removed from the Protection Group and are not protected from the next protection run. The existing backup for the namespace is preserved until the snapshot expires.

When the Auto Protect option is disabled, you need to manually:

- Add newly added namespaces to the Protection Group.
- From the Protection Group, remove existing namespaces that are deleted on Red Hat OpenShift. The existing backups of the namespaces are preserved until the snapshot expires.

Enable or Disable Auto Protect at Red Hat OpenShift cluster level

To enable or disable Auto Protect at the Red Hat OpenShift cluster level:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for the Red Hat OpenShift cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the Kubernetes page, click the edit icon in the **Objects** section.

5. Click the icon in the cluster level to enable or disable Auto Protect and click **Save Selection**.
6. Click **Save**.

Enable or Disable Auto Protect at Label level

To enable or disable Auto Protect at the Kubernetes label level:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for the Red Hat OpenShift cluster.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. On the **Kubernetes** page, click the edit icon in the **Objects** section.
5. Select the label icon ( ) available on the top-right of the **Objects** section.
6. From the **label** drop-down list, select the label(s) for which you want to enable **Auto Protect**.
7. Click the icon in the enable or disable Auto Protect and click **Save Selection**.
8. Click **Save**.

Attach New Protection Policy to the Protection Group

A Protection Policy can be attached to multiple Protection Groups, and any changes in the policy settings impact all the Protection Groups that are associated with the Protection Policy. You can create and attach a new Protection Policy with the required policy settings to your Protection Group. The new policy is applied to the Protection Group from the next scheduled protection run.

To attach a new Protection Policy to the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Red Hat OpenShift.
3. Click the Actions menu (⋮) next to the Protection Group and select **Edit**.
4. In the **Policy** drop-down list, select the new Protection Policy for the Protection Group.
5. Click **Save**.

Start an On-demand Protection Run

You can manually start a protection run of the Protection Group that you have created to back up the namespaces of the Red Hat OpenShift cluster. It provides you with options to:

- Back up all the objects or only selected objects in the Protection Group.
- Select an incremental or full backup.

- Add replication.
- Add Archival.

To start a protection run:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Red Hat OpenShift.
3. Click the Actions menu (⋮) next to the Protection Group and select **Run Now**.
4. In the pop-up window, perform the following:
    a. In the **Backup** drop-down list, select any one of the following:
        - **Backup all Objects in the Protection Group**
        - **Backup only selected Objects in the Protection Group**
    b. In the **Backup Type** drop-down list, select the backup type. By default, an incremental backup will be performed.
    c. Optional. In the **Replicate** section, select the replication target and the retention duration.
    d. Optional. In the **Archive** section, select the Archival target and the retention duration.
5. Click **Run Now**.

Pause Future Protection Runs

You can pause future protection runs for the Protection Group. However, if the Protection Group is currently executing a protection run, the current run continues to execute and only future runs of the Protection Group are paused.

To pause future protection run:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Red Hat OpenShift.
3. Click the Actions menu (⋮) next to the Protection Group and select **Pause Future Runs**.

Resume Paused Protection Runs

You can resume a paused protection run. When you resume a protection run, if the backup snapshots are expired then a full backup is triggered.

To resume a paused protection run:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had paused.
3. Click the Actions menu (⋮) next to the Protection Group and select **Resume**.

Delete the Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups and the local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- Delete Protection Group Only: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- Delete Protection Group and Snapshots: Allows you to delete the Protection Group and all the local and archived Snapshots generated by the Protection Group.

  After deleting the Protection Group and Snapshots, the reclaimed storage space details are not displayed immediately. It will take some time to reflect on the Cohesity Dashboard.

To delete the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the Protection page, select the Protection Group that you had created for Red Hat OpenShift.
3. Click the Actions menu (⋮) next to the Protection Group and select **Delete**.
4. In the **Delete Protection Group?** dialog box, select any one of the following options:
   - **Delete Protection Group Only**
   - **Delete Protection Group and Snapshots**

Backup of Selected VMs

Cohesity supports the backup of selected VMs for Red Hat OpenShift Virtualization present in the Kubernetes namespaces. You can back up individual VMs or select groups of VMs, rather than backing up the entire namespace.

Cohesity uses the VM Snapshot API that enables taking a single snapshot of all volumes associated with a VM in one backup operation.

To protect the selected VMs:

1. In the Cohesity cluster UI, perform one of the following:
   - Select **Data Protection > Protection**. Click **Protect** displayed on the top-right of the page and then select Red Hat OpenShift Cluster.
   - Select Data Protection > Sources. Click the Actions menu (⋮) next to the Red Hat OpenShift source and then select Protect.

2. From the **Registered Source** drop-down list, select the registered Red Hat OpenShift source or click **Register Source** to register a Red Hat OpenShift source. For more information, see Register Red Hat OpenShift Cluster as a Source.
3. From the **Objects** list, select the VMs present in the namespace that you plan to backup and click **Save Selection**.

> **Note:** You can use the advanced filters that are available to search and find the namespaces you plan to protect. Also, you can click the shield icon to enable Auto Protect at the Red Hat OpenShift cluster level or based on the namespace label(s). For more information, see Backup of Selected VMs.

Exclude or Include VMs

Cohesity enables you to choose which VM to back up or exclude for Red Hat OpenShift Virtualization in a Kubernetes Namespace or Protection Group. Select the VM that is currently configured with a PVC. You can choose to either enable VM inclusion/exclusion or enable PVC inclusion/exclusion.

To include or exclude VMs for Red Hat OpenShift Virtualization in a Kubernetes namespace:

> **Note:**
> Ensure that you meet the following prerequisites:
>
> - On the Red Hat OpenShift cluster, enable Virtualization.
> - The OADP operator must install the kubevirt plugin.
> - The OAPD operator must install the Cohesity DataProtect plugin.

a. In the **New Protection** dialog, click the **Edit** ( ) icon next to that selected namespace.

> **Note:** To turn ON the **VM Inclusion/Exclusion** option to an existing Protection Group, click the Actions menu (⋮) for that Protection Group and select **Edit**. In the **Edit Protection** dialog, under **Objects**, select the namespace and click the **Edit** icon.

b. In the **Options** dialog, turn ON the **Enable VM Inclusion/Exclusion** toggle.
c. You can choose to either include or exclude the VMs. Select **Include** or

Exclude, and then select the VMs from the drop-down.

d. Click **Save**.

4. When you select the VM, the **Leverage CSI snapshot** toggle is enabled by default. This option is necessary for making a VM selection, as snapshots are essential for ensuring a consistent backup of the VM.
5. Click **Protect** to complete the configuration of the protection group.

   More information about the selected VMs is available in the protection run details.

> **Note:**
> - If any of the volumes attached to the VM are not backed up correctly or if it fails, then the backup for the VM will also fail, and a warning message will appear.
> - When an entire namespace is protected, individual VMs or groups can still be selected during recovery, provided that certain plugin and filter requirements are met, and CSI snapshots are enabled.
> - We can check the protection run details to determine if granular backups are being performed or skipped. If it indicates "Skipping," this means some requirements were not met.

Advanced Settings

Advance Settings	Description
Pause Future Runs	Enable **Pause Future Runs** to stop protection runs of the Protection Group from executing. Once you enable this option, no protection runs will be scheduled for this Protection Group.
End Date	Enable **End Date** and select the date when the Protection Group stops capturing snapshots. A Protection Group run that starts prior to this date runs until completion even if it completes after this date.

Advance Settings	Description
QoS Policy	Select an appropriate quality of service (QoS) policy. **Note:** Cohesity recommends specifying **Backup HDD**, which is the default. • **Backup HDD**: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. • **Backup SSD**: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. • **Backup Auto**: (Applicable only for Cohesity C6000 series) The Cohesity cluster writes the data to both SSDs and HDDs. The distribution of data will be based on the current usage of SSD and HDD. This policy tries to achieve a similar backup performance as the **Backup SSD** policy and reduces the SSD wear-out compared to the Backup SSD policy. For best performance, Cohesity recommends the Backup Target SSD policy. If necessary, you can change the policy at any time later. But it doesn't change or take effect on the currently running task.
Alerts	Select one or more of the following settings if you want alerts to be created for the following triggers: • **Success**: Creates an informational alert when a Protection Group completes successfully. Emails are not sent when informational alerts are created. • **Failure**: Creates a critical alert if the Protection Group fails to complete. Emails are sent when critical alerts are created. • **SLA Violation**: Creates a warning alert if the takes longer than the time period specified in the SLA field. Emails are sent when warning alerts are created. For more information, see Alerts.
Email Recipients	You can add email addresses to a Protection Group to notify the email recipients when alerts are triggered for the protection run.
Priority	Select a priority for the Protection Group eProtection Groupxecution. Cohesity supports concurrent backups. However, if the number of protection runs exceeds the ability to process runs, this will be the priority of implementation of the runs: a. High-priority protection runs b. Medium-priority protection runs c. Low-priority protection runs

Advance Settings	Description
SLA	The service-level agreement (SLA) defines how long the administrator expects a Protection Group run to take. • **Full**: Enter the number of minutes you expect a full backup Protection Group run takes to complete. A full backup captures the entire object (all blocks). • **Incremental**: Enter the number of minutes you expect an incremental backup Protection Group run takes to complete. An incremental backup captures only the differences (changed blocks) since the last run. > **Note:** Users with similar SLAs can be grouped together and a custom policy and Protection Group can be created to achieve different SLAs.
Description	Enter a brief description about the Protection Group.
Include or Exclude Labels	Cohesity enables you to choose which PVCs or VMs to include or exclude in the Protection Group. To select the PVCs or VMs: a. Turn ON the **Enable PVC Inclusion/Exclusion** toggle. b. In the **Logical Rule** drop-down, select "Match Any of the following labels" or "All of the following labels". c. Select either **PVCs** or **VMs**. d. Select **Include** or **Exclude**. e. Click **+ Add**. f. Enter the PVC's or VMs label. The label is the unique identifier assigned to that PVC or VM in the Red Hat OpenShift cluster. You can add more PVC or VM until you have all the PVCs or VMs you want to be included or excluded for that Protection Group. > **Note:** > • PVC or VM Inclusion or Exclusion settings of the Namespace take precedence over the settings of its Protection Group. > • Do not modify these settings while protecting VMs.

Source-Side Deduplication and Cache Optimization for VMs and PVCs

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Source-side deduplication is a Cohesity feature designed to optimize backup efficiency by eliminating redundant data at the source (VMs and PVCs) before transferring it to the Cohesity cluster. This process reduces network usage and speeds up backup operations, ensuring efficient data protection for your VMs and PVCs.

With source-side deduplication enabled for a VM or PVC during Protection Group creation, unique data blocks that have not been previously transferred are sent to the Cohesity cluster. This significantly minimizes network bandwidth utilization and reduces data transfer times.

Additionally, source-side deduplication integrates seamlessly with Cohesity's built-in deduplication, further optimizing storage by processing only new data blocks.

Advantages of Source-Side Deduplication

- **Reduced Network Traffic**: This feature minimizes the amount of data transferred between VMs or PVCs and the Cohesity cluster, conserving network bandwidth.
- **Faster Backups**: Speeds up backup processes, particularly over slower networks such as WANs, enabling quicker recovery point objectives (RPOs).
- **Efficient Resource Utilization**: Offloads deduplication processing from the Cohesity cluster, allowing the system to handle a higher volume of backup workload sessions.
- **Lower Storage Costs**: Reduces overall storage consumption by transmitting only new or changed data blocks.

Considerations

After enabling source-side deduplication, if you back up a VM that uses a sparse filesystem-type PVC, recovering the PVC converts it to a thick provisioned format. This results in increased storage consumption on the Red Hat OpenShift cluster. For large recovery operations, the increased space usage may lead to insufficient storage capacity on the Red Hat OpenShift side.

Recover Red Hat OpenShift Virtualization

After you protect your Red Hat OpenShift Cluster, you can recover namespaces to the same or an aternate location.

Recovering Namespaces to the Same Location

When you are recovering the namespaces to the same or original location, then only the VMs, resources, PVCs, or the metadata of the Kubernetes namespaces that are missing or deleted from the original location are restored from the namespace snapshot you selected for recovery. Cohesity does not overwrite or restore the VMs, resources, PVCs, or the metadata of the Kubernetes namespace that already exists in the original location.

For example, a Kubernetes namespace you want to restore to the same location contains a deployment resource and a service account. When Cohesity initiates the restore to the

original location and if the service account is missing from the original location, then Cohesity will restore only the service account. The restore of the deployment resource is skipped.

Cohesity supports the recovery of PVCs with the volumeMode set to Filesystem or Block.

To restore the namespaces of a Red Hat OpenShift cluster to the same location:

1. In the Cohesity cluster UI, navigate to **Data Protection > Recoveries**.
2. Click **Recoveries** on the top-right of the page and then select **Red Hat OpenShift Cluster** as the recovery source.

   The Kubernetes Namespaces page appears.
3. Search for a namespace or Protection Group that contains the snapshots to recover from by entering characters of the namespace name, Protection Group name, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, Storage Domain, or date.
4. From the list of search results, select one or more objects (namespaces) or Protection Groups that contain the snapshots to recover.

   > **Note:**
   > - If you have selected a Protection Group, then all the namespaces backed up in the latest protection run of this Protection Group will be selected for recovery. The option to select these namespaces individually will be greyed out in the search result list.
   > - You can only select Protection Groups and namespaces with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a namespace or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, perform the following steps depending on the object selected:

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Protection Group	a. From the selected lists, mouse over the snapshot, and then click the **edit** icon. b. From the **Protection Run** drop-down menu, select the snapshot from which you want to recover the namespaces. By default the snapshot is recovered from the local cluster. c. Optionally, you can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6.
Namespace	a. From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <namespace_name> page appears. By default, the page displays the snapshots available in the last one month. b. Optionally, you can also select a snapshot from a different date range or by selecting the date range from the **Date** drop-down list. Once you select the date range and then click **Apply**, the Edit recovery point page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. By default, the snapshot is recovered from the local cluster. d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the namespace to the original location.
8. From the **PVC Selection** option, select the PVCs of the namespace you want to recover.

> **Note:** Do not modify these settings while protecting VMs.

9. Optional. Edit the default settings of the following **Recovery Options**:

> **Note:**
> - When you select the recovery of a namespace, existing resources will not be overwritten. Only the deleted resources will be replaced by the recovered VMs.
> - Cohesity recommends that you follow the default settings.

- **Rename**: Add Prefix and/or Suffix strings to the names of the new namespaces created by this task. For example, you could specify Prefix Test_ and Suffix _QA to the namespace's name WinOS and when the namespace is recovered, the new namespace that is created is called Test_WinOS_QA.
- **Task Name**: Change the default name of the recovery task.
- **Storage Class**: Enable the **PVCs will be recovered to mapped Storage Class.** option if you want to map the PVCs to be recovered to a different storage class. Click the **+Add** option and from the **Current Storage Class** drop-down, select the storage class you want the PVCs to be mapped.

10. Click **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recovering Namespaces to an Alternate Location

You can recover the namespaces of the Kubernetes to an alternate Red Hat OpenShift cluster that is registered with the Cohesity cluster. If you are recovering the Kubernetes namespace to an alternate location, then the entire namespace is restored. However, if you are restoring to this location again, then Cohesity will not restore the VMs, resources, PVCs, or the metadata of the Kubernetes namespace that already exist in the location.

To restore the namespaces of a Red Hat OpenShift cluster to the same location:

1. In the Cohesity cluster, navigate to **Data Protection > Recoveries**.
2. Click **Recoveries** on the top-right of the page and then select **Red Hat OpenShift Cluster** as the recovery source.

   The Kubernetes Namespaces page appears.
3. Search for a namespace or Protection Group that contains the snapshots to recover from by entering characters of the namespace name, Protection Group name, or by specifying the wildcard character *.

   You can optionally filter the search results by source, Protection Group, Storage Domain, or date.
4. From the list of search results, select one or more namespaces or Protection Groups that contain the snapshots to recover.

> **Note:**
> - If you have selected a Protection Group, then all the namespaces backed up in the latest protection run of this Protection Group will be selected for recovery. The option to select these namespaces individually will be greyed out in the search result list.
> - You can only select Protection Groups and namespaces with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a namespace or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

By default, the latest snapshot is selected for recovery

5. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Protection Group	a. From the selected lists, mouse over the snapshot, and then click the **edit** icon. b. From the **Protection Run** drop-down menu, select the snapshot from which you want to recover the namespaces. By default the snapshot is recovered from the local cluster. c. Optionally, you can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a..	Then..
Namespace	a. From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <namespace_name> page appears. By default, the page displays the snapshots available in the last one month. b. Optionally, you can also select a snapshot from a different date range or by selecting the date range from the **Date** drop-down list. Once you select the date range and then click **Apply**, the Edit recovery point page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. By default, the snapshot is recovered from the local cluster. d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** to recover the namespace to a new location.
8. From the **Registered Source** drop-down list, select the registered Red Hat OpenShift source to which you want to recover the namespaces or click **Register Source** to register an Red Hat OpenShift source. For more information, see Register Red Hat OpenShift Cluster as a Source.
9. From the **PVC Selection** option, select the PVCs of the namespace you want to recover.

   > **Note:** Do not modify these settings while protecting VMs.

10. Optional. Edit the default settings of the following **Recovery Options**:

> **Note:**
> - When you select the recovery of a namespace, existing resources will not be overwritten. Only the deleted resources will be replaced by the recovered VMs.
> - Cohesity recommends that you follow the default settings.

- **Rename**: Add Prefix and/or Suffix strings to the names of the new namespaces created by this task. For example, you could specify Prefix Test_ and Suffix _QA to the namespace's name WinOS and when the namespace is recovered, the new namespace that is created is called Test_WinOS_QA.
- **Task Name**: Change the default name of the recovery task.
- **Storage Class**: Enable the **PVCs will be recovered to mapped Storage Class.** option if you want to map the PVCs to be recovered to a different storage class. Click the **+Add** option and from the **Current Storage Class** drop-down, select the storage class you want the PVCs to be mapped.

11. Click **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Selected VMs Present in the Namespace

Cohesity supports recovery of selected VMs for Red Hat OpenShift Virtualization present in the Kubernetes namespaces. You can recover individual VMs or a group of VMs from a backup, rather than recovering the entire namespace.

Recovering Selected VMs Present in the Namespace to the Original Location

You can restore the VMs present in the namespace of a Red Hat OpenShift cluster to the same location.

1. In the Cohesity cluster UI, navigate to **Data Protection > Recoveries**.
2. Click **Recoveries** on the top-right of the page and then select **Red Hat OpenShift Cluster** as the recovery source.

   The Kubernetes Namespaces page appears.
3. Search for a namespace that contains the snapshots to recover from by entering characters of the namespace name or by specifying the wildcard character *.

   You can optionally filter the search results by source, Storage Domain, or date.
4. From the list of search results, select one or more objects (namespaces) that contain the snapshots to recover.

> **Note:** You can only select namespaces with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a namespace, the **Storage Domain** and **Source** filters will automatically get locked.

By default, the latest snapshot is selected for recovery.

5. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Namespace	a. From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <namespace_name> page appears. By default, the page displays the snapshots available in the last one month. b. Optionally, you can also select a snapshot from a different date range or by selecting the date range from the **Date** drop-down list. Once you select the date range and then click **Apply**, the Edit recovery point page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. By default, the snapshot is recovered from the local cluster. d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **Original Location** to recover the VMs to the original location.
8. Optional. Edit the default settings of the following **Recovery Options**:

> **Note:**
> - When you select the recovery of a namespace, existing resources will not be overwritten. Only the deleted resources will be replaced by the recovered VMs.
> - Cohesity recommends that you follow the default settings.

- **Rename**: Add Prefix and/or Suffix strings to the names of the new namespaces created by this task. For example, you could specify Prefix Test_ and Suffix _QA to the namespace's name WinOS and when the namespace is recovered, the new namespace that is created is called Test_WinOS_QA.
- **Task Name**: Change the default name of the recovery task.
- **Namespace Resources**: Allows you to change the namespace resources you want to recover. Enable this option to customize the following settings:
  - **Resources**: In the Resources you can select **VM Inclusion/Exclusion** and select the VMs to include or exclude from recovery. When you select a VM for recovery, you have the following options:
    a. **Instant Recovery**: The VMs will be instantly available after recovery in the target location. The data will be moved to the target storage later. The VM can be used in production during Instant Recovery when the data is being copied to the VMs target data volume from the Source PVC (backed by Cohesity NFS view). In such scenarios, the VM performance may be low as the Storage vMotion is still in progress.

       Cohesity recommends that you choose Copy Recovery for powered-off VMs for better performance.

    b. **Copy Recovery**: The VMs will be available in the target location only after all the data is copied to the target storage from the source location (Cohesity cluster or cloud). Once the restore is complete, the VM will be available to use.
  - **Storage Class**: Enable the **PVCs will be recovered to mapped Storage Class.** option if you want to map the PVCs to be recovered to a different storage class. Click the **+Add** option and from the **Current Storage Class** drop-down, select the storage class you want the PVCs to be mapped.

9. Click **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page. More information about the selected VMs is available in the recovery details.

Recovering Selected VMs Present in the Namespace to an Alternate Location

You can restore the VMs present in the namespace of a Red Hat OpenShift cluster to a new location:

1. In the Cohesity cluster UI, navigate to **Data Protection > Recoveries**.
2. Click **Recoveries** on the top-right of the page and then select **Red Hat OpenShift Cluster** as the recovery source.

   The Kubernetes Namespaces page appears.
3. Search for a namespace that contains the snapshots to recover from by entering characters of the namespace name or by specifying the wildcard character *.

   You can optionally filter the search results by source, Storage Domain, or date.
4. From the list of search results, select one or more objects (namespaces) that contain the snapshots to recover.

   > **Note:** You can only select namespaces with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a namespace, the **Storage Domain** and **Source** filters will automatically get locked.

   By default, the latest snapshot is selected for recovery.
5. To recover from a different snapshot, perform the following steps depending on the object selected:

   > **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a..	Then..
Namespace	a. From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <namespace_name> page appears. By default, the page displays the snapshots available in the last one month. b. Optionally, you can also select a snapshot from a different date range or by selecting the date range from the **Date** drop-down list. Once you select the date range and then click **Apply**, the Edit recovery point page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. By default, the snapshot is recovered from the local cluster. d. Optionally, you can also select snapshots archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

6. Click **Next: Recover Options**.
7. From the **Recover To** options, select **New Location** to recover the VMs to the new location.
8. From the **Registered Source** drop-down list, select the registered Red Hat OpenShift source to which you want to recover the VMs present in the namespace or click **Register Source** to register an Red Hat OpenShift source. For more information, see Register Red Hat OpenShift Cluster as a Source.
9. Optional. Edit the default settings of the following **Recovery Options**:

> **Note:**
> - When you select the recovery of a namespace, existing resources will not be overwritten. Only the deleted resources will be replaced by the recovered VMs.
> - Cohesity recommends that you follow the default settings.

- **Rename**: Add Prefix and/or Suffix strings to the names of the new namespaces created by this task. For example, you could specify Prefix Test_ and Suffix _QA

to the namespace's name WinOS and when the namespace is recovered, the new namespace that is created is called Test_WinOS_QA.

- **Task Name**: Change the default name of the recovery task.
- **Namespace Resources**: Allows you to change the namespace resources you want to recover. Enable this option to customize the following settings:
  - **Resources**: In the Resources you can select **VM Inclusion/Exclusion** and select the VMs to include or exclude from recovery. When you select a VM for recovery, you have the following options:
    a. **Instant Recovery**: The VMs will be instantly available after recovery in the target location. The data will be moved to the target storage later. The VM can be used in production during Instant Recovery when the data is being copied to the VMs target data volume from the Source PVC (backed by Cohesity NFS view). In such scenarios, the VM performance may be low as the Storage vMotion is still in progress.

       Cohesity recommends that you choose Copy Recovery for powered-off VMs for better performance.

    b. **Copy Recovery**: The VMs will be available in the target location only after all the data is copied to the target storage from the source location (Cohesity cluster or cloud). Once the restore is complete, the VM will be available to use.

       > **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

  - **Storage Class**: Enable the **PVCs will be recovered to mapped Storage Class.** option if you want to map the PVCs to be recovered to a different storage class. Click the **+Add** option and from the **Current Storage Class** drop-down, select the storage class you want the PVCs to be mapped.

10. Click **Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page. More information about the selected VMs is available in the recovery details.

Recover Files or Folders for VMs

After you protect your Red Hat OpenShift cluster, you can recover selected files or folders from the VM backups to either the original location on the VM or an alternate location, eliminating the need for a full VM restore.

To learn more, see About Recovering VMs.

For more information about the recovery options:

- Recover Files or Folders to the Original Location within VMs
- Recover Files or Folders to the New Location within VMs

Indexing and Recovery

Note the following about indexing and recovering files:

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the lvol_N directory, where N is a unique number such as 1.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as C:.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Server Type	Volume Type	

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

- Cohesity supports recovering files/folders from NTFS (Windows VMs) to Windows VMs, and from Linux VMs to Linux VMs only.
- When recovering a file or instantly mounting a volume from a Windows VM or Physical Server Backup Source that has Windows deduplication installed and enabled for one or more volumes, you must choose a target machine that also has Windows deduplication installed (it does not have to be enabled for any volume). (However, this rule does not apply to Nutanix AHV VMs. If AHV VMs are enabled with Windows deduplication, the only supported recovery option is full VM recovery.)

  If the target does not have Windows deduplication installed:

  - File level recovery will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."
  - Instant volume mounting will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."

  To determine if Windows deduplication is installed on the Source or target machine, follow the steps given below:

1. Open **Server Manager**.
2. Select **Roles and Features > File and Storage Services > File and iSCSI Services**.
3. Select the **Data Deduplication** check box, if necessary.
4. Click **Next** until the **Install** button is enabled and then click **Install**.

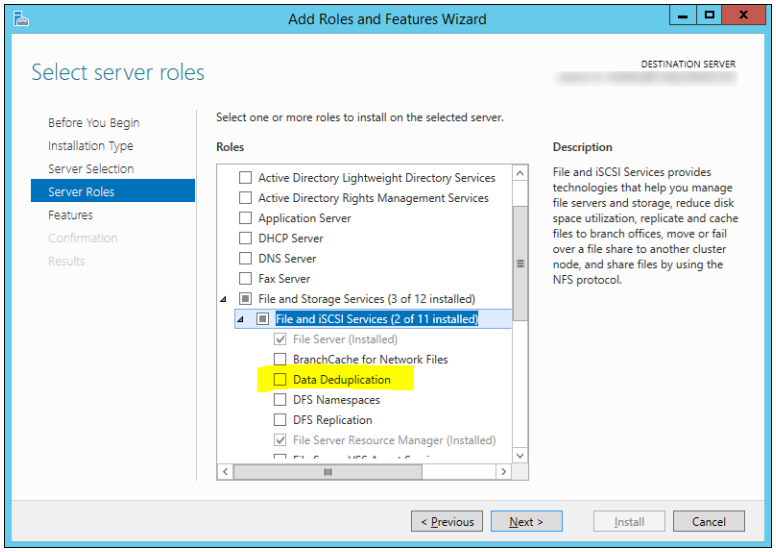

- Before you can recover a file from a snapshot, the snapshot must exist. To recover files from a snapshot, the files must be indexed during the backup process.
- When unzipping a zip file that was created by downloading files and folders from an archived Snapshot, if the file or folder name has encoded characters, unzip the zip file using the corresponding encoding. For example if a file name in the zip file has a UTF-8 character, unzip the file using the following command:

```
unzip -O UTF-8 Download-Files_Sep_20_2018_3-17pm_3090.zip
```

Prerequisites

QEMU Guest Agent Installation

The QEMU Guest Agent must be installed and running inside the RHOV virtual machine. The guest agent is required to execute commands inside the VM for file-level recovery operations.

Linux VMs

1. Install the guest agent:

```
sudo apt-get install -y qemu-guest-agent
```

2. Enable and start the service:

```
sudo systemctl enable qemu-guest-agent
sudo systemctl start qemu-guest-agent
```

Windows VMs

Install the VirtIO drivers and QEMU Guest Agent and run the qemu-ga-x86_64.msi installer. The service starts automatically after installation.

Verify Guest Agent Connectivity

1. Verify that the QEMU Guest Agent is connected from outside the VM:

```
oc get vmi <vm-name> -n <namespace> -o jsonpath='{.status.conditions}'
| python3 -m json.tool
```

2. Confirm that the following condition is present:

   "type": "AgentConnected"
   "status": "True"

Enable guest-exec Command

The guest-exec and guest-exec-status RPC commands must be enabled in the QEMU Guest Agent configuration.

Linux VMs

1. Edit the configuration file:

```
sudo vi /etc/sysconfig/qemu-ga
```

2. Update RPC settings:
   - If using BLOCK_RPCS, remove:

```
guest-exec, guest-exec-status
```

   - If using AL, add LOW_RPCS:

```
guest-exec, guest-exec-status
```

3. Restart qemu-guest-agent.

Windows VMs

1. Update the QEMU Guest Agent service configuration to ensure guest-exec is not blocked.

1. Restart-Service QEMU-GA.

Verify guest-exec Enablement

- Run the following command

```
oc exec -it <virt-launcher-pod> -n <namespace> -c compute -- \
virsh qemu-agent-command <namespace>_<vm-name> \
'{"execute":"guest-exec"}'
```

- A successful response returns:

```
{"return":{"pid":<pid>}}
``` [1](https://cohesity-
```

SELinux Configuration (Linux Only)

Configure SELinux to allow QEMU Guest Agent operations:

```
sudo semanage permissive -a virt_qemu_ga_t
```

Considerations

- Dependency on QEMU Guest Agent
  - File-level recovery requires the QEMU Guest Agent to be installed and connected.
  - Recovery operations depend on successful communication with the agent.
- guest-exec Command Requirement
  - File-level recovery depends on the guest-exec and guest-exec-status commands being enabled.
  - These commands are disabled by default and must be explicitly configured.
- SELinux Restrictions
  - SELinux in enforcing mode can restrict operations required for file-level recovery, such as mounting backup disk partitions.
- Platform-Specific Behavior
  - SELinux configuration applies only to Linux VMs.
  - Windows VMs do not require SELinux-related configuration.
- Agent Connectivity Requirement
  - The QEMU Guest Agent must report a connected state (AgentConnected = True) for recovery operations to proceed.
- Mount Points Without /etc/fstab Entry
  Indexed paths depend on how the cluster associates backup content with guest mount metadata.
  Runtime-only mounts or volumes that are not defined persistently (for example, no /etc/fstab entry on Linux or no drive letter on Windows) might not appear under their original paths in the UI.
  In such cases, content may be grouped under generic or synthetic prefixes (for example, lvol_N, where N is a unique number).
  Recommendation: Use clear and persistent mount definitions for volumes that must be browsable and recoverable.
- LVM Handling During Recovery
  For Linux guests using LVM, recovery relies on importing and activating volume groups so backup volumes can be mounted within the guest.
```

This requires LVM packages, vgimportclone, and activation using vgchange -ay. If these utilities are not available, recovery might fail.

- No Cross-OS Recovery
  Cross-operating system file-level recovery is not supported.
  Windows → Linux: Not supported
  Linux → Windows: Not supported.
- Windows Data Deduplication
  If Windows Data Deduplication is enabled, recovery to a VM without deduplication drivers might fail or result in incomplete recovery. Use Download Files, full VM recovery, or other supported workflows if recovery fails.
- Storage Replica on Windows with GPT Disks
  If Storage Replica is enabled with GPT disks, original-location recovery might be unsupported or unreliable. Alternatives include recovering to another path, full VM recovery, or using download.
- Long File Paths (Windows Targets)
  Paths longer than approximately 200 characters might fail during restore operations due to Windows limitations.
  Mitigation: Use shorter paths, restore smaller sets, or rename files after restore.
- Hard Links, Reparse Points, and Alternate Data Streams
  Recovery does not guarantee preservation of NTFS features such as alternate data streams, reparse points, or hard link relationships.
- Sparse File Support on Windows
  Sparse file handling is not supported during recovery. Sparse files may be restored as fully allocated files.
- Copy Tools: rsync and robocopy
  Recovery relies on rsync on Linux and robocopy on Windows. If these tools are missing or fail, metadata such as ACLs and attributes may not be preserved.
- Supported Operating System Versions
  Support for backup and file-level recovery follows Cohesity RHOV interoperability guidelines, including requirements for QEMU Guest Agent, guest-exec, and SELinux where applicable.

Troubleshooting

QEMU Guest Agent Not Connected

1. Check agent status inside the VM

   ```
   sudo systemctl status qemu-guest-agent
   ```

2. Verify connectivity from KubeVirt:

   ```
   oc get vmi <vm-name> -n <namespace> -o jsonpath='{.status.conditions[?(@.type=="AgentConnected")].status}' ``` [1](https://cohesity-
   ```

guest-exec Still Blocked

1. Verify configuration.

   ```
   cat /etc/sysconfig/qemu-ga
   ```

2. Restart the agent:

   ```
   sudo systemctl restart qemu-guest-agent
   ```

> **Note:** qemu-ga -b=help displays potential blocked commands, not the current active configuration.

Mount fails with "Permission denied"

Check current mode

```
getenforce
```

Recover Files or Folders to the Original Location within VMs

To recover files and folders to the original server:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Files or Folders**.

   The Files page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select...	Then...
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select...	Then...
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select...	Then...
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.

- **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

  Character replacement explanation for : and \ in file and directory names in zip files

  If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
  Example:

  - The original name is "a:b" then the archived name is "a_COLON_b"
  - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

  If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
  Example:

  - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
  - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **Original location** to recover the files or folders to the original server.
6. To recover to a different location in the original server, disable the **Recover to Original Path** option, and then provide the location to which the files or folders are to be recovered in the **Recover To** field. By default, the files and folders will be recovered to the original location.
7. Optional. Edit the default settings of the following **Recovery options**:

   - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

     Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

     Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

     - C:\restore\1.txt
     - C:\restore\cohesity<taskid>_0\1.txt
     - C:\restore\bar
     - C:\restore\cohesity<taskid>_1\bar

If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.

- **Preserve File/Folder Attributes**: File attributes are preserved when rsync for (Linux) or robocopy (Windows) is available and functioning correctly. If these tools are not present, preservation of file attributes is not guaranteed.

  By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only, files, then files will receive the new timestamps.
- **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
- **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
- **Task Name**: Change the default name of the recovery task.

8. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Recover Files or Folders to the New Location within VMs

To recover files and folders to the new location:

1. Navigate to **Data Protection > Recoveries**.
2. Click **Recover** located at the top right of the page and then select **Files or Folders**.

   The Files page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select...	Then...
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select...	Then...
Browse	1. Enter the full or partial server or Protection Group name. You can by specify the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select...	Then...
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lvol_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	

Linux Virtual Machine	Simple Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

4. Once you select the files and folder to recover, select one of the following options:
   - **Next Recover Option**: If you select this option, then continue with the step 5 to configure the file recovery options.

- **Download Files**: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

  Character replacement explanation for : and \ in file and directory names in zip files

  If the original file or directory name contains a ':' , it is replaced with '_COLON_' in the archive.
  Example:

  - The original name is "a:b" then the archived name is "a_COLON_b"
  - The original name is "a:b:c" then the archived name is "a_COLON_b_COLON_c"

  If the original file or directory name contains a '\', it is replaced with '_BACK_SLASH_' in the archive.
  Example:

  - The original name is "x\y" then the archived name is "x_BACK_SLASH_y"
  - The original name is "x\y\z" then the archived name is "x_BACK_SLASH_y_BACK_SLASH_z"

5. From the **Recover To**, select **New location** to recover the files or folders to the original server.
6. From the **Registered Source** drop-down menu, select a registered server to which you want to recover the files and folders. You can also register a new source by clicking **Register Source** from the **Source** drop-down menu.
7. From the **Target**drop-down menu, select a VM or volume to which the files and files are to be recovered.
8. In the Recover To field, change the default location to which the files or folders are recovered. By default, the files and folders are recovered to the /tmp/ directory for VMs.

   Example for Windows alternate location:

   D:\cohesity_restore\1.txt

   D:\cohesity_restore\bar

   These requirements may need to be met to recover files to a Windows VM.

   If recovering to a standalone Hyper-V host or cluster, you must add the guest VM to the reverse lookup zone in the DNS server to complete the recovery successfully.

   Files to be recovered: "C\foo\1.txt" and "C:\1stfloor\bar". Specified alternate location: "D:\cohesity_restore". At the end of the copy, the target folder looks like this:

9. Optional. Edit the default settings of the following **Recovery options**:
    - **Overwrite Existing File/Folder**: By default, this option is enabled to overwrite the existing files and folders. Disable this option to create the files and folders in the specified location. Any duplicate files are skipped.

      Example for a Windows alternate location with **Overwrite Existing File/Folder** enabled and the target already exists:

      Files to be recovered: "C:\foo\1.txt", "C:\foo\test\1.txt", "C:\bar", "C:\foo\bar". Specified alternate location: "C:\restore". At the end of the copy, the target folder looks like this:

        - C:\restore\1.txt
        - C:\restore\cohesity<taskid>_0\1.txt
        - C:\restore\bar
        - C:\restore\cohesity<taskid>_1\bar

      If **Overwrite Existing File/Folder** is enabled, recovering a file to source when the file is in use may cause the open file to be overwritten. Whether overwriting occurs depends on the application using the file.
    - **Preserve File/Folder Attributes**: File attributes are preserved when rsync for (Linux) or robocopy (Windows) is available and functioning correctly. If these tools are not present, preservation of file attributes is not guaranteed.

      By default, this option is enabled and the ACLs, permissions, and timestamps are preserved for all files and folders. If you disable this option, then ACLs and permissions are not preserved. If recovering both folders and files, then folders will receive the new timestamps, but files retain their original timestamps. If recovering only, files, then files will receive the new timestamps.
    - **Continue on Error**: Enable this option if you want to continue the recovery even if one of the objects encounters an error. By default, this option is disabled and the recovery operation will fail if one of the objects encounters an error.
    - **Cluster Interface**: By default, the Auto Select option is enabled and the recovery task automatically selects the correct VLAN. If you disable this option, then select a configured interface group from the Interface Group drop-down.
    - **Task Name**: Change the default name of the recovery task.
10. Click **Recover**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Browse and Download Files and Folders in Red Hat OpenShift Virtualization

Cohesity supports file and folder browsing and download capabilities for Red Hat OpenShift Virtualization. This feature allows you to browse VM content and download single or multiple files or folders through the web interface.

Download Files and Folders

To donwload files and folders:

1. In the Cohesity cluster UI, navigate to **Data Protection > Recoveries**.
2. Click **Recoveries** on the top-right of the page and then select **Kubernetes Cluster > Files or Folders**.

   The Recover Files or Folders page appears.
3. Select any one of the following options to search for the file or folder:
   - **Files and Folders**: Allows you to search and recover files using file or folder names from any backup available on the Cohesity cluster.
   - **Browse or Specify Path for Search**: Allows you to search and recover files using the name of the server or Protection Group from any backup available on the Cohesity cluster.

If you select...	Then...
Files and Folders	1. Enter characters of the file name, path, or by specifying the wildcard character *. A list of items displays. You can optionally filter the search results by object type, source, object, Protection Group, Storage Domain, files only, or folders only. > **Note:** When entering the path of a Windows VM, be sure to reformat as shown in the example: Instead of *C:\Windows\System32\winevt* use *C:/Windows/System32/winevt/*. 2. From the list of search results, select the required files or folders to recover. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot: > **Note:** The option to change the recover point is not available if multiple items were selected for recovery. 1. From the selected list, mouse over the snapshot, and then click the **edit** icon. The Edit recovery point page appears. 2. By default, only the snapshots that are indexed are displayed. To display all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 3. Select a date range and the snapshot point for recovery. 4. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. 5. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up

If you select...	Then...
	to 150 MB/s of retrieval throughput. 4. Click **Select Recovery Point** and go to step 4.

If you select...	Then...
Browse	1. Enter the full or partial server or Protection Group name. You can do this by specifying the wildcard character *. You can also optionally filter the search results by object type, source, object, Protection Group, Storage Domain. 2. From the list of search results, select the object whose files or folders are to be recovered. 3. By default, the latest snapshot is selected for recovery. To recover from a different snapshot, click the snapshots drop-down in the top-left corner and select the snapshot you need. 4. **Browse on Indexed Data**: Toggle on this setting to browse the indexed data. It helps to speed up browsing of recovery data. By default, it only displays the files and folders that are indexed. Hidden files, folders, paths such as /usr, /var, /sys , /Windows, /ProgramData etc. are excluded from indexing. To display all the available files and folders, including hidden files and folders toggle off the **Browse on Indexed Data** setting. **Note:** • Changing the snapshot after selecting the items (files or folder) removes the selected items from the cart. • Enabling or disabling the **Browse on Indexed Data** option after selecting the items (files or folder) removes the selected items from the cart. • Deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process. 5. Browse to the file or folder that you want to recover by clicking a folder and subfolders. 6. Click **Save**. 7. After saving the selection, if you want to recover from a different snapshot or change the location from where the snapshot is recovered, perform the following steps: 1. From the selected list, mouse over the snapshot, and click the edit icon. 2. The Edit recovery point page appears. 3. By default, only the snapshots that are indexed are displayed. To display

If you select...	Then...
	all the available snapshots, toggle off the **Show Indexed Snapshots Only** option. 4. Select a date range and the snapshot point for recovery. 5. By default, the snapshot is recovered from the local Cohesity cluster. You can also select a snapshot archived to the cloud (icon) by clicking this icon if available in the **Location** option. When recovering from an archive, consider the following points: • The archive (external target) must be currently registered for the task to succeed. • (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput. 8. Click **Select Recovery Point** and go to step 4.

4. Once you select the files and folder to recover, choose the **Download Files** option.

> **Note:** The **Next Recover Option** is not supported. Use **Download Files** to download files or folders.

Download Files: If you are recovering a single file, this option downloads the file to your browser's download folder. For all other selections, this creates a recover task. When the task is completed, from the Recoveries page, click the task name and then click **Download Files** to download the generated zip file.

Chapter 22: Private Cloud Platforms

A private cloud is a dedicated computing environment that offers greater security, control, and compliance. Cohesity provides data protection for modern applications deployed in private cloud environments.

VMware Cloud Foundation (VCF) Automation

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Cohesity supports data protection for VMware Cloud Foundation (VCF) Automation environments, enabling backup and recovery operations. VCF Automation is a full-stack cloud platform that delivers a consistent private cloud experience across on-premises data centers, edge environments, and public cloud deployments. It integrates compute, storage, networking, and lifecycle management into a unified software-defined platform supporting both traditional Virtual Machines and modern workloads.

VCF Automation is the self-service automation layer of VMware Cloud Foundation.It turns the underlying vSphere infrastructure into a self-service private cloud that teams consume through declarative APIs and UIs, with policy-based governance and multi-tenancy.

Virtual Machines

Virtual Machines

Cohesity provides a simple and fast backup, recovery, and data management solution for VCF Automation Virtual Machines.

Plan and Prepare for VCF Automation Spurce Protection

Key Concepts for VCF Automation Source Protection

Register a VCF Automation Source

Backup VCF Automation Source

Recover VCF Automation Source

Plan and Prepare for VCF Automation Virtual Machines Protection

This topic provides information on minimum permissions, prerequisites and considerations to protect VCF Automation Virtual Machines with the Cohesity cluster.

Minimum Permissions

Ensure you have the required minimum permissions to register VCF Automation Virtual Machines as a source with the Cohesity, and to perform backup and recovery. For more information, see Ensure Adequate Privileges for Cohesity on the Source.

Support Matrix

Before you register your VCF Automation Virtual Machines with the Cohesity cluster, ensure that Cohesity supports the VCF Automation Virtual Machines versions. For more information, see Supported Software.

Ports Used for Communication

Ensure all necessary ports are open in the firewall between the Cohesity cluster and VCF Automation Virtual Machines. For more information, see Manage Firewall Ports.

Prerequisites and Considerations

Prerequisites

Before you register the VCF Automation source and protect virtual machines, ensure that the following prerequisites are met:

Environment Requirements

- The VCF Automation environment is deployed and accessible.
- The Supervisor is configured on the target vCenter in the workload domain.
- Virtual machines must be created using the VM Service in VCF Automation.

Access and Credentials

- Obtain the VCF Automation endpoint (FQDN or IP address).
- Ensure you have valid credentials for:
  - VCF Automation
  - Associated vCenter instances
- Ensure the user account has the required privileges to:
  - Discover VCF Automation entities
  - Access and manage Virtual Machines

Considerations

- Only VM Service Virtual Machines that are created within the 'All Apps' VCF Automation Organization type are supported.
- Only NBDSSL transport is supported for backup and recovery.
- The following recovery options are not supported:
  - Instant recovery, Clone, Disk recovery, and File & Folder recovery.
  - Instant recovery is not yet supported because Broadcom have not yet implemented storage vmotion for VCF Automation VM Service VMs. Cohesity relies on storage vmotion for instant recovery operations.
  - Recovery to an alternate location.
  - Differential copy recovery.

Key Concepts for VCF Automation Source Protection

Cohesity discovers the following VCF Automation entity hierarchy when you register a source:

Organization

An Organization represents the top-level tenant in VCF Automation. It can correspond to an end customer for a service provider or a department or line of business within an enterprise. An Organization:

- Acts as the primary administrative boundary.
- Is mapped to one or more vSphere Supervisor clusters through Regions to provide infrastructure resources.
- Contains one or more Projects for logical resource organization.

Projects

A Project resides within an Organization and provides a logical grouping of resources based on teams, environments, or business units. A Project:

- Exists under a single Organization.
- Can be created in multiples within an Organization.
- Represents different teams or lines of business.
- Manages user access, roles, and governance policies at the project level.
- Contains one or more Namespaces for workload deployment.

Namespaces

A Namespace resides within a Project and acts as the execution and resource boundary for workloads. A Namespace:

- Exists under a Project, with multiple namespaces supported per project.
- Is instantiated on the vSphere Supervisor cluster associated with the Organization.
- Defines resource allocations such as CPU, memory, storage, and network.
- Carves out a dedicated resource pool for quota management.
- Serves as the deployment location for workloads, including virtual machines and container-based applications.

Virtual Private Cloud (VPC)

A Virtual Private Cloud (VPC) provides networking capabilities similar to public cloud environments, enabling secure and isolated communication for workloads. A VPC:

- Provides network isolation and connectivity for workloads.
- Can be associated with one or multiple Namespaces.
- Enables a cloud-like networking experience within VCF Automation.

Supervisor

A Kubernetes-style control plane embedded in ESXi and vCenter. VCF Automation creates VMs through the Supervisor (VM Service). Each workload-domain vCenter requires at least one Supervisor; one NSX instance supports up to 16 vCenters.

Region

Maps to a Supervisor. In the Cohesity entity hierarchy, each namespace shows the Region it runs on.

VM Service VM

Virtual Machines created within the All Apps organization, directly through VCF Automation.

Register a VCF Automation Source

1. On the **Dashboard** page in the Cohesity cluster, navigate to **Data Protection** > **Sources**.
2. Click **Register** at the top-right of the page and then select **VCF Automation**.

   The **Register VCF Automation** page appears.
3. The workload option **Virtual Machines** is selected by default.
4. Enter the following VCF Automation details:

- **Hostname**: Specify the hostname of the source that contains the VMs to protect.
- **Username**: Enter admin@SYSTEM.
- **Password**: Enter the password for the specified user.

5. Click **Connect** to authenticate the credentials. This Initiates connection and discovery. A list of discovered vCenters such as Management and Workload domain vCenters appears.
6. Enter the following vCenter details:
   - **Username**: You must specify a VMware vCenter user with adequate privileges on the VMware vCenter server. The role assigned to the specified VMware vCenter user should include privileges to turn on Change Block Tracking (CBT) and other privileges required by the Cohesity cluster. For detailed information about the necessary privileges, see Create a Role with Necessary vCenter Server Privileges.

     When specifying a Windows Active Directory username, use the following format: *<domain>\<username>* or *<username>@domain*. For all other cases, just enter: *<username>*.
   - **Password**: Password: Enter the password for the specified vCenter user. Repeat for each discovered vCenter (management and workload domains).
7. Click **Register**.

After source registration finishes successfully, the new source is displayed on the **Sources** page.

View the Registered VCF Automation Source

Entity Hierarchy

The VCF Automation hierarchy is structured as follows:

Organization → Project → Namespace → Workloads (VMs)

- An Organization is mapped to one or more vSphere Supervisors via Regions.
- Each Organization contains one or more Projects.
- Each Project contains one or more Namespaces.
- Namespaces are instantiated on Supervisors and serve as the runtime environment. The namespace on the registered source view entity hierarchy page also shows the region that the namespace is mapped to.
- Workloads, including VMs are deployed within Namespaces.

The source is registered successfully, and the entity hierarchy is discovered. You can view the discovered entities in the hierarchy and list views.

To View VCF Automation Entities Using Hierarchy View:

1. Navigate to **Data Protection** > **Source**.

   The **Sources** page displays all the sources registered with that Cohesity cluster.
2. Select the registered source under VMware.
3. Select the Hierarchy View to browse entities in the following order: Organization > Project > Namespace > VMs. The namespace includes the region details.

To View VCF Automation Entities Using List View:

1. Navigate to **Data Protection** > **Sources**.

   The **Sources** page displays all the sources registered with that Cohesity cluster.
2. Select the registered source under VMware.
3. Select the Hierarchy View to browse entities in the following order: Organization > Project > Namespace > VMs. The namespace includes the region details.

Backup VCF Automation Source

After registering your VCF Automation source, you can use Cohesity DataCloud to protect the source.

To protect VCF Automation source:

1. Navigate to **Data Protection** > **Protection**.
2. Click **Protect** on the top-right of the page and then select **Virtual Machines**.
3. Click **Add Objects** and from the **Registered Source** drop-down list, select the VCF Automation source you want to back up.

   You can also register a new source with the Cohesity cluster so you can protect it.
4. From the objects list, browse the hierarchy or list view. Select one or more VMs under the desired namespace you want to protect.

   > **Note:** You can autoprotect entities at multiple levels in the hierarchy, including Organization, Project, Namespace, and Workloads (VMs).
5. Click **Continue**.
6. In the **Protection Group** field, enter a name for the Protection Group. The name can contain alphanumeric characters, underscores, dashes, and periods.
7. From the **Policy** drop-down list, select an existing Protection Policy or create a new Protection Policy. For more information, see Create or Edit a Standard Policy.
8. If you want to change or configure any of the additional setings, then select **More Options** and perform the below steps or else, click **Protect**.
9. Optionally, from the **Policy** drop-down list, you can change the policy you selected earlier or create a new Protection Policy. For more information, see Create or Edit a

Standard Policy.

10. In the **Storage Domain** field, select the appropriate storage domain.
11. Configure additional options if required.
12. Click **Protect**.

    The protection of VCF Automation Virtual Machines is now active and running.

For more information see, Add or Edit a Protection Group for Virtual Servers.

Recover VCF Automation Source

Cohesity provides the ability to recover Protected Objects (such as Virtual Machines) from a Snapshot created earlier by a Protection Group. You can recover Virtual Machines from a Cohesity cluster. You can recover Virtual Machines to their original location.

Follow these procedures to recover Virtual Machines from a Cohesity cluster:

Recover Virtual Machines to its Original Location

Recover Virtual Machines to its Original Location

To recover Virtual Mchines to its original location:

1. Navigate to **Data Protection > Recoveries**.
2. Click on **Recover** located at the top right of the page and then select **Virtual Machines > VMs**.

   The Virtual Machines page appears.
3. Search for a Protected Object (VM) or Protection Group that contains the snapshots to recover from by entering characters of the VM name, Protection Group name, or by specifying the wildcard character *.

   A list of VMs displays.

   You can optionally filter the search results by Source, Protection Group, Storage Domain, or Date.
4. From the list of search results, select one or more VMs or Protection Groups that contain the snapshots to recover.

> **Note:**
> - If you have selected a Protection Group, then all the VMs backed up in the latest protection run of this Protection Group will be selected for recovery and the option to select those VMs individually will be greyed out in the search result list.
> - You can only select Protection Groups and VMs with snapshots stored in the same Storage Domain and belonging to the same source. Once you select a VM or Protection Group, the **Storage Domain** and **Source** filters will automatically get locked.

5. In the **VM Protection Group Selection** dialog box, review the list of Virtual Machines in the Protection Group and click **Confirm**.
6. By default, the latest snapshot of the VMs is selected for recovery. To recover from a different snapshot, perform the following steps depending on the object selected:

> **Note:** In this workflow, deleted snapshots may be displayed as valid recover points for a brief period until they are removed from the search index by a background process.

To choose a different snapshot for a...	Then...
Protection Group	1. From the selected lists, mouse over the snapshot, and then click the **edit** icon. 2. From the Protection Run drop-down menu, select the snapshot from which you want to recover the VMs. By default the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. 3. Click **Select Recovery Point** and then go to step 6.

To choose a different snapshot for a...	Then...
VM	From the selected lists, mouse over the snapshot, and then click the edit icon. The Edit recovery point for <VM_name> page appears. By default, the option to select the snapshot is displayed in **Timeline** mode. You can also choose to view the option in **List** mode by clicking the List tab in the top-right corner of the page. • To select a snapshot from the Timeline mode: a. From the **Choose a date** drop-down, select a date. The snapshots available on the selected date are displayed in blue dots on the timeline bar. b. Select the snapshot to be recovered by clicking the corresponding blue dot. For CDP, select the point-in-time restore. You can specify the exact point-in-time recovery by sliding the timeline bar. c. By default, the snapshot is recovered from the local Cohesity cluster. You can also select the snapshot archived to the cloud by clicking the cloud ( ) icon if available in the **Location** option. d. Click **Select Recovery Point** and then go to step 6. • To select a snapshot from the **List** mode: a. Select the List tab from the top-right corner of the Edit recovery point for <VM_name> page. By default, the page displays the snapshots available in the last one month. b. You can also select a snapshot from a different date range by selecting the date range from the **Date** drop-down list. The page displays the available snapshots in the specified range. c. Select the snapshots to be recovered. d. By default, the snapshot is recovered from the local Cohesity cluster. You can also select snapshots archived to the cloud by clicking the cloud (

To choose a different snapshot for a...	Then...
	) icon if available in the **Location** option. e. Click **Select Recovery Point** and then go to step 6.

7. Click **Next: Recover Options**.
8. The **Original Location** is selected by default, as it is currently the only available option.
9. VM recovery is performed as a copy recovery by default.

   > **Note:** Only copy recovery is supported, and this is the default recovery method.

10. **Existing VM Handling** options list
    - **VMware**: From the **Existing VM Handling** options list, choose how you want to handle the original VM when recovering to the original location:
      - **None**: Recovers the VM as a new VM and retains the original VM. In the **Rename** field available in the **Recovery Options**, you can specify the name of the recovered VM by adding prefix and/or suffix strings to the name of the original VM. By default, the recovered VM is named as `copy-<original_VM_name>`.
      - **Overwrite Existing VM**: Recovers the VM by deleting the original VM. The recovered VM will have the original VM name. The original virtual machine is deleted prior to the recovery process. Therefore, a recovery failure or cancelation of recovery will also lead to the loss of the original VM.
        Once you select this option, a pop-up dialog box will be displayed. Review the message, type **YES**, and click **Confirm**.
11. Optional. Edit the default settings of the following **Recovery Options**:

Recovery Options	Description
Network	By default, the VMs that are to be recovered do not have a virtual Network Interface Card (vNIC) attached. Enable the **Attach** option to attach a virtual Network Interface Card (vNIC) to each VM that is to be recovered. With the **Attach** option enabled, each VM will connect to its original network and the following option is displayed: • **Start Connected**– Enable this option to connect to the original network when the VMs reboots. If this option is disabled, the VMs are not connected to any network on reboot. While recovering VMs to the original location, there is no explicit control to preserve the MAC address. If the **Attach** option is enabled and the VM is configured to have a manual MAC address assigned at the time of backup, the same address will be recovered. **Note:** When recovering AHV VMs to the original location, the recovered VM will have the same number of NICs as the source VM. If the VM is configured to have an automatic MAC address at the time of backup, the same setting will be applied to the recovered VM.
Rename	(*Applies for VM recovery*) Add Prefix and/or Suffix strings to the names of the new VMs created by this task. For example you could specify Prefix `Test_` and Suffix `_QA` to the VM name WinOS and when the VM is recovered, the new VM that is created is called `Test_WinOS_QA`.
Power State	Disable **Power On** if you want the recovered VMs to remain powered off after they are created. If you recover VMs from a Storage Domain that has Cloud Tier enabled, Cohesity recommends leaving the VMs powered off. By default, the VMs are powered on after they are created.

Recovery Options	Description
Continue on Error	Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cohesity cluster will still attempt to create the other VMs. If this option is disabled, then the recovery task is discontinued if VM creation fails. However, if the VM creation succeeds, but VM storage migration fails, the task continues. By default this option is disabled.
Task Name	Change the default name of the recovery task.

12. Click **Recover**. Review the Recovery Summary and click **Start Recovery**.

The recovery task is initiated. You can monitor the recovery task from the Recoveries page.

Chapter 23: About Bare Machine Recovery

A Bare Machine Recovery (BMR) is the ability to recover a system, including the operating system, its configuration, applications running on the operating system, and its data back to a state it was at a given point.

Cohesity integrates with Cristie software to protect your servers for bare machine recovery (BMR). Cohesity supports two types of BMR:

- **Cristie Bare Machine Recovery (CBMR)**: In this integration, a separate schedule for BMR backup is required to be added to the protection policy.

  First, a block-based protection run is executed followed by a separate BMR run. The BMR run is always Full (no incremental). For more information, see CBMR.
- **Cohesity Bare Machine Recovery (CoBMR)**: In this integration, CoBMR leverages Cohesity's file-based and block-based physical server protection. No separate BMR schedule or protection run is required. For more information, see CoBMR.

> **Note:**
> - For BMR-based protection of a physical server, you must first install the Cohesity agent on the physical server followed by the BMR agent and register the physical server with the Cohesity cluster.
> - You can also deploy the Cohesity agent on a virtual machine and protect the VM as a physical server to perform BMR. For more information, see Deploy Cohesity Agent on Virtual Servers.

CBMR and CoBMR Comparison

Feature	CBMR	CoBMR
Supported physical servers	• Windows • Linux For more information, see Supported Software.	• Windows • Linux For more information, see Supported Software.
Protection Policy	You must add a BMR schedule to the policy.	Default Policy - No BMR schedules.
Protection Group	Block-based backup	File-based backup Block-based backup
Protection Run	First performs block-based protection run followed by a BMR run.	Leverages the physical file-based and block-based backup run (No separate BMR schedule)
System State Backup	For block-based backup: First Full followed by incremental backup. For BMR backup: Always full backup.	First full backup followed by incremental backup.
Recovery ISO	For Windows servers, you must install CRISP software and create a separate ISO. For Linux servers, Cristie provides pre-build recovery ISO.	Cristie provides pre-build recovery ISO.
Recovery Workflow	Initiate BMR recovery from Cohesity to fetch the recovery path (the SMB share has the “system.vtd” file, the Custom Virtual Device Driver file). Then mount the system.vtd file in the CBMR recovery wizard to complete the recovery.	Recovery is performed through CoBMR Recovery Environment.

CBMR

CBMR is required and must be licensed and installed separately from Cohesity. Cristie can be used for all supported platforms, but Cohesity's Cristie bare machine recovery (BMR) integration is supported for Windows and Linux. For Cristie product documentation, see CBMR Documentation.

Setup Protection with CBMR and Cohesity

The steps involved in setting up protection with CBMR and Cohesity are as follows:

- Plan and prepare the server for Cristie-Cohesity integration. For more information, see Plan and Prepare for CBMR Protection.
- Register your server as a Cohesity source. For more information, see Register the Server as a Cohesity Source for CBMR Protection.
- Create a block-based Protection Group for the server. For more information, see Protect the Server Using CBMR.
- Recover your server using the CBMR software. For more information, see Recover the Server Using CBMR.

The following figure illustrates the CBMR workflow:

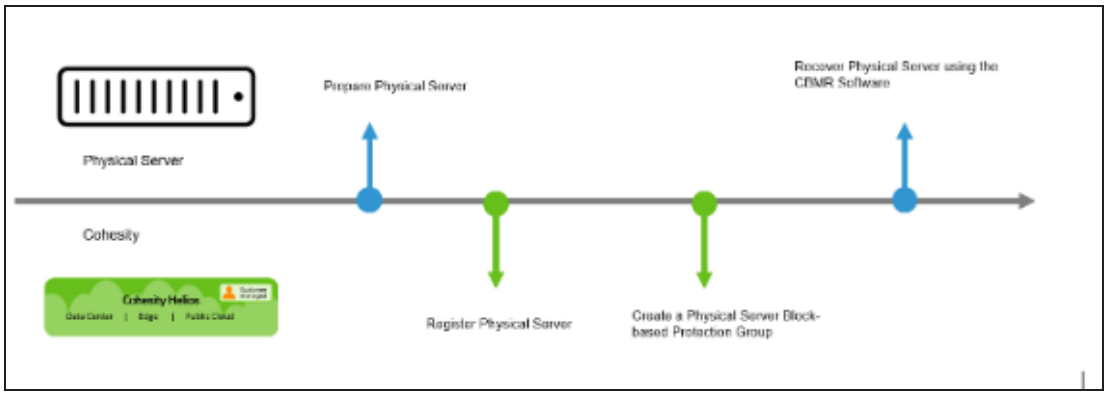

Plan and Prepare for CBMR Protection

Before you register and protect your server for bare machine recovery, ensure that:

- You register as a user with Cristie and activate your Cristie license before your fully functional 30-day trial period expires
- Cohesity supports the server marked for BMR protection. For more information, see Supported Software.
- The firewall ports required for the communication between:
  - The Cohesity cluster and the server are open.
  - The Cohesity cluster and the recovery environment are open.

  For more information, see Firewall Ports Used for Communication.
- You have installed the Cohesity physical agent and CBMR software on the server. For more information, see Prepare the Server for Cristie-Cohesity Integration.

Supported Software

For more information on supported CBMR versions, see Supported Software.

Firewall Ports Used for Communication

Ensure all necessary ports are open in the firewall between Cohesity cluster and VMware VMs and Hypervisors. For more information, see Firewall Ports Used for Communication.

Prepare the Server for Cristie-Cohesity Integration

You must install and run the following software on the physical server you want to protect for BMR:

- Install Cohesity Agent
- Install CBMR

Install Cohesity Agent

Installing the Cohesity Agent on the server helps to establish the network path between the system image backup and Cohesity.

- For information on installing the Cohesity Agent on Windows servers, see Install and Manage the Agent on Windows Servers.
- For information on installing the Cohesity Agent on Linux servers, see Install and Manage the Agent on Linux Servers.
- For information on deploying Cohesity Agent on VMs to protect them as physical servers, seeDeploy Cohesity Agent on Virtual Servers.

Install CBMR

To enable the physical server for bare machine recovery, you must install the CBMR software on the server. You can download the CBMR installer to the server from the CBMR website. Registration is required to access the Cristie portal. For more information, see Cristie Portal Registration.

Install CBMR on Windows Server

To install CBMR on a Windows server:

1. On the Windows server, run the CBMR installer you downloaded.

   The **InstallShield Wizard for CBMR VA Suite <version>** page appears.
2. Select the language for the installation and click **Next>**.

   The **Welcome** page appears.
3. Select **Install CBMR Suite**.

   The **License Agreement** page appears.
4. Read the license agreement carefully and select the **I accept the terms in the license agreement** radio button to proceed with the installation steps and click **Next>**.

   The **Select install component** page appears.

5. Select **Install CBMR** to install the CBMR components.
6. Click **Install** to install the CBMR component. By default, the CBMR software is installed at C:\Program Files\Cristie\. Optionally, you can select a different location by clicking **Change**.

   > **Note:** Cristie recommends installing the software in the default installation location.

   You can also enter the IPv4 address of a Cristie Virtual Appliance if you want the CBMR to be managed by the virtual appliance.

When the installation completes, click **Finish** and reboot the server. For more information, see CBMR Windows Installation and Licensing Guide.

Install CBMR on Linux Server

To install CBMR, you must have the login credentials with root/sudo for the user on the Linux server.

To install CBMR:

1. Copy the downloaded installation file to a temporary location, for example, '/tmp.

   The CBMR installation binaries are available in three formats: RPM package (for RHEL & CentOS), Debian package (for Ubuntu), and Tar file (to install from the source).

   - To install CBMR from an RPM package, run the `rpm -ivh --nodeps` command.

     For example:

     ```
     rpm -ivh --nodeps cbmr-8.7.1008-1.x86_64.rpm
     ```

   - To install CBMR from a Debian package, run the `dpkg` command.

     For example:

     ```
     dpkg -i cbmr-8.7.1008-1_amd64.rpm
     ```

   - To install CBMR from a Tar file:
     - Extract the Tar file using the `tar xvzf` command.

       For example:

       ```
       tar xvzf cbmr-8.7.1008-1.linux.x86_64.tar.gz
       ```

     - Navigate to the extracted folder and run the following command to install CBMR:

       ```
       ./install
       ```

2. Check the status of the installation in the `/CBMRCFG/cbmrcfg.log` file.

3. Once the installation is complete, run the `licmgr` command to check the license and validate the CBMR installation.

   For example:

   ```
   licmgr -p cbmr
   ```

4. Copy the `ubax_monitor` file inside the `tar` package to `/usr/bin` to monitor the BMR protection status.

   For example:

   ```
   # cd cbmr
   ```

   ```
   cbmr]# cp ubax_monitor /usr/bin
   ```

For more information, see CBMR Linux CBMR Installation and Licensing Guide.

Register the Server as a Cohesity Source for CBMR Protection

For performing CBMR, you must first protect your server. To protect your server data, you need to register the physical server as a source on the Cohesity cluster. For more information see, Register or Edit a Physical Server.

You can also register VMware VMs and deploy the Cohesity agent on VMs to protect the VMs as physical servers and perform CBMR.

For information on registering VMware VMs, see Register or Edit a Hypervisor Source.

For information on deploying Cohesity agent, see Deploy Cohesity Agent on Virtual Servers.

Protect the Server Using CBMR

You can protect a server using a Protection Group and a Protection Policy.

A Protection Policy is a reusable set of settings that define how and when objects are protected and replicated.

A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

When creating a Protection Group for protecting your server for CBMR, you must ensure to:

- Create a block-based Protection Group.
- Enable the **Take BMR backup** option for all the volumes you select to protect.
- If you want to back up all the volumes in a system, ensure to enable **Protect all local volumes** while selecting the physical servers for protection.
- Attach a Protection Policy with the **Add BMR Backup (Physical Server)** option configured to the Protection Group.

Recover the Server Using CBMR

To perform a CBMR from a Cohesity protected physical server, you require the following:

- A CBMR recovery ISO having the Cohesity agent and Cristie software installed.

  For Windows Server, you can use the Windows Preinstallation Environment (WinPE) 10 based recovery media.

  For Linux Server, you can download the recovery ISO from the Cristie website.
- A recovery target you can:
  - **Recover the server to the original machine**: You can choose this option if a malware attack deletes all the data on your server but the server is still accessible.
  - **Recover the server to a new machine**: You can choose this option if your server becomes unavailable due to a complete hardware failure.
  - **Recover the server as a VM**: You can choose this option if your physical server becomes unavailable due to a hardware failure.

For information on recovering Windows Server with CBMR, see the **Recover Your Server Using Cohesity with CBMR** section in the Bare Machine Recovery Using Cohesity & Cristie Solution Guide.

For information on recovering Linux Server with CBMR, see the **Recover Your Server Using Cohesity with CBMR** section in the Bare Machine Recovery for Linux Using Cohesity & Cristie.

CoBMR

Cohesity Bare Machine Recovery (CoBMR) is a licensed BMR software suite. Cohesity integrates with CoBMR to perform the bare machine recovery of the physical servers (Windows and Linux) by leveraging Cohesity's file-based backup and block-based backup for physical server protection.

With CoBMR, the system state information of the physical server is captured first, followed by the file-based or block-based backup for all drives, volumes, and partitions. This way, CoBMR recovers the physical server's original state using the recovery image and the protected data.

To obtain the CoBMR software and license, contact Cristie.

Setup Protection with CoBMR and Cohesity

Steps involved in setting up protection with CoBMR and Cohesity are as follows:

- Plan and prepare the server for Cristie-Cohesity integration. For more information, see Plan and Prepare for CoBMR Protection.

- Register your server as a Cohesity source. For more information, see Register the Server as a Cohesity Source for CoBMR Protection.
- Create a file-based or block-based Protection Group for the server. For more information, see Protect the Server using CoBMR.
- Recover your server using the Cristie CoBMR software. For more information, see Recover the Server Using CoBMR.

The following figure illustrates the CoBMR workflow:

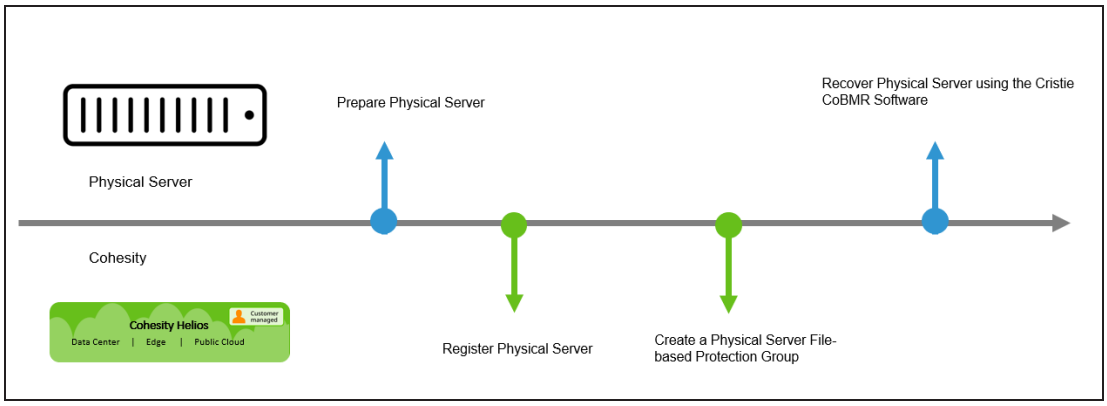

Plan and Prepare for CoBMR Protection

Before you register and protect your server for bare machine recovery, ensure that:

- Cohesity supports the physical server marked for BMR protection. For more information, see Supported Software.
- The firewall ports required for the communication between:
  - The Cohesity cluster and the server are open.
  - The Cohesity cluster and the recovery environment are open.

  For more information, see Firewall Ports Used for Communication.
- You have installed the Cohesity physical agent and CoBMR software on the physical server. For more information, see Prepare the Server for Cristie-Cohesity Integration.
- You have added the executable file in the **user_scripts** directory to capture the system state. For more information, see Add Executable File in user_scripts Directory.
- You have reviewed and understood the considerations of the CoBMR process. For more information, see Consideration.

Supported Software

For more information on supported CoBMR versions, see Supported Software.

Firewall Ports Used for Communication

Ensure all necessary ports are open in the firewall between the Cohesity cluster and VMware VMs and Hypervisors. For more information, see Firewall Ports Used for Communication.

Prepare the Server for Cristie-Cohesity Integration

You must install and run the following software on the physical server you want to protect for BMR:

1. Install Cohesity Agent.
2. Install CoBMR Software.

Install Cohesity Agent

Installing the Cohesity agent on the server helps to establish the network path between the system image backup and Cohesity.

For information on installing the Cohesity agent on Windows servers, see Install and Manage the Agent on Windows Servers.

For information on installing the Cohesity agent on Linux servers, see Install and Manage the Agent on Linux Servers.

For information on deploying Cohesity agent on VMs to protect then as physical server, see Deploy Cohesity Agent on Virtual Servers.

Install CoBMR Software

> **Note:** The CoBMR installation can take some time so you may need to plan accordingly.

To enable the server for bare machine recovery, you must install the CoBMR software on the server. You can download the CoBMR installer to the server from the CoBMR site. Registration is required to access the Cristie portal. For more information, see Cristie Portal Registration.

> **Note:** To use CoBMR in DoDIN mode, contact Cristie team to obtain a suitable version of the CoBMR software.

Install CoBMR on Windows Server

To install CoBMR on a Windows server:

1. On the Windows server, run the CoBMR installer you downloaded.

   The **Welcome to the InstallShield Wizard for CoBMR VA Suite <version>** window appears.

2. Select **Install CoBMR Suite**.

   The **License Agreement** window appears.

3. Read the license agreement carefully and select the **I accept the terms in the license agreement** radio button to proceed with the installation steps and click **Next>**.

   The **Select install component** window appears.

4. Select **Install CoBMR** to install the CoBMR components.

5. Click **Install** to install the CoBMR component. By default, the CoBMR software is installed at `C:\Program Files\Cristie\`. Optionally, you can select a different location by clicking **Change**.

   > **Note:** Cristie recommends installing the software in the default installation location.

At the end of the installation, detailed machine configuration information is recorded in the **CoBMRcf.exe** executable file located at `:\Program Files\Cristie\CoBMR` on the Windows System drive. This executable captures the system recovery data.

Install CoBMR on Linux Server

For Linux servers, Cristie shares the CoBMR installation binary in the following formats:

- RPM package for RHEL & CentOS.
- DEB package for ubuntu machines.
- As gunzip tar file to install from the source.

These packages can be installed on the Linux servers using the command line. For example, the rpm package can be installed using the following command.

```
# rpm -i cobmr-9.1-1.x86_64.rpm
```

At the end of the installation, detailed machine configuration information is recorded in the cobmrcfg executable saved in `the /usr/bin` directory. This executable captures the system recovery data.

Add Executable File in user_scripts Directory

Once the CoBMR installation is complete, the **cobmrcfg** executable gets installed at the following location:

- For Windows Server: `\Program Files\Cristie\CoBMR`
- For Linux Server: `/usr/bin/`

This executable file must be run before every protection run to capture the latest System Recovery Data information. You can run this executable as a pre-script in the Cohesity's Protection Group by adding this executable file in the **user_scripts** directory available in

the Cohesity agent installation directory on the physical server. This pre-script captures the system state as a system recovery data file.

To add the executable in the user_scripts directory:

For Windows Server:

Create a batch file to trigger **CoBMRcfg.exe** from its installed location.

For example:

```
echo "C:\Program Files\Cristie\CoBMR\CoBMRCfg.exe" > c:\"Program Files"\Cohesity\user_scripts\<file_name>.bat
```

> **Note:**
> - You must execute the echo command using MS-DOS prompt. Executing the echo command using PowerShell inserts a special character causing the command to fail.
> - You must not copy the **CoBMRcfg.exe** script to the **INSTALL_DIR/user_scripts** directory.

For Linux Server:

Create a shell script under the `<cohesity_agent_install>/user_scripts` directory.

Within the shell script, invoke `/usr/bin/cobmrcfg`.

In the pre-script, you have to mention the path as the script name.

Consideration

When recovering a Linux server using CoBMR, the option for verifying the SSL certificate validity is auto-enabled. However, you must disable this option for achieving seamless recovery.

Register the Server as a Cohesity Source for CoBMR Protection

For performing CoBMR, you must first protect your server. To protect your physical server data, you need to register the server as a source on the Cohesity cluster. For more information see, Register or Edit a Physical Server.

Protect the Server using CoBMR

Protecting a physical server for CoBMR involves creating a **Protection Group** and a **Protection Policy**.

- A **Protection Policy** defines how and when objects are backed up, archived, or replicated. For details on creating a policy, see Create or Edit a Standard Policy.
- A **Protection Group** applies the schedules and settings defined in the policy to execute backups. For instructions, see Add or Edit a Protection Group for Physical Servers.

When setting up a Protection Group for CoBMR, follow the platform-specific guidelines below:

For Windows Servers

1. **Choose Backup Type**: Select either a file-based or block-based Protection Group.

   > **Note:** Recovery steps vary by backup type. File-based backups do not require local SMB user configuration, while block-based backups do. Refer to the CoBMR Windows User Guide for details.

2. **Include All Volumes**: If backing up all system volumes, enable Protect all local volumes during server selection.
3. **Enable Crash Consistent Backups**: Check the Crash Consistent Backups option in the Protection Group settings.
4. **Add Pre-Script**: Add CoBMRcfg.bat as a Pre Script in the Advanced Settings. In the Script Path field, specify the batch file that triggers CoBMRcfg.exe to capture the system state.

   For instructions on adding executable files, see Add Executable File in user_scripts Directory.
5. **Set Timeout**: Set Timeout (mins) to at least 45 minutes to account for script execution time.
6. **Disable Script Failure Option**: Uncheck Continue Backup if script fails.

For Linux Servers

1. **Choose Backup Type**: Select either a file-based or block-based Protection Group.

   > **Note:** Recovery steps vary by backup type. Refer to platform-specific documentation for details.

2. **Include All Volumes**: If backing up all system volumes, enable Protect all local volumes during server selection.
3. **Add Pre-Script**: Configure a pre-script that invokes cobmrcfg with the --cobmr_boot_backup flag to ensure the /boot partition is backed up correctly.

   For guidance on adding scripts, see see Add Executable File in user_scripts Directory.

4. **Set Timeout**: Set Timeout (mins) to at least 45 minutes to account for script execution time.
5. **Disable Script Failure Option**: Uncheck Continue Backup if script fails.

Recover the Server Using CoBMR

To perform a CoBMR from a Cohesity protected physical server, you require the following:

- A CoBMR recovery media (ISO) provided by Cristie having the Cohesity agent and Cristie software.
- A recovery target. you can:
  - **Recover the server to the original machine**: You can choose this option if a malware attack deletes all the data on your server but the server is still accessible.
  - **Recover the server to a new machine**: You can choose this option if your server becomes unavailable due to a complete hardware failure.
  - **Recover the server as a VM**: You can choose this option if your physical server becomes unavailable due to a hardware failure.

  During the recovery process, the CoBMR recovery environment connects to the Cohesity cluster to fetch the source details and the desired snapshot from Cohesity. Then, the discovery process starts to fetch the system data file that captures the system state and restores the system state. This is followed by restoring the boot drives and other mounted drives.

Minimum Permissions

The following table lists the minimum permissions required for CoBMR recovery.

Roles	Permissions
Cluster Management	View Alert Details
Data Protection	View Protection Groups Manage Sources
Recovery Management	View Recover Tasks Manage Recover Tasks
Source Management	Manage Physical Servers

Recover Windows Server

> **Note:** For information on Cristie Windows Server requirements, see CoBMR for Windows.

To recover the server using CoBMR, perform the following steps:

1. Boot the recovery ISO to start the CoBMR Recovery Environment.

   Depending on the recovery target speed and the ISO recovery media, the CoBMR window might take some time to appear.

   The **CoBMR Recovery Environment** window appears. Click **Recover System**.
2. (Optional) In the **Log Files Path** field, specify the path either as a share in UNC format or as a drive letter and path where the log files must be saved at the end of the recovery. By default, the log files are saved in the /COBMRCFG on the root drive.
3. Specify the **Domain, Username**, and **Password** of the UNC path to which the log files must be saved.
4. Click **Next**.
5. In the **Cohesity Server** field, enter the IPv4 address or the FQDN of the Cohesity cluster.
6. In the **Authentication** field, enter the Domain name, user name, and password of the Cohesity cluster. For more information on the permissions required for authentication, see Minimum Permissions.
7. In the **Registered Source** field, enter the IPv4 address or the FQDN of the physical server source as registered with the Cohesity cluster.
8. Enable **Point-in-time (PIT)** restore and select a time slot corresponding to the snapshot you want to restore. The CoBMR software will select the latest snapshot up to the time specified.
9. Click **Next**.

   The Window displays the status of the Cristie software attempting to connect to the Cohesity cluster, registering the recovery environment with the Cohesity cluster, and then extracting the Cohesity backup snapshots. The **Next** button is enabled once the BMR configuration is extracted.
10. Click **Next**.

    The CoBMR - Volume Layout window appears that displays a list of disks and partitions to be recovered. Depending on the volume layout (similar or dissimilar hardware) of the target server, you can perform one of the following actions:

    - **Similar Volume Layout**: If the recovery is back to the original server or a different server with the same volume layout (similar hardware), the disk

mapping is similar. You can select the Proposed layout and click **Next**.

- **Dissimilar Volume Layout**: When recovering to a different server with dissimilar hardware, the mapping is complex with the different disk geometry and capacity. In such cases, you can right-click the disks displayed in **Proposed layout** to select the disk to be partitioned or not. Then click **Next**.

The **CoBMR - Select Volumes to Restore** window appears.

11. Select the datasets to be restored from the backup and click **Next**.

    The **CoBMR - Clone Settings** page appears.

12. Change the NeTBIOS name, host name, or IP address of the target server. You can use DHCP or enter a valid static IP address. Click **Finish**.

    The **Restore Status** window displaying the progress through the recovery process appears.

    The CoBMR recovery process transfers the system data and user volumes that were backed up on Cohesity and restores them on the target server with the system state maintained as that of the original server.

Recover Linux Server

> **Note:** For information on Cristie Linux Server requirements, see CoBMR for Linux.

To recover the server using CoBMR, perform the following steps:

1. Boot the CoBMR recovery environment from the ISO recovery media. You can recover back on the same physical server or onto a new physical server.

   The **CoBMR Recovery Environment** window appears.

2. Read and accept the end-user license agreement (EULA) to proceed with the recovery process.
3. Select **CoBMR** from the **Choose which product to recover** drop-down and then click **OK** to continue.
4. From Cristie Software menu options, select **Automatic Recovery**.

   The **Setup Cohesity VE Server information** window appears.

5. Enter the following information for accessing the Cohesity cluster:
   - IPv4 address or the FQDN of the Cohesity cluster
   - (Optional) Port number
   - Domain name
   - User name
   - Password

> **Note:** The **Verify SSL certificate validity** option is enabled by default. However, you must disable this option to achieve seamless recovery.

6. Click Network Setup and validate the network setup and click **OK**.
7. Click **Next**.

   The **Get Configuration** window appears.
8. In the **Select Source** field, select the Linux server you want to recover.
9. In the **Select Backup** field, select a date and a time slot corresponding to the snapshot you want to restore. The CoBMR software will select the latest snapshot up to the time specified.
10. Click **Next**.

    The **Copy Configuration From Backup** window displays the status of the Cristie software downloading and installing the Cohesity Agent, connecting to the Cohesity cluster, registering the recovery environment with the Cohesity cluster, and then extracting the Cohesity backup snapshots. The **Next** button is enabled once the BMR configuration is extracted.
11. Click **Next**.

    The **Start Recovery** window appears.
12. Specify the restore settings and then click **Next**.

    The **Recovery Status** window appears.
13. Once the recovery is complete, a pop-up window will appear to remind you to copy the log files. Click **OK**.
14. From the Cristie Software menu options, click **Reboot** to restart your system. Once the system restarts, validate the recovered volumes.

The CoBMR recovery process transfers the system data and user volumes that were backed up on Cohesity and restores them on the target server with the system state maintained as that of the original server.

Chapter 24: Manage External Targets

The Cohesity cluster can move rarely used and inactive data to a designated external target, such as tape or cloud storage.

From the **External Targets** page, you can view, manage, and unregister the external targets that have been registered on the Cohesity cluster. For detailed instructions, see Register an External Target.

From the **Targets** tab of the external targets page, the following fields and options are available:

Column Name	Description / Action
Name	Drill-down on this link to view the following details about the target.
Purpose	The target's purpose, either Archival or Cloud Tier.
Type	The target's type, such as Google Nearline, AWS Glacier or Azure Standard.
Actions	To manage an existing target, use the following options: • **Edit** ( )–Configure the settings of the existing Target. For detailed instructions, see Register an External Target. • **Unregister**–You can unregister an external target. By unregistering, the associated metadata will be deleted from the target and will no longer be available for retrieval. For details, see Unregister an External Target.

From the **Settings** tab of the of the external targets page, you can:

- Register an External Target for Archival.
- Set Bandwidth throttling at Cluster-level for Tiering to External Targets.

CloudArchive Direct

CloudArchive Direct lets you archive data directly to an external target without retaining a copy of the data on the Cohesity cluster. In CloudArchive Direct, the Cohesity cluster acts as a local cache to stream the dataset to the external target, eliminating the need to store a full backup copy on the Cohesity cluster first.

This approach reduces the capacity requirements on the local Cohesity cluster to a small cluster footprint, as only the metadata and index, and not the dataset itself are stored on the local cluster. Storing the metadata and index on the local cluster enables much quicker search and recovery down the road. The external target also stores the metadata and index, in addition to the data itself. On the Cohesity cluster, you can register a cloud storage, S3-compatible storage, or NFS-mounted servers as external targets.

> **Note:** For the first archival, depending on the size of data, the Cohesity cluster might experience significant capacity utilization, because the data first gets stored on the local storage (Storage Domain) and then aggressively gets down tiered to the external target. During this time, the capacity utilization will increase till the data is garbage collected upon the successful transfer.

> **Important:**
> The option to enable CloudArchive Direct (native mode - no dedupe, no compression) from the NAS Protection Group page is deprecated in the 6.7 version.

How CloudArchive Direct Works

When you run a Protection Group for a registered source with CloudArchive Direct enabled, then according to the policy set for the Protection Group, the file system from NAS is read and stored on the external target as a streaming backup using the Cohesity cluster as a temporary cache.

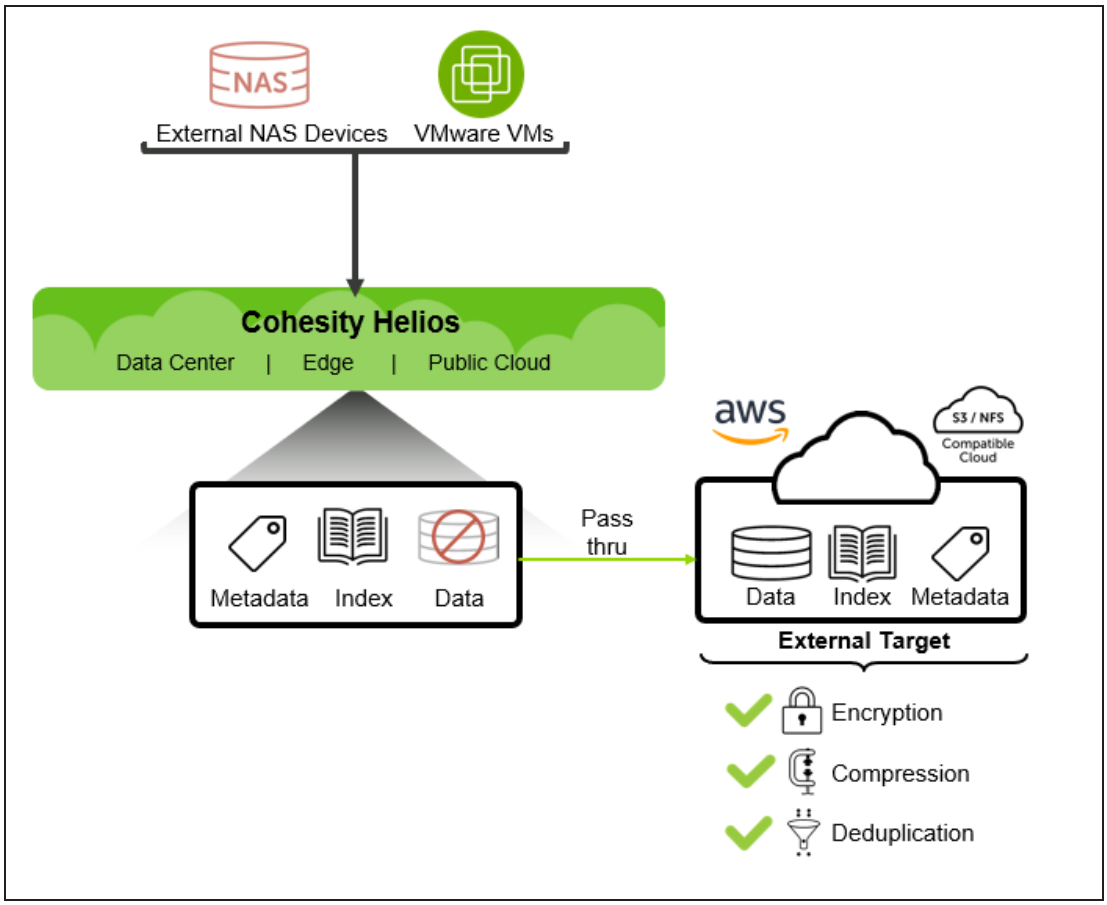

CloudArchive Direct Characteristics

Feature	CloudArchive Direct
Supported Sources	• VMware • NAS sources • S3-compatible
Mode of enablement	While creating or editing a Protection Policy.
Deduplication	Supported
Number of Objects per Protection Group	Multiple
Data archival	Incremental forever

Feature	CloudArchive Direct
Metadata archival	Incremental forever
Granularity level of incremental archival	Block-level–Need to perform only incremental archival of a file if the data or metadata of the file is updated.
Native format as primary	No
Tiering of objects using Cohesity Data Movement	Supported–can choose to tier the archived data to a cheaper storage tier.
CloudRetrieve	Supported
Vault unregistration	Not supported

Supported External Targets

For the list of supported external targets, see Supported Software.

Supported Workflows for CloudArchive Direct

The following table lists the workflows that are supported workflows for the sources, based on the external target (Standard Storage, Archive Storage, and Hot tiers) to which you CloudArchive Direct:

CloudArchive Direct Workflow	VMware			NAS			S3-Compatible
	Standard Storage Tier	Archive Storage Tier	Hot Tier	Standard Storage Tier	Archive Storage Tier	Hot Tier	Non-Archive Storage Tier
Direct Archival	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Restore	Yes	Yes	Yes	Yes	Yes	Yes	Yes
File-level restore	Yes	No	Yes	Yes	Yes	Yes	No**
Indexing	Yes*	No	Yes	Yes	Yes	Yes	No
CloudRetrieve	Yes	Yes	Yes	Yes	Yes	Yes	Yes

*Indexing for VMware VMs in CloudArchive Direct-enabled Protection Groups may cause data egress from the external target, resulting in charges from the cloud provider.

**Cohesity supports restore of S3-compatible only at prefix-level.

Differences Between CloudArchive Direct and Cloud Archive

Term	Description
CloudArchive Direct	CloudArchive Direct copies the data from the supported source directly to any registered external target such as public cloud providers (AWS, Azure, Google Cloud Platform) or any S3-compatible external target to reduce storage costs and eliminate the need to store a full backup copy on-premises.
Cloud Archive	Cloud Archive sends a copy of the backup data (it contains the backup data, metadata, dedupe fingerprint, and index) stored on the Cohesity Cluster to any registered external target. To perform Cloud Archive, at least one full copy of the primary data needs to be stored on the cluster.

Figure X: CloudArchive Direct and Cloud Archive workflow

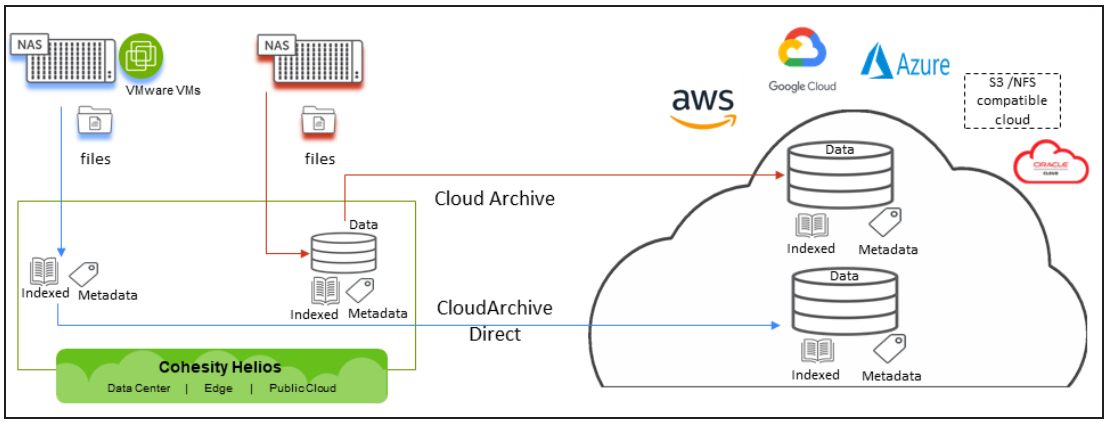

Considerations

- CloudArchive Direct is supported for both physical clusters and virtual clusters. CloudArchive Direct is not supported for Cohesity Cloud Edition (legacy). Cohesity Cloud Edition (formerly known as NGCE cluster) uses a CloudArchive Direct architecture, in which backup data is directly down tiered to object storage, and only metadata is stored on the cluster's cache tier.
- If **Encryption** is enabled at the cluster level, you must also enable encryption while registering the external target for CloudArchive Direct. Archiving directly to an external target may fail if encryption is enabled at the cluster level but not at the

external target level.

- Review and understand the following considerations if you are performing CloudArchive Direct of NAS sources:
  - When you upgrade the Cohesity cluster from 6.6 or below versions to the latest version, the existing Protection Group created with CloudArchive Direct native mode will not automatically switch to the new CloudArchive Direct format. Therefore, you must pause the existing Protection Group and then create a new Protection Group for the objects with CloudArchive Direct new format.
- Review and understand the following considerations if you are performing CloudArchive Direct of VMware VMs to **Archive Storage tier** (AWS S3-Glacier, AWS S3-Glacier-Deep-Archive, AWS S3-Glacier-Deep-Archive, AWS S3-Glacier-Deep-Archive Gov):
  - File and folder recovery is not supported.
  - Instant recovery of VMware VMs is not supported.
  - Instant Volume Mount is not supported.
  - Test & Dev (Cloning) is not supported.
  - Virtual Disks (VMDK) recovery is not supported.
  - Indexing is not supported.

Enable CloudArchive Direct

This topic provides information on how to enable CloudArchive Direct.

To enable CloudArchive Direct, perform the following steps:

1. Register any of the following sources:
   - NAS

     For more information, see Register or Edit NAS.
   - VMware Vcenter

     For more information, see Register or Edit a Hypervisor Source.
2. Register an external target with the **Archival Format** option as **Incremental Forever**. For more information, see Register an External Target for Archival.
3. Create a Policy with CloudArchive Direct Settings.
4. Create a Protection Group and Attach the CloudArchive Direct Policy.

Create a Policy with CloudArchive Direct Settings

A policy is a reusable set of settings that define how and when objects are protected, replicated, and archived. You can select which standard policy to use when configuring a Protection Group. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

For CloudArchive Direct, you must create a policy with the following settings:

1. Navigate to **Data Protection** > **Policies**.
2. Click **Create Policy** located at the top right of the page.

   To edit an existing policy in the list, click the actions menu (⋮) and select **Edit** (✎).
3. **Policy Name**: Specify a protection policy name. The name can contain alphanumerics, underscores, hyphens, periods, and spaces. This field is required and can be changed later.
4. Specify when backups are captured:
   - **Backup every xx**: Select the time unit for how often backups are captured by the Protection Group. This defines the protection schedule for capturing backups.
5. Under **Primary Copy**, specify the following:
   - **Keep On**: Select the external target that you have registered. You can also register a new target by clicking **Register External Target**.

     > **Note:** For CloudArchive Direct, you should not keep the primary copy on the Cohesity cluster (Local copy).

   - **Retain for xx days**: Specify how many days backups are retained on the external target before they are deleted.
6. If you want to change or configure any of the additional policy settings, then select **More Options** and perform the below steps, or else, click **Create**.
   1. **Periodic Full Backup**: Select this option from the **Backup Options** floating menu and specify the frequency for periodic full backup:
      - Day
      - Week
      - Month
   2. **Continuous Data Protection**: (*Applies only for VMware*) Select this option from the **Backup Options** floating menu and specify the time unit (hours) to capture and retain logs on the Cohesity cluster. This value defines the CDP schedule for capturing logs from the VMs.

      You can use this option to restore the VM applications to a point-in-time as opposed to periodic snapshots available with regular backups.
   3. **Quiet Times:** Select this option from the **Backup Options** floating menu to add a quiet time. Quiet Times define time periods when new protection runs are not started. For example, you may want to configure hourly backups that run

during weekdays but not on Saturdays or Sundays. Quiet Times only prevent new protection runs from starting during the specified time period. Protection runs that started before the quiet time period are not affected and will continue to run. For example, if a protection run starts at 12:45 AM on Saturday morning and there is a weekend quiet time period that starts at 1 AM on Saturday morning, the protection run will continue to run past 1 AM. No new protection runs are started during the quiet times.

You can add multiple quiet times. For example, you could define a quiet time period for Saturday and Sunday and another quiet time period every morning from 9:00 AM to 11:00 AM.

> **Note:** If you update the Quiet Time periods in a policy for a running Protection Group, only new protection runs using the policy are affected and any existing Protection runs that have already started will use the original Quiet Time periods.

4. **Retry Options**: By default, the Cohesity cluster attempts to capture snapshots three times before the Job fails. The default time between retries is five minutes. You can optionally customize the number of retries and how long to wait between each attempt.
5. (Optional) Click the **Primary Copy** tile and configure the following **Backup Options** from the floating menu:
   - **Extended Retention**: Select this option to retain a subset of snapshots (backups) for a longer time period than defined by the protection schedule. Specify when to take the first snapshot and how long to retain it.
   - **Data Movement**: This option applies only when you have selected AWS Standard Storage Tier as the external target to store the primary copy. With this option, you can move the data (downtier) from the selected AWS Standard Storage Tier to an AWS Archive Storage Tier target after a specified time period. Data Movement helps in space reclamation on the external target. From the **Move Data** drop-down, you can select the data S3 target and specify after how many days, weeks, months, or years you want to move the data from the external target to the S3 bucket.

     > **Note:** As an AWS consideration, you cannot downtier from S3 Glacier to S3 Deep Archive. After downtiering to either S3 Glacier or S3 Deep Archive, you cannot perform any further downtiering.
6. Click **Create**.

Create a Protection Group and Attach the CloudArchive Direct Policy

A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured.

For performing CloudArchive Direct on the supported sources, you must create a Protection Group and attach the policy created with the CloudArchive Direct settings.

- For information on creating a Protection Group for VMware VMs, see Add or Edit a Protection Group for Virtual Servers.
- For information on creating a Protection Group for NAS, see Protect NAS.

Use Cases for CloudArchive Direct

Use-Case 1: Perform CloudArchive Direct to a cold storage tier

1. On the AWS console, create an S3 bucket. Ensure that you do not set lifecycle management policy on the S3 bucket.
2. Log in to the Cohesity cluster, and perform the following steps:
    1. Register the S3 bucket as an S3-Glacier or S3-Glacier-Deep-Archive as an external target. While registering:
        - Select the **Archival Format** as **Incremental Forever**.
        - Download and execute the CloudFormation template in your AWS account to create the required resources for Garbage collection.

        For more information, see Register an External Target for Archival.
    2. Create a policy with the **Primary Copy** option set to the registered external target.

        For more information, see Create a Policy with CloudArchive Direct Settings.
    3. Create a Protection Group and Attach the CloudArchive Direct Policy. For more information, see Create a Protection Group and Attach the CloudArchive Direct Policy.

Use-Case 2: Perform CloudArchive Direct to hot storage tier (S3) with tiering to a cold storage tier (S3-Glacier or S3-Glacier-Deep-Archive)

1. On the AWS console, create an S3 bucket. You must not set the lifecycle management policy on the S3 bucket.
2. Log in to the Cohesity cluster, and perform the following steps:
    1. Register the S3 bucket as an S3-Glacier or S3-Glacier-Deep-Archive as an external target. While registering:

- Select the **Archival Format** as **Incremental Forever**.
- Execute the Cloud Formation template in your AWS account to create the required resources for Garbage collection.

  For more information, see Register an External Target for Archival.

2. Create a policy with the following settings :
   - The **Primary Copy** option is set to the registered external target.
   - Add the **Data Movement** option in the policy and select the desired storage class to tier the snapshots.

   For more information, see Create a Policy with CloudArchive Direct Settings.
3. Create a Protection Group and Attach the CloudArchive Direct Policy. For more information, see Create a Protection Group and Attach the CloudArchive Direct Policy.

Recover CloudArchive Direct Data

> **Note:** See the Considerations section before recovering the data directly archived to the external target.

For NAS, Cohesity supports both file-based and volume-based recovery of filesystem that is directly archived to the external target.

- For information about file-based recovery, see Recover Files or Folders.
- For information about volume-based recovery, see Recover NAS Storage Volumes.

For VMware VMs, Cohesity supports the recovery of VMs and virtual disks that are archived to the external target.

For more information, see:

- Recover VMs.
- Recover Virtual Disks.

CloudRetrieve of CloudArchive Direct Data

Cohesity supports retrieving the CloudArchive Direct data to an alternate Cohesity cluster. For more information on how to retrieve data to an alternate cluster, see About CloudRetrieve.

Archive Object Lock

> **Note:** This feature applies only to CloudArchive.

Cohesity stores immutable data at both the source cluster and external target levels. By configuring Object Lock, you can lock archives on external targets to prevent data from being modified, deleted, or overwritten. Object Lock can be configured during external target registration. For registering an external target, see Register an External Target for Archival.

Prerequisites for Object Lock

Before registering an external target with Object Lock enabled, ensure the following:

- **Confirm Object Lock support**: Verify that the target you select supports Object Lock. For the list of targets that support Object Lock, see Supported Workflows and External Targets.
- **For AWS external targets**: The Object Lock-supported AWS external target must have versioning and Object Lock enabled. For more information, see Using S3 Object Lock.
- **For Azure external targets**: The Object Lock-supported Azure external target must have version-level immutability enabled on the storage container, which indicates that the storage account must also have version-level immutability. For more information, see Enable version-level immutability support on a container.
- **For GCP external targets**: The Object Lock-supported GCP external target must have versioning and Object Retention Lock enabled. For more information, see Object Retention Lock.

> **Note:**
> - The above properties must be set at the time of target creation and cannot be enabled after creation.
> - Avoid setting default retention on S3 buckets and Azure - this is not required as Cohesity locks the objects for the duration required. Large default retentions affect garbage collection and increase costs.
> - For AWS and S3-compatible targets, by default, the archived objects will be locked in the Governance retention mode. Support for enabling Object Lock in **Compliance** mode via the Cohesity UI will be introduced in an upcoming release.
> - For Azure targets, by default, the objects will be locked in the Locked mode.

Considerations for Object Lock

- If you registered an Object Lock-capable external target on a Cohesity cluster running version 6.8 or earlier and later upgraded the cluster to a version that supports Object Lock, you will not be able to enable Object Lock on those previously registered targets.
- When archiving data to versioned buckets, ensure that you do not edit the storage class of objects directly. Editing the storage class of objects directly will create a new version of the object while the original version remains in the previous storage tier. Object-level locks will not be carried over to the new version.
- The Cohesity Data Movement feature, which can enabled during protection policy creation cannot coexist with Object Lock. An archival target can only have one of the settings enabled.
- The expiry of the object lock will match the expiry set for archived snapshots. However, the object lock will not update if you modify the expiry for archived snapshots after the Protection Group has run.
- You cannot disable Object Lock once you enable it on a target.
- **For archival with periodic full format**:
  - For the list of targets (S3-compatible, AWS, and Azure) that support Object Lock for archival with periodic full,

    see Supported Workflows and External Targets.
  - If Object Lock is not met for an archive, the archive is still considered successful and can be recovered.
  - After archival, to view the Object Lock status, perform the following:
    a. Click the **Protection** menu.
    b. Select the required Protection Group.
    c. Select the required run.
    d. Click the **Cloud Archive** tab.

       The **Object Lock Compliance** status and **Archive Object Lock Expiry Time** will be displayed.
- **For archival with incremental forever format**:
  - For the list of targets (S3-compatible, AWS, Azure, and GCP) that support Object Lock for archival with incremental full, see Supported Workflows and External Targets.
  - If Object Lock compliance is not met for an archive, then the archival will fail.
  - For on-prem clusters, Cohesity does not support Object Lock for CloudArchive Direct. However, for Cohesity Cloud Edition clusters in AWS, GCP, and Azure, where the primary copy of data is down-tiered directly to an external target via

CloudArchive Direct, support Object Lock for CloudArchive Direct of primary data.

- With Object Lock enabled, the storage utilization of external targets may be higher.

Archive Data Over AWS PrivateLink

AWS PrivateLink allows AWS users to access Amazon services (S3, EC2, etc) through a private network. With AWS PrivateLink, you can create an interface endpoint in your VPC. These endpoints are directly accessible from applications that are on premises over VPN and AWS Direct Connect, or in a different AWS Region over VPC peering. For archiving your data to AWS S3, you can configure AWS PrivateLink to securely connect to the AWS external target over your private network.

The following diagram depicts how AWS PrivateLink is used to access Amazon S3 through a private network:

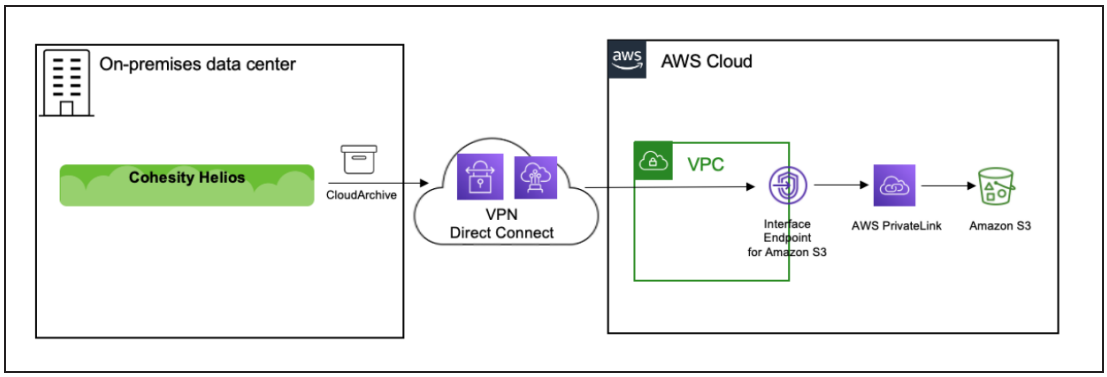

Configuring Amazon PrivateLink for Archive

Prerequisites

Before configuring Amazon PrivateLink for archive, ensure to have:

- AWS Direct Connect or AWS Virtual Private Network (AWS VPN) connection between the on-premise data center and AWS.
- AWS Virtual Private Cloud (VPC), VPC subnet, and security group.

To configure Amazon PrivateLink for archive, perform the following steps:

1. Create a PrivateLink Interface Endpoint for S3 in Your VPC.
2. Create an S3 Bucket in the AWS Console.
3. Register S3 Bucket as an External Target.

Create a PrivateLink Interface Endpoint for S3 in Your VPC

To support on-premise traffic for S3, you must create an interface VPC endpoint. To create an interface endpoint for S3, see AWS documentation.

Once you create an interface endpoint, Amazon S3 generates two types of endpoint-specific S3 DNS names that you can use to communicate with the AWS service from your VPC: *Regional* and *zonal*. Select **Regional S3 DNS name** to communicate with the AWS service.

Create an S3 Bucket in the AWS Console

To archive your data to Amazon S3, you must create an S3 bucket. While creating a bucket, Cohesity recommends you update the S3 bucket policy to restrict access only to the VPC endpoint. For more information on creating an SR bucket, see AWS documentation.

Register S3 Bucket as an External Target

Once you map the VPC endpoint IP address with the FQDN of the S3 bucket, for archiving your data over the private network, you can register the AWS S3 bucket with Cohesity as an external target. For more information, see Register an External Target for Archival.

If you have configured Amazon PrivateLink interface endpoints with the Cohesity cluster to access the AWS S3 external targets for archival, then you must provide values for the following parameters while executing CFT:

- **NetworkSecurityGroup**: The network security group ID to use for the lambda function.
- **VPCSubnets**: One or more subnet IDs in the VPC to use for the lambda function.

For more information, see CloudFormation Template.

Archive Data Over Azure Private Endpoints

This topic provides the requirements and the procedure to configure Incrementation Forever Archival using an Express Route and/or VPN from an on-premises Cohesity cluster to Azure Blob and Azure Function App using Azure Private Endpoints.

Prerequisites

Before configuring Azure Private Endpoint for archival, ensure the following:

- Azure ExpressRoute or Virtual Private Network (VPN) between the on-premises Cohesity cluster and Azure.
- Azure Resources:
    - Azure Resource Group
    - Virtual Network (VNET)
    - VNET subnets for private endpoints (Storage Account and Azure Function App)

- Outbound subnet for Azure Function App
- Network Security Group

Configuring Azure Private Endpoint for Archive

Perform the following configuration steps for archiving data over Azure Private Endpoints. The process includes creating necessary Azure resources and configuring the Cohesity cluster.

1. Create an App Service Plan.
2. Configure Subnet for Outbound Azure Function App.
3. Create an Azure Storage Account with Private Endpoint.
4. Create an Azure Function App with Private Endpoint.
5. Configure Azure Function App for Incremental Forever Archival
6. Configure the Cohesity cluster to use Azure Private Endpoint.

Create an App Service Plan

The App Service Plan defines the compute resources for your Azure Function App.

Perform the following steps to create an App Service Plan:

1. In the Azure Portal, use the **Create App Service Plan** wizard to create the App Service Plan. Enter the following values:
   - **Operating System**: Linux
   - **Pricing Plan**: Select Basic, Standard, Premium, Isolated, or Elastic to support Azure Private Endpoint for App.
2. Click **Next: Tags** to optionally add tags to the App Service Plan.
3. Click **Review + create**.
4. Click **Create** to provision the App Service Plan.

Configure Subnet for Outbound Azure Function App

Perform the following steps to configure a subnet for outbound traffic from your Azure Function App:

1. In the Azure Portal, select the Virtual Network (VNET).
2. On the **Virtual Network** page, select **Subnets** from the left navigation under **Settings**. On the Subnets page, click **+ Subnet**.
3. On the **Add Subnet** screen, enter or select the required values for the subnet settings:

- **Subnet Delegation**: Set to **Microsoft.Web/serverFarms** to enable outbound connectivity for the Azure Function App.

4. Click **Add** to create the subnet.

Create an Azure Storage Account

Create an Azure Storage Account with a private endpoint to archive data to Azure by following these steps:

1. In the Azure Portal, use the **Storage Account** wizard to create the storage account. Enter or select the required values for the storage account settings, then click **Next**.
2. In the **Advanced** tab, configure the advanced options as needed, then click **Next**.
3. On the **Networking** tab, configure the network settings for the storage account, then click **Next**:
   - **Network Connectivity**: Select the appropriate network access option. Refer to the Microsoft documentation for details.
   - **Add a Private Endpoint**: Enter or select the required values, then click **OK**.

     > **Note:** : The private endpoint must be included in DNS. You can use Azure DNS or your own DNS.

4. Configure **Data Protection** (optional), **Encryption**, and **Tags** (optional), then click **Review + Create**.
5. On the **Review + Create** tab, verify your settings and click **Create** to create the storage account.

Create Azure Function App

To archive data to an Azure target using the Incremental Forever archival format, you must create an Azure Function App in the Azure Portal. The Azure Function App enables cost-efficient garbage collection.

Perform the following steps to create an Azure Function App with private endpoints:

1. In the Azure Portal, use the **Create Function App** wizard to create a new Azure Function App. Choose either **App Service** or **Functions Premium** hosting plan, and then click **Select**.

   > **Note:** : Consumption or Flex Consumption hosting plans are not supported with private endpoints.

2. On the **Basics** tab, enter or select the following values, and then click **Next: Storage**:

Setting	Value
Try a secure unique default hostname	Toggle off
Do you want to deploy code or container image	Code
Operating System	Linux
Runtime stack	Python
Version	3.9 or newer
Linux Plan	Select the Azure Service Plan that you created.

3. On the **Storage** tab, select the storage account you created, and then click **Next: Networking**.
4. On the **Networking** tab, enter or select the following values:

Setting	Value
Enable public access	Off
Virtual Network	Select the VNET
Enable private endpoints	On
Private endpoint name	Enter name
Inbound subnet	Select subnet
DNS	Select Azure DNS or your own DNS
Enable VNET integration	On
Outbound subnet	Select the outbound subnet created previously.

5. (Optional) On the **Monitoring** tab, configure monitoring settings, and then click **Next: Durable Function**.
6. (Optional) On the **Durable Functions** page, select the desired option, and then click **Next: Deployment**.
7. On the **Deployment** tab, configure **Basic authentication**, then click **Next: Authentication**.
8. On the **Authentication** tab, enter or select the following values, and then click **Next: Tags**:

Setting	Authentication type
Host storage (AzureWebJobsStorage)	Secrets

9. (Optional) On the **Tags** tab, enter values, and then click **Next: Review + create**.
10. On the **Review + create** tab, review the configuration, and then click **Create**.

Configure Azure Function App for Incremental Forever Archival

Perform the following steps to configure environment variables for your Azure Function App:

1. In the Azure Portal, select the Azure Function App.
2. On the **Function App** page, select Environment variables from the left navigation under **Settings**.
3. In **App settings**, click **Add** to add the required environment variables, then click **Apply**:

Name	Value	Notes
SCM_DO_BUILD_DURING_DEPLOYMENT	true	Required
ENABLE_ORYX_BUILD	true	Required
WEBSITE_DNS_SERVER	<IP address>	Optional - Set this variable to use Azure DNS for DNS resolution in the Azure Function.

4. Click **Confirm** to save the changes.

Configure the Cohesity Cluster to Use Azure Private Endpoint

To use the Azure Storage Account and Azure Function App through their Private Endpoints for archiving data, you must add the Private Endpoint IP addresses to the host mapping on the Cohesity cluster.

Perform To add host mapping for the VPC endpoint on the Cohesity cluster:

1. Navigate to **Settings > Networking > Host Mapping**.
2. Click **Add Host Mapping**.
3. In the **IP Address** field, enter the IP addresses of the private endpoint of Azure Storage account and Function App.
4. In the **Hostname** field, enter the FQDNs of the private endpoint of Azure Storage account and Function App.

   Example:

   - FQDN of Azure Storage account: <Azure Storage Account>.blob.core.windows.net
   - FQDN of Function App:

     <Azure Function App Name>.scm.azurewebsites.net

     <Azure Function App Name>.azurewebsites.net

Register Azure Account as an External Target

After mapping the Azure Storage Account and Azure Function App private endpoint IP addresses with their FQDNs, you can register the Azure Storage Account with Cohesity as an external target. This allows you to archive data over the private network using the Incremental Forever archival format.

For more information, see Register an External Target for Archival.

Register an External Target

You can register an external target with the Cohesity cluster for Cloud Archive, CloudArchive Direct, or for Cloud Tiering.

Supported External Target

The following table provides the mapping of external target types that can serve as Cloud Archive, CloudArchive Direct, and Cloud Tier:

Cloud Provider	Class of Storage	Cloud Archive			CloudArchive Direct	Cloud Tier
		With Periodic Full	With Incremental Forever	Always Full		
AWS	Glacier (Legacy) Standard	Supported	Not Supported	NA	Not Supported	Not Supported
	Glacier (Legacy) Gov	Supported	Not Supported	NA	Not Supported	Not Supported
	Glacier (Legacy) C2S	Supported	Not Supported	NA	Not Supported	Not Supported
	S3 Standard	Supported	Supported	NA	Supported	Supported
	S3 Gov	Supported	Not Supported	NA	Not Supported	Supported
	S3 C2S	Supported	Not Supported	NA	Not Supported	Supported
	S3-Glacier Standard (Currently S3 Glacier Flexible Retrieval)	Supported	Supported	NA	Supported	Not Supported
	S3-Glacier Gov	Supported	Supported	NA	Supported	Not Supported
	S3-Glacier-Deep-Archive Standard	Supported	Supported	NA	Supported	Not Supported
	S3-Glacier-Deep-Archive Gov	Supported	Supported	NA	Supported	Not Supported
	S3-IA Standard	Supported	Supported	NA	Supported	Not Supported

Cloud Provider	Class of Storage	Cloud Archive			CloudArchive Direct	Cloud Tier
		With Periodic Full	With Incremental Forever	Always Full		
	S3-IA Gov	Supported	Supported	NA	Supported	Not Supported
	S3-IA C2S	Supported	Not Supported	NA	Not Supported	Not Supported
	S3-Intelligent Standard	Supported	Supported[4]	NA	Supported[4]	Supported
	S3-Intelligent Gov	Supported	Supported[4]	NA	Supported[4]	Not Supported
	S3-One-Zone Standard	Supported	Supported	NA	Supported	Not Supported
	S3-One-Zone Gov	Supported	Supported	NA	Supported	Not Supported
	Snowball Edge[1]	Supported	Not Supported	NA	Not Supported	Not Supported
	S3 Glacier Instant Retrieval	Supported	Supported	NA	Supported	Not Supported
	S3 Glacier Instant Retrieval Gov	Supported	Supported	NA	Supported	Not Supported

Cloud Provider	Class of Storage	Cloud Archive			CloudArchive Direct	Cloud Tier
		With Periodic Full	With Incremental Forever	Always Full		
Azure	Archive Blob	Supported	Supported	NA	Not Supported	Not Supported
	Archive Blob-Gov	Supported	Not Supported	NA	Not Supported	Not Supported
	Cool Blob-Standard	Supported	Supported	NA	Supported	Not Supported
	Cool Blob-Gov	Supported	Not Supported	NA	Not Supported	Not Supported
	Hot Blob-Standard	Supported	Supported	NA	Supported	Supported
	Hot Blob-Gov	Supported	Not Supported	NA	Not Supported	Supported
	Cold Blob-Standard	Not Supported	Supported	NA	Supported	Not Supported

Cloud Provider	Class of Storage	Cloud Archive			CloudArchive Direct	Cloud Tier
		With Periodic Full	With Incremental Forever	Always Full		
Google Cloud Platform	Coldline	Supported	Supported (with Object Lock enabled)[5]	NA	Not Supported	Not Supported
	Multi-Regional	Supported	Not Supported	NA	Not Supported	Supported
	Nearline	Supported	Supported (with Object Lock enabled)[5]	NA	Not Supported	Not Supported
	Regional	Supported	Not Supported	NA	Not Supported	Supported
	Standard	Supported	Supported (with Object Lock enabled)[5]	NA	Not Supported	Supported
Google Next-Gen Cloud Edition	Standard	Not Supported	Not Supported	NA	Supported	Not Supported
Oracle	Archive Storage	Supported	Not Supported	NA	Not Supported	Not Supported
	Object Storage	Supported	Not Supported	NA	Not Supported	Supported

Cloud Provider	Class of Storage	Cloud Archive			CloudArchive Direct	Cloud Tier
		With Periodic Full	With Incremental Forever	Always Full		
Other External Targets	NAS	Supported	Supported[2]	NA	Supported[2]	Not Supported
	QStar Tape	Not Supported	Not Supported	Supported[3]	Not Supported	Not Supported
	S3 Compatible Object Storage - Standard Tier (Regular)	Supported	Supported[2]	NA	Supported[2]	Supported
	S3 Compatible Object Storage -Tape Based	Not Supported	Not Supported	Supported[3]	Not Supported	Not Supported

[1]**AWS Snowball and Snowball Edge**

- Cohesity does not support AWS legacy Snowball as an external target.
- AWS Snowball Edge can be used as a local S3 target by selecting **Snowball Edge** as the **Storage Class** while registering the AWS external target.
- Contact Cohesity Support if you want to ship the Snowball Edge device to AWS.

[2]**S3-Compatible Standard Tier and NAS**

With incremental forever archival format, the Cohesity cluster will download a portion of the data from the S3-compatible Standard tier and NAS external targets for the space reclamation process and reupload it after compaction. This leads to increased network bandwidth usage. Cohesity recommends having sufficient network bandwidth between the cluster and the external target, for example, hosting the external target within the same data center or network as the Cohesity cluster.

[3]**S3-Compatible Tape Based**

The archival format is Always Full and does not support deduplication and incremental forever archival.

[4]**AWS - S3-Intelligent**

Cloud Archive with incremental forever and CloudArchive Direct is supported for AWS - S3-Intelligent starting from the 7.2.2_u2 release.

[5] **GCP with Incremental Forever Archival**

You can archive data to GCP using the Incremental Forever archival format only when Object Lock is enabled. This is an Early Access feature. Contact your Cohesity account team to enable the feature.

S3-Compatible Object Storage - Standard Tier

Cohesity supports any S3-compatible device as an S3-Standard tier (Regular as storage class) external target with similar capabilities to an AWS S3-Standard. Any S3-compatible targets that use S3-Standard must be registered as a regular storage class. The S3-compatible uses S3-Standard APIs for archival and recovery. This Standard tier is supported for both incremental with periodic full and incremental forever archival formats. Compression and deduplication are supported for both archival formats.

With incremental forever archival format, the Cohesity cluster will download a portion of the data from the S3-compatible Standard tier external target for the space reclamation process and reupload it after compaction. This leads to increased network bandwidth usage. Cohesity recommends having sufficient network bandwidth between the cluster and the external target; for example, hosting the external target within the same data center or network as the Cohesity cluster.

> **Note:** Any S3-compatible targets that use Tape storage in the backend are not supported with the Standard tier; instead, you must register them as S3-compatible with the Tape-based Storage Class.

The following table provides the list of S3-Compatible Object Storage - Standard Tier and their details:

The S3-compatible table below contains a subset of supported S3-compatible targets. Refer to the S3 External Target Compatibility Tool to validate the S3-compatible targets.

Object Storage Vendor/ Cloud Provider	Product Names	Service URL	Port Number	Transport	AWS Signature Version	Additional Details
Backblaze	B2 Cloud Storage with v2 API version	s3.us-west-000 .backblazeb2.com	443	HTTPS	AWS Sign ver4	Requires V4 signing to provide S3 compatibility.
Cohesity	Cohesity Data Cloud	-	3000	HTTPS	AWS Sign ver4	-
Dell EMC	ECS	<fqdn-or-IP>: <portnumber>	9021	HTTPS	AWS Sign ver4	It is recommended to use AWS Sign ver4 with HTTPS port 9021.
			9020	HTTP	AWS Sign ver2	
Hitachi Vantara	Hitachi Content Platform	<fqdn-or-IP>	Not Required	HTTPS	AWS Sign ver4	Both v4 and v2 are supported for HTTPS.
		<fqdn-or-IP>:8080	8080	HTTP	AWS Sign ver2	-
Huawei Software Technologies Company	Object Storage Service	obs.ap-southeast-3.myhuaweicloud.com	Not required	HTTPS	AWS Sign ver2	-
IBM	Cloud Object Storage	s3.private.<region>-<site>.cloud-object-storage.appdomain.cloud	Not required	HTTPS	AWS Sign ver4	Both v4 and v2 work; v4 is the preferred version.

Object Storage Vendor/ Cloud Provider	Product Names	Service URL	Port Number	Transport	AWS Signature Version	Additional Details
11:11 Systems	11:11 Cloud Object Storage	us-central-1a.object.ilandcloud.com	Not required	HTTPS	AWS Sign ver4	-
		us-central-1b.object.ilandcloud.com				
		eu-west-1a.object.ilandcloud.com				
		eu-west-2a.object.ilandcloud.com				
		ap-southeast-1a.object.ilandcloud.com				
MInIO	MInIO Object Storage	<fqdn-or-IP>:9000	9000	HTTP	AWS Sign ver2	HTTP is the default protocol, but can be configured for HTTPS.
NetApp	StorageGrid	<fqdn-or-IP>:8082	8082	HTTPS	AWS Sign ver2	8082 via storage node.
	StorageGRID Webscale	<fqdn-or-IP>:18082	18082	HTTPS	AWS Sign ver2	18082 via gateway node.

Object Storage Vendor/ Cloud Provider	Product Names	Service URL	Port Number	Transport	AWS Signature Version	Additional Details
Pure Storage	Pure Flashblade	-	443	-	AWS Sign ver2	-
Wasabi	s3.us-east-1.wasabisys.com s3-<region>.wasabisys.com s3.wasabisys.com	-	Not required	HTTPS	AWS Sign ver4	Both v4 and v2 work; v4 is the preferred version.
Netdepot	-	us-central-1-s3.netdepot.com	Not required	HTTPS	AWS Sign ver4	-
IBM	IBM Storage Ceph	s3.cephlabs.com	443	HTTPS	AWS Sign ver4	-

S3-Compatible Object Storage - Tape-Based

Cohesity supports any S3-compatible targets that use tape storage in the backend as an external target and have capabilities similar to those of AWS S3 Glacier. Any S3-compatible targets that use Tape storage in the backend must be registered as Tape-based Storage Class. The tape-based Storage Class uses S3 Glacier APIs for archival and recovery. This supports always full with compression enabled by default and does not support deduplication and incremental archives.

> **Note:** S3-compatible Tape-based targets that deviate in their capabilities compared to AWS S3 Glacier tier are not supported. For the supported workflows of Tape-based, see Supported Workflows and External Targets.

Object Storage Vendor/ Cloud Provider	Product Names	Service URL	Port Number	Transport	AWS Signature Version	Additional Details
IBM	Storage Protect	<fqdn-or-IP>:9000	9000	HTTPS	V4	
Spectra Logic	Spectra Logic	<fqdn-or-IP>:443	443	HTTPS	V4	Solution Guide
PoINT Software & Systems	PoINT Archival Gateway	<fqdn-or-IP>:4443	4443	HTTPS	V4	Solution Guide

S3 External Target Compatibility Tool

The S3 external target compatibility tool validates whether the provided S3 compatible configuration has all the necessary permissions to execute the required APIs for successful archive and restore operations. This tool can validate S3-compatible storage targets and Tape-based S3 targets.

The S3 external target compatibility tool can be run from any linux environment and does not require a Cohesity cluster.

The S3 external target compatibility tool validates the following features:

- Incremental with periodic full
- Incremental with periodic full - versioning support
- Incremental with periodic full - object lock support
- Forever incremental
- Forever incremental - versioning support
- Forever incremental - object lock support

Install and Configure S3 Compatible Tool

To install the S3 compatibility tool:

1. Download the tool from the Cohesity Tools section on the Cohesity Download portal.
2. Run the following command to install the tool:

   ```
   ./s3_external_target_compatibility_tool_installer
   ```

   The installer is a self-extracting package that installs all required files under the following directory:

```
/tmp/s3_external_target_compatibility_tool
```

> **Note:** The dll_libs directory contains all dependent libraries. The installer automatically sets the correct interpreter and rpath.

3. Change the directory to `/tmp/s3_external_target_compatibility_tool`.
4. Run the following command to start the validation:

```
./s3_external_target_compatibility_tool --vault_config_file_path=<path_to_vault_config_file>
```

5. Run the following command to view the list of available options in the tool:

```
./s3_external_target_compatibility_tool
```

Configuration File Format

The S3 External Target compatibility tool requires the external target configuration file in the following format:

```
{
"access_key_id": "<Provide access key id>",
"secret_access_key": "<Provide secret access key>",
"host_name": "<Enter the hostname>",
"bucket_name": "<Enter the bucket name>",
"port_number": "<Optional: Default is 443. Specify only if different>",
"signature_version": <Optional: Default is 4. Values can be 2 or 4 (integer)>,
"use_https": <Optional: Default is true. Values can be true or false (boolean)>
"region": <Optional>
}
```

Parameters

The following table describes the external target configuration file parameters.

Parameter	Description
access_key_id	The access key ID for authentication.

Parameter	Description
secret_access_key	The secret access key for authentication.
host_name	The hostname of the S3-compatible storage system.
bucket_name	The bucket name where archives and restores will be performed.
port_number	(Optional) The port number to connect to the S3-compatible target (default: 443).
signature_version	(Optional) The signature version for API authentication (default: 4, possible values: 2 or 4)
use_https	(Optional) Boolean flag to indicate whether HTTPS should be used (default: true).
region	(Optional) Region of the bucket.

Related topics:

- Register an External Target for Archival
- Register an External Target for Cloud Tier

Prerequisites and Considerations for Registering External Targets

Before registering an external target, read and understand the following points:

- Ensure the external target is reachable. Verify by pinging the target from the Cohesity CLI.
- Check the network infrastructure and confirm it is not blocking connectivity.
- Check firewall policy settings. You may need to add the Cohesity node IP addresses to firewall rules.
- For AWS S3 or Glacier external targets, add the DNS names for the S3 and Glacier endpoints for the region to the allow list. For more information, see AWS service endpoints.
- For GCP, ensure that the Cloud Resource Manager API is enabled on the GCP project.
- Before registering GCP as external targets, ensure the following endpoints are accessible to the Cohesity cluster: "www.googleapis.com", "accounts.google.com".
- Cloud Tier targets cannot be unregistered currently.
- The target type cannot be modified once you register the target.

- You can register an external target either for Cloud Archive, CloudArchive Direct, or for Cloud Tier. You cannot register the same external target for Cloud Archive, CloudArchive Direct, and Cloud Tier purposes.
- The external targets that are registered for Cloud Tiering must have low latency. Therefore, high latency storage classes like glacier vault etc are not supported for Cloud Tiering.
- The same bucket can be registered as different external targets depending on the storage class.
- If you are using an S3 bucket with object versioning enabled as the external target, you must set up a Lifecycle Management Policy on the bucket to permanently delete non-current versions of objects. This is required to garbage-collect object versions. If there is a malware attack on the bucket that overwrites object versions, this policy could result in data loss in the absence of WORM locks on the object versions.
- To configure an HTTP proxy server for outbound internet connectivity from a Cohesity cluster to an external target, see this KB article.
- From the Cohesity UI, you cannot unregister an external target registered for CloudTier.
- For Cloud Archive with incremental forever and for CloudArchive Direct, Cohesity cluster may download a portion of the data from the S3-compatible and NAS external targets during each garbage collection process and reupload it after compaction. This can potentially lead to increased network bandwidth usage. To prevent additional network usage, Cohesity recommends using S3-compatible and NAS targets hosted within the same data center or network as the Cohesity cluster.
- If an external target you want to register for archival has object lock with a default retention period configured, then you must not select "incremental forever" as the archival format while registering this external target with the Cohesity cluster.

  Cohesity may require additional permissions to validate whether such a data lock policy is configured on the external target. For example, to validate the data lock configuration on AWS S3 targets, you must have the `s3:GetBucketObjectLockConfiguration` permission.
- If the **Object Key Pattern** of the S3 View is **Object ID**, then you must use the **Incremental Forever** archival format while archiving the S3 View. You cannot archive the S3 View using the **Incremental with Periodic Full** format.
- If you are planning to archive your data to Azure hot tier using the incremental forever archival format, then ensure that the storage account you create is accessible through a public network and not a private network.
- To avoid early deletion charges, garbage collection of objects archived to an external target is performed only after the minimum retention period is met. As a result, storage utilization at the target level may increase; however, the associated costs

typically remain the same or may even be lower.

- For archiving data to **S3 Compatible - Tape Based Storage Class**:
  - The archival format is *Always Full* with compression enabled by default.
  - Cohesity does not support deduplication and incremental archive.
  - Cohesity does not support file-level recovery.
  - The S3 cache must have sufficient storage to hold the largest Protection Group while the cache is filled with the restored data.
  - The recovery workflow retains the restored data from the tape on the S3 cache for seven days. If you want to reduce the retention period of restored data on the S3 cache, contact Cohesity Support.
  - If multiple recoveries are triggered within a span of seven days, the cache size must be large enough to hold all the recovered snapshots.
  - By default, the number of parallel streams per cluster is set to 10. If you want to configure these values, contact Cohesity Support.
  - By default, the number of restore tasks per Tape Storage is set to 1.
  - Cohesity recommends not to add more objects in a Protection Group as the recovery workflow could hydrate the complete archive while restoring a single object.
  - To protect IBM Storage Protect S3 Compatible, you must contact IBM support.
- You must not set up a Lifecycle Management (LCM) Policy on the external target registered for Cloud Tier.
- If you are archiving data using Cloud Archive or CloudArchive Direct with the Incremental Forever archival format to Amazon S3 Intelligent-Tiering, ensure you do not explicitly configure a bucket lifecycle policy that down-tiers data to the Archive or Deep Archive tiers. Doing so will cause the external target registration to fail.

  Additionally, if such a policy is configured after the S3 Intelligent-Tiering target has been registered with the cluster, any data archived to that bucket may become unrecoverable
- Prerequisites for Registering AWS S3 Buckts as External Targets Using IAM Roles Anywhere:

  Before registering an AWS S3 bucket as an external target on the Cohesity cluster for archival using IAM Roles Anywhere, follow these steps:

  1. Create a Certificate Authority (CA):
     - Ensure that the certificates are X.509 v3 compliant.
     - Set Basic Constraints to "CA" with the value "false".
     - Include the key usage for Digital Signature.

- Use a signing algorithm such as SHA256 or stronger. Avoid MD5 and SHA1 algorithms.

2. Create a Trust Anchor in AWS:
   - Establish the relationship between the certificate authority and IAM Anywhere by creating a Trust anchor.
3. Create an IAM Role:
   - This role should have access to the AWS S3 bucket.
   - Implement a custom trust policy specifically for the IAM Anywhere service principal.
4. Create an IAM Anywhere Profile:
   - Specify which roles IAM Roles Anywhere can assume. This will define the permissions and restrictions associated with the roles used by IAM Roles Anywhere.

For more information about creating a Trust Anchor and Profile role, see the AWS documentation.

Minimum Permissions for Registering External Targets

This topic covers the minimum permissions required for registering the AWS, Azure, and GCP as external targets.

Minimum Permissions for registering AWS External Targets

The following tables list the AWS permissions required for registering AWS source as an external target. The description column in the table explains why the permissions are needed. The permissions are for both AWS and AWS GovCloud.

You can create an AWS policy to grant permissions. For more information, see Assign AWS Permissions.

A JSON file is also provided for each set of permissions.

The permissions listed for Cloud Archive and CloudTier are required for both the IAM user and the IAM role used for registering the external target.

Cloud Archive (Incremental with Periodic Full) and CloudTier - S3

Open the JSON (or right-click and save as a file).

Storage Class	Resource Type	Permissions	Description
• S3-Standard • S3-IA-Standard • S3-Intelligent-Tiering • S3-One Zone • S3-One Zone-IA • S3-GOV • S3-C2S • S3-IA-C2S • S3-Glacier-Deep-Archive • S3-Glacier • Gov - S3-Glacier • Gov - S3 Glacier Deep Archive	S3	AbortMultipartUpload DeleteObject DeleteObjectVersion GetBucketLocation GetBucketObjectLockConfiguration GetBucketVersioning GetLifecycleConfiguration GetObject GetObjectAttributes GetObjectVersion GetObjectVersionAttributes ListBucket ListBucketMultipartUploads ListMultipartUploadParts PutObject PutObjectRetention RestoreObject ListBucketVersions GetIntelligentTieringConfiguration	Cohesity creates objects in the bucket given while archiving/tiering data to AWS. "s3:ListBucket" is required to list all the buckets in the AWS account and identify the archival/cloud tier target. "s3:GetObject", "s3:RestoreObject" and "s3:GetObjectVersion" are required during recovery from archive/cloud tier target. "s3:DeleteObject" and "s3:DeleteObjectVersion" are required to delete the object from the bucket when the policy's retention period for the archive/cloud tier is over. "s3:GetBucketLocation" is required to get the location of the bucket for cloud tier workflow. "s3:ListBucketMultipartUploads" and "s3:ListMultipartUploadParts" are needed to list out the parts uploaded to the bucket using multi part upload. "s3:AbortMultipartUpload" is needed to abort upload if Cohesity encounters an error. "s3:GetLifecycleConfiguration" is needed if user has lifecycle in the bucket. "s3:GetObjectAttributes" is used to fetch metadata of objects. "s3:GetBucketVersioning" is used to check whether buckets have versioning enabled. "s3:GetObjectVersionAttributes" is used to fetch metadata of objects in versioned objects.

Storage Class	Resource Type	Permissions	Description
			"s3:GetBucketObjectLockConfiguration" is used to get the object lock configuration on a bucket. "s3:PutObjectRetention" is used to set object-level locks in case of WORM enabled archives. "s3:ListBucketVersions" grants permission to list metadata about all the versions of objects in an Amazon S3 bucket.
	KMS	GenerateDataKey Decrypt	The "kms:GenerateDataKey" and "kms:Decrypt" permissions are required only when Server-Side Encryption with AWS Key Management Service keys (SSE-KMS) is enabled on the S3 bucket.

Cloud Archive (Incremental with Periodic Full) - Glacier

Open the JSON (or right-click and save as a file).

Storage Class	Resource Type	Permissions	Description
• Glacier-Standard • Glacier-C2S	Glacier	AbortMultipartUpload CompleteMultipartUpload DeleteArchive DescribeJob DescribeVault GetJobOutput InitiateJob InitiateMultipartUpload ListParts ListProvisionedCapacity UploadMultipartPart	Cohesity creates objects in the bucket given while archiving data to AWS. "glacier:InitiateJob" is needed to initiate a job in Glacier to list objects in Glacier. "glacier:DescribeVault" is needed to get details about the vault. "glacier:GetJobOutput" is needed to get output of jobs initiated by Cohesity. "glacier:ListParts" is needed to list parts. "glacier:DescribeJob" is needed to describe the job on the vault. "glacier:CompleteMultipartUpload", "glacier:InitiateMultipartUpload" and "glacier:UploadMultipartPart" permissions are required to upload objects to the archive bucket. "glacier:AbortMultipartUpload" is needed to abort upload if Cohesity encounters an error. "glacier:DeleteArchive" is required for garbage collection of the data in the cloud. "glacier:ListProvisionedCapacity" is required to check before issuing an expedited request to ensure a successful file/folder download.

Cloud Archive and CloudArchive Direct (Incremental Forever) - S3

Open the JSON (or right-click and save as a file).

Storage Class	Resource Type	Permissions	Description
• S3-Standard • S3-One-Zone • S3-IA • S3 Glacier IR • S3 Glacier • S3-Gov • S3-Gov Glacier • S3-Gov Glacier Deep Archive	S3	AbortMultipartUpload DeleteObject DeleteObjectVersion GetBucketLocation GetBucketObjectLockConfiguration GetBucketVersioning GetLifecycleConfiguration GetObject GetObjectAttributes GetObjectVersion GetObjectVersionAttributes ListBucket ListBucketMultipartUploads ListMultipartUploadParts PutObject PutObjectRetention RestoreObject ListBucketVersions	Cohesity creates objects in the bucket given while archiving/tiering data to AWS. "s3:ListBucket" is required to list all the buckets in the AWS account and identify the archival/cloud tier target. "s3:GetObject", "s3:RestoreObject" and "s3:GetObjectVersion" are required during recovery from archive/cloud tier target. "s3:DeleteObject" and "s3:DeleteObjectVersion" are required to delete the object from the bucket when the policy's retention period for the archive/cloud tier is over. "s3:GetBucketLocation" is required to get the location of the bucket for cloud tier workflow. "s3:ListBucketMultipartUploads" and "s3:ListMultipartUploadParts" are needed to list out the parts uploaded to the bucket using multi part upload. "s3:AbortMultipartUpload" is needed to abort upload if Cohesity encounters an error. "s3:GetLifecycleConfiguration" is needed if user has lifecycle in the bucket. "s3:GetObjectAttributes" is used to fetch metadata of objects. "s3:GetBucketVersioning" is used to check whether buckets have versioning enabled. "s3:GetObjectVersionAttributes" is used to fetch metadata of objects in versioned objects. "s3:GetBucketObjectLockConfiguration" is used to get the object lock configuration on a bucket. "s3:PutObjectRetention" is used to set object-level locks in case of WORM enabled archives. "s3:ListBucketVersions" grants

Storage Class	Resource Type	Permissions	Description
			permission to list metadata about all the versions of objects in an Amazon S3 bucket.
	Lambda	InvokeFunction UpdateFunctionCode	"lambda: InvokeFunction" is required for garbage collection. "lambda:UpdateFunctionCode" is required for supporting upgrade in the future.
	KMS	GenerateDataKey Decrypt	The "kms:GenerateDataKey" and "kms:Decrypt" permissions are required only when Server-Side Encryption with AWS Key Management Service keys (SSE-KMS) is enabled on the S3 bucket.

CloudRetrieve - S3

Open the JSON (or right-click and save as a file).

Storage Class	Resource Type	Permissions	Description
• S3-Standard • S3-IA-Standard • S3-Intelligent-Tiering • S3-One Zone • S3-One Zone-IA • S3-GOV • S3-C2S • S3-IA-C2S • S3-Glacier-Deep-Archive • S3-Glacier	S3	GetLifecycleConfiguration GetObject GetObjectVersion ListBucket RestoreObject	"s3:ListBucket" is required to list all the buckets in the AWS account and identify the archival target. "s3:GetObject" and "s3:GetObjectVersion" and "s3:RestoreObject" are required during retrieve from archive. "s3:GetLifecycleConfiguration" are required for LifeCycle.

CloudRetrieve - Glacier

Open the JSON (or right-click and save as a file).

Storage Class	Resource Type	Permissions	Description
• Glacier-Standard • Glacier-C2S	Glacier	DescribeJob DescribeVault GetJobOutput InitiateJob ListParts	While retrieving objects from archive, "glacier:InitiateJob" is needed to initiate a job in Glacier to list objects in Glacier. "glacier:DescribeVault" and "glacier:DescribeJob" are required to describe the vault details and job details. "glacier:ListParts" is needed to list parts while retrieving objects.

Minimum Permissions for registering Azure External Targets

When registering Azure as an external target for archival or tiering, if authentication is performed using storage access keys, no additional permissions are required.

When registering Azure as an external target for archival or tiering, if authentication is performed using managed identity, then the following permissions are required:

Resource Provider	Operation Name
Microsoft.Storage	• Microsoft.Storage/storageAccounts/blobServices/containers/* • Microsoft.Storage/storageAccounts/blobServices/containers/blobs/*

Optionally, you can assign the predefined **Storage Blob Data Owner** role to the managed identity. This role includes the permissions required to access the storage account.

Minimum Permissions for registering GCP External Targets

The following tables list the GCP permissions required for registering GCP source as an external target.

Action	Permissions

Cloud Archive	**If using Client Private Key authentication:**
	`storage.buckets.get`
	`storage.multipartUploads.abort`
	`storage.multipartUploads.create`
	`storage.multipartUploads.list`
	`storage.multipartUploads.listParts`
	`storage.objects.create`
	`storage.objects.delete`
	`storage.objects.get`
	`storage.objects.list`
	`storage.objects.update`
	If using Service Account Email Address authentication:
	`storage.buckets.get`
	`storage.multipartUploads.abort`
	`storage.multipartUploads.create`
	`storage.multipartUploads.list`
	`storage.multipartUploads.listParts`
	`storage.objects.create`
	`storage.objects.delete`
	`storage.objects.get`
	`storage.objects.list`
	`storage.objects.update`
	Optionally, you can assign the Storage Object Admin role (roles/storage.objectAdmin) to the service account. This predefined role includes the permissions required to access Google Cloud Storage
	The service account attached the Cohesity cluster node has the `iam.serviceAccounts.getAccessToken` permission to assume role of the service account to be impersonated.

Recommended Azure Storage Account Configuration for External Targets

When using Microsoft Azure as an external target with Cloud Archive or Cloud Archive Direct, we recommend specific Azure Storage Account settings to support reliable operation and cost effective storage.

Before estimating Azure storage costs, determine the required storage capacity for the external target by obtaining the total capacity from the Cohesity sales team, based on the internal sizing tools.

This topic describes the recommended Azure Storage Account configuration options and explains why those selections are typically used. With this information, you can

independently use Microsoft's Azure Pricing Calculator to estimate storage costs based on their own capacity and retention requirements.

> **Note:** Azure options and defaults may change over time. Always refer to the official Azure documentation and Azure Pricing Calculator for the most current information.

Recommended Settings for Azure Storage Account

We recommend the following settings when configuring an Azure Storage Account for use as an external target:

- Region: Select the Azure region where the cluster runs or where the external target data will be stored.
- Storage Type: Select **Azure Blob Storage** as the storage type.
- Storage Account Type: Select **General Purpose v2** as the storage account type so that Hot, Cool, Cold, and Archive access tiers are available.
- Performance Tier: Select **Standard** as the performance tier, because standard performance is typically sufficient for Cloud Archive workloads.
- Access Tier: Select **Hot**, **Cool**, **Cold**, or **Archive** based on access frequency and retention requirements, accounting for Azure minimum retention and early deletion charges.
- Redundancy: Select **Locally Redundant Storage (LRS)** as the redundancy option to keep storage and network costs low. Note that Geo-Redundant Storage (GRS) is selected by default.
- Namespace: Select a flat namespace for the file structure, because hierarchical namespaces are usually not required.
- Capacity: Enter the required storage capacity obtained earlier and review the estimated monthly and yearly storage costs shown by the Azure Pricing Calculator.

Additional Considerations

- Azure Pricing Calculator estimates reflect Azure storage costs only and exclude Cohesity licensing and Azure compute costs.
- Exact read, write, list, and delete operation costs vary; some teams account for approximately 5% additional cost beyond raw storage charges.
- Configure the actual Azure Storage Account using the Azure portal, CLI, or APIs after planning.
- Some Azure features (for example, immutability) are configured at account creation but do not affect pricing calculator inputs.

- If you want to use Cohesity Data Lock for immutability, you must enable the Object Lock on the Azure Storage Account.

Register an External Target for Archival

Before register an external target for archival, ensure you have:

- Met the pre-requirements to register an external target.
- Reviewed the considerations.
- The minimum permissions to register the external target.

To register an external target for Cloud Archive, perform the following steps:

1. Navigate to **Infrastructure > External Target**.
2. Click **Add External Target** located at the top right of the page.
3. Select **Archival** as the external target's **Purpose**.
4. From the **Storage Type** drop-down list, select the external target you want to register with the Cohesity cluster for Cloud Archive.
5. From the **Storage Class** drop-down, select the storage class of the storage type you selected.

   > **Note:** For Cloud Archive with incremental forever and for CloudArchive Direct, Cohesity cluster may download a portion of the data from the S3-compatible and NAS external targets during each garbage collection process and reupload it after compaction. This can potentially lead to increased network bandwidth usage. To prevent additional network usage, Cohesity recommends using S3-compatible and NAS targets hosted within the same data center or network as the Cohesity cluster.

   AWS S3 Standard or AWS Glacier(Legacy) Standard or AWS S3-Glacier Standard (Currently S3 Glacier Flexibel Retrieval) or AWS S3-Glacier-Deep-Archive Standard or AWS S3 Intelligent Standard or AWS S3 One Zone Standard

   For IAM user, if you do not already have an access key to use with Cohesity, create a new one in the AWS Management Console (root account credentials required). In AWS, go to **IAM** (AWS Identity and Access Management) and create a new access key ID and secret access key ID. Use this access key ID and secret access key to complete the corresponding fields below.

> **Warning:** Do not apply an AWS Object Lifecycle Policy with the Expiration action to an AWS S3 External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Bucket Owner Account Id	The AWS account ID of the Amazon storage you archive the data. This field is not applicable for AWS Glacier(Legacy) Standard.
Authentication Method: **IAM User**	
Access Key ID	Access key ID used to access Amazon storage.
Secret Access Key	Secret access key address to access Amazon storage.
Authentication Method: IAM Anywhere • Ensure you meet the prerequisites before registering the AWS S3 buckets using **IAM Anywhere Roles**. • Cohesity does not support IAM anywhere authentication for registering Glacier(Legacy) Standard with the Cohesity cluster.	
Trust Anchor ARN	ARN of the trust anchor.
Profile ARN	ARN of the IAM Roles Anywhere profile.

Field Name	Description
Role ARN	The ARN of the IAM role that trusts the IAM Roles Anywhere service principal.
Private Key	Private key of the Certificate Authority.
Authentication Method: **IAM Role** **Note:** To register the AWS Cloud Source using the IAM Role, you need to download the CloudFormation Template and execute on the AWS account. For more information, see Execute the CloudFormation Template.	
Account ID	The account ID where the CE is running
IAM Role	The IAM Role for external target registration. Example: `arn:aws:iam::<account>:role/<role-name-with-path>`.

For most regions, the Cohesity uses `s3-<region>.amazonaws.com` to construct the S3 endpoints instead of `s3.<region>.amazonaws.com`.

Example

`s3-us-east-2.amazonaws.com`, where `s3-us-east-2.amazonaws.com` is the CNAME for `s3.us-east-2.amazonaws.com`.

However, the Cohesity uses `s3.<region>.amazonaws.com` for regions listed below:

- China
- us-east-1
- S3 acceleration

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for AWS.

You can archive data to Azure Data Box by registering Data Box as NAS External Target on Cohesity cluster. For more information, see How to use Cohesity Microsoft Azure Data Box Integration.

AWS S3 Gov or AWS Glacier(Legacy) Gov or AWS S3-Glacier Gov or AWS S3-Glacier-Deep-Archive Gov or AWS S3 Intelligent Gov or AWS S3 One Zone Gov

> **Note:** IAM Role for AWS S3 Gov is not supported.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Bucket Owner Account Id	The AWS account ID of the Amazon storage you archive the data. This field is not applicable for AWS Glacier(Legacy) Gov.
Access Key ID	Access key ID used to access Amazon storage.
Secret Access Key	Secret access key address to access Amazon storage.

AWS S3 C2S or AWS Glacier(Legacy) C2S

You must contract with a Commercial Cloud Services (C2S) agency to provide a private instantiation of the AWS commercial cloud.

Field Name	Description
Base URL	URL of the Commercial Cloud Services (C2S) server that contains the Amazon storage provided by the C2S agency.
Agency	The name the agency that is providing the Commercial Cloud Services (C2S) server.
Mission	The account name for accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Role	The IAM role to use when accessing the Amazon storage which is provided by the C2S agency.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Bucket Owner Account Id	The AWS account ID of the Amazon storage you archive the data. This field is not applicable for AWS Glacier(Legacy) C2S.
Client Certificate Password	The password to read the Client Private Key when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
CA Trusted Certificate	Upload the Certificate Authority Trusted Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Certificate	Upload the Client Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Private Key	Upload the Client Private Key file to use when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for AWS.

AWS S3-IA Standard or AWS-S3-Glacier-IR, or AWS-S3-Glacier-IR Gov

For IAM user, if you do not already have an access key to use with Cohesity, create a new one in the AWS Management Console (root account credentials required). In AWS, go to **IAM** (AWS Identity and Access Management) and create a new access key ID and secret access key ID. Use this access key ID and secret access key to complete the corresponding fields below.

> **Warning:** Do not apply an AWS Object Lifecycle Policy with the Expiration action to an AWS S3 External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Bucket Owner Account Id	The AWS account ID of the Amazon storage you archive the data.
Authentication Method: **IAM User**	
Access Key ID	Access key ID used to access Amazon storage.
Secret Access Key	Secret access key address to access Amazon storage.
Authentication Method: IAM Anywhere • Ensure you meet the prerequisites before registering the AWS S3 buckets using **IAM Anywhere Roles**. • Cohesity does not support IAM anywhere authentication for registering **AWS-S3-Glacier-IR Gov** with the Cohesity cluster.	
Trust Anchor ARN	ARN of the trust anchor.
Profile ARN	ARN of the IAM Roles Anywhere profile.

Field Name	Description
Role ARN	The ARN of the IAM role that trusts the IAM Roles Anywhere service principal.
Private Key	Private key of the Certificate Authority.
Authentication Method: **IAM Role** **Note:** To register the AWS Cloud Source using the IAM Role, you need to download the CloudFormation Template and execute on the AWS account. For more information, see Execute the CloudFormation Template.	
Account ID	The account ID where the CE is running
IAM Role	The IAM Role for external target registration. Example: `arn:aws:iam::<account>:role/<role-name-with-path>`.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for AWS.

AWS S3-IA Gov

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Bucket Owner Account Id	The AWS account ID of the Amazon storage you archive the data.

Field Name	Description
Access Key ID	Access key ID used to access Amazon storage.
Secret Access Key	Secret access key address to access Amazon storage.

AWS S3-IA C2S

You must contract with a Commercial Cloud Services (C2S) agency to provide a private instantiation of the AWS commercial cloud. IA refers to an Amazon storage class that is infrequently accessed.

> **Warning:** Do not apply an AWS Object Lifecycle Policy with the Expiration action to an AWS S3 External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Base URL	URL of the Commercial Cloud Services (C2S) server that contains the Amazon storage provided by the C2S agency.
Agency	The name the agency that is providing the Commercial Cloud Services (C2S) server.
Mission	The account name for accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Role	The IAM role to use when accessing the Amazon storage which is provided by the C2S agency.
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.

Field Name	Description
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. > **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Bucket Owner Account Id	The AWS account ID of the Amazon storage you archive the data.
Client Certificate Password	The password to read the Client Private Key when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
CA Trusted Certificate	Upload the Certificate Authority Trusted Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Certificate	Upload the Client Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Private Key	Upload the Client Private Key file to use when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for AWS.

Azure Archive Blob

After creating a storage account in Azure, if you have not done so already, go to Blob service and create a new container to use with Cohesity. Provide the container name, storage account name and either of the keys generated by Azure to the Cohesity Dashboard.

> **Note:** Cohesity recommends you to use the **Standard** (optimized for high capacity and high throughput) performance tier and not the Premium performance tier when creating a new Azure Blob storage account and container to register with the Cohesity cluster as an external target.

> **Warning:** Microsoft supports Azure Archive for block blobs in Blob storage accounts with replication set to Locally-redundant storage (LRS).

Field Name	Description
Container Name	Name of the existing container where you will store objects. This cannot be modified after you register the Target.
Storage Account Name	Account name used to access Azure storage.
Authenitcation method: **Storage Access Key**	
Storage Access Key	Storage access key to access Azure storage.
Authenitcation method: **Managed Identity** **Important:** This authentication method is supported only on Cloud Edition clusters.	
Client ID	Provide the client ID of the Managed Identity that is attached to the cluster VM. Ensure the Managed Identity has the following required permission to access the storage account you are registering. • Microsoft.Storage/storageAccounts/blobServices/containers/* • Microsoft.Storage/storageAccounts/blobServices/containers/blobs/* Optionally, you can assign the predefined **Storage Blob Data Owner** role to the managed identity. This role includes the permissions required to access the storage account.

> **Note:** When using archive tier on Azure Blob, Cohesity rehydrates the data from archive tier to hot tier instead of cool tier. Cool tier has a minimum retention duration of 30 days after which the data automatically moves to the archive tier. Hence, Cohesity uses hot tier. The data moves back to the archive tier after the restore job is deleted/expired.

> **Warning:** Do not apply an lifecycle management policy with the Expiration action to an Azure External Target used for archiving data. Applying a policy with this action may result in data loss.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for Azure.

Azure Cold Blob Standard or Gov

After creating a storage account in Azure, if you have not done so already, go to Blob service and create a new container to use with Cohesity. Provide the container name, storage account name and either of the keys generated by Azure to the Cohesity Dashboard.

> **Note:** Cohesity recommends you to use the **Standard** (optimized for high capacity and high throughput) performance tier and not the Premium performance tier when creating a new Azure Blob storage account and container to register with the Cohesity cluster as an external target.

> **Warning:** Do not apply a lifecycle management policy with the Expiration action to an Azure External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Container Name	Name of the existing container where you will store objects. This cannot be modified after you register the Target.

Field Name	Description
Storage Account Name	Account name used to access Azure storage.
Authenitcation method: **Storage Access Key**	
Storage Access Key	Storage access key to access Azure storage. **Note:** If you are planning to archive your data to Azure using the incremental forever archival format, then ensure that the storage account you create is accessible through a public network and not a private network.
Authenitcation method: **Managed Identity** **Important:** This authentication method is supported only on Cloud Edition clusters.	
Client ID	Provide the client ID of the Managed Identity that is attached to the cluster VM. Ensure the Managed Identity has the following required permission to access the storage account you are registering. • Microsoft.Storage/storageAccounts/blobServices/containers/* • Microsoft.Storage/storageAccounts/blobServices/containers/blobs/* Optionally, you can assign the predefined **Storage Blob Data Owner** role to the managed identity. This role includes the permissions required to access the storage account.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for Azure.

Azure Hot Blob Standard or Gov

After creating a storage account in Azure, if you have not done so already, go to Blob service and create a new container to use with Cohesity. Provide the container name, storage account name and either of the keys generated by Azure to the Cohesity Dashboard.

> **Note:** Cohesity recommends you to use the **Standard** (optimized for high capacity and high throughput) performance tier and not the Premium performance tier when creating a new Azure Blob storage account and container to register with the Cohesity cluster as an external target.

> **Warning:** Do not apply a lifecycle management policy with the Expiration action to an Azure External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Container Name	Name of the existing container where you will store objects. This cannot be modified after you register the Target.
Storage Account Name	Account name used to access Azure storage.
Authenitcation method: **Storage Access Key**	
Storage Access Key	Storage access key to access Azure storage. **Note:** If you are planning to archive your data to Azure Hot Blob using the incremental forever archival format, then ensure that the storage account you create is accessible through a public network and not a private network.
Authenitcation method: **Managed Identity** **Important:** This authentication method is supported only on Cloud Edition clusters.	

Field Name	Description
Client ID	Provide the client ID of the Managed Identity that is attached to the cluster VM. Ensure the Managed Identity has the following required permission to access the storage account you are registering. • Microsoft.Storage/storageAccounts/blobServices/containers/* • Microsoft.Storage/storageAccounts/blobServices/containers/blobs/* Optionally, you can assign the predefined **Storage Blob Data Owner** role to the managed identity. This role includes the permissions required to access the storage account.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for Azure.

GCP Coldline, GCP Multi-Regional, GCP Nearline or GCP Regional or GCP Standard

In the Google Cloud Platform Console, create a new bucket to use with Cohesity if you have not done so already. To create credentials to use with Cohesity, from the Dashboard go to **IAM & Admin > IAM** and create a service account. Give the service account the Role of Project Owner and select Google Nearline or Coldline. After you create the service account, download and open the credential (JSON) file and use the contents (copy the text in the fields between the double quotes) to complete the fields below.

> **Warning:** Do not apply a Google Cloud Object Lifecycle Policy with the Delete action to any GCP External Target used for archiving data. Applying a policy with this action may result in data loss.

> **Note:** Before registering GCP as external targets, ensure the following endpoints are accessible to the Cohesity cluster: "www.googleapis.com", "accounts.google.com".

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.

Field Name	Description
Project ID	Project ID of the existing project where you will store objects. This cannot be modified after you register the Target.
Authenitcation method: **Client Private Key** **Important:** This authentication method is supported only on Cloud Edition clusters.	
Client Email Address	Email address used to access Google Cloud Storage.
Client Private Key	Private key used to access Google Cloud Storage.
Authenitcation method: **Service Account Impersonation** **Important:** This authentication method is supported only on Cloud Edition clusters.	

Field Name	Description

Service Account Email Address	Email address of the service account to be impersonated.
	Ensure:
	• The service account to be impersonated has the following permissions required to access the Google Cloud Storage:
	`storage.buckets.get`
	`storage.multipartUploads.abort`
	`storage.multipartUploads.create`
	`storage.multipartUploads.list`
	`storage.multipartUploads.listParts`
	`storage.objects.create`
	`storage.objects.delete`
	`storage.objects.get`
	`storage.objects.list`
	`storage.objects.update`
	Optionally, you can assign the **Storage Object Admin** role (`roles/storage.objectAdmin`) to the service account. This predefined role includes the permissions required to access Google Cloud Storage.
	• The service account attached the Cohesity cluster node has the "`iam.serviceAccounts.getAccessToken`" permission to assume role of the service account to be impersonated.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for GCP.

You can archive data to Google Transfer Appliance by registering Google Transfer Appliance as NAS External Target on Cohesity cluster.

Oracle Archive Storage or Oracle Object Storage

A tenancy must be created by Oracle for your organization. In addition, an administer must set up the Tenancy and possibly create a compartment to hold your resources or a new user if required. Determine the best Oracle storage tier to use for your organization: Oracle Archival Storage (Archive Storage tier) or Oracle Object Storage (Object Storage tier). In the Oracle Cloud console, create a new bucket to use with Cohesity if you have not done so already. The Storage Tier specified for the bucket must match the specified Cohesity External Target Type (Oracle Archive or Oracle Object).

The user with which the bucket is registered must be associated with the following Oracle policy statements:

- Allow group <*group-name*> to manage buckets in tenancy
- Allow group <*group-name*> to manage objects in tenancy

To register OCI as an external target, Cohesity utilizes Amazon S3 Compatibility API keys for access to the OCI bucket. Use the OCI console to create the "**Customer Secret keys**" and then use the Access Key ID and Secret Access Key to complete the corresponding fields below.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target. In the Oracle Cloud console, create a new bucket to use with Cohesity if you have not done so already. The Storage Tier specified for the bucket must match the specified Cohesity External Target Type (Oracle Archive or Oracle Object).
Region	Geographic area to use for Oracle storage. For more information, see Data Regions.
Access Key ID	Access key ID used to access the storage. For more information about create and copying the key, see create a secret key.
Secret Access Key	Secret access key address to access the storage. For more information about create and copying the key, see create a secret access key.
Tenancy	Enter the namespace to which the bucket belongs. You can obtain the **Namespace** name from the Bucket Details page of the Oracle Cloud Infrastructure console.

NAS

Only NFS is currently supported.

Field Name	Description
NAS Host	Hostname or host IP of the existing NAS host.
Mount Path	The mount path of the NAS host.
Share Type NFS	Only NFS is currently supported.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for NAS.

QStar Tape

To use QStar with Cohesity, QStar must be configured with the settings described in the Archive Data to Tape with QStar & Cohesity.

Field Name	Description
QStar Host	Hostname or host IP of the existing QStar host.
QStar Web Services Port	The QStar port that Cohesity communicates with. Port number 18082 is set on the QStar host by default. If your QStar host is not using the default Web Services Port, update the port number.
QStar Username	Account name used to access the QStar host.
QStar Password	Password used to access the QStar host.
QStar Integral Volumes	Names of the existing Integral Volumes where you will store objects. Cohesity supports the LTFS and TDO file systems for Integral Volumes. The Integral Volumes must have already been created on the QStar host with the necessary minimal configuration as recommended in the QStar Administrator's Guide. If you attempt to edit this list after the External Target was initially registered, be advised that the Cohesity cluster can only archive to and recover from Integral Volumes that are currently registered with Cohesity, that is, if you edit the External Target by adding an Integral Volume, it cannot be used by the Cohesity cluster. In this case, you must register a new External Target that includes the Integral Volume.

S3 Compatible - Standard Storage Class

- Cohesity supports any S3-compatible device as an S3-Standard tier (Regular as storage class) external target with similar capabilities to an AWS S3-Standard.
- With incremental forever archival format, the Cohesity cluster will download a portion of the data from the S3-compatible Standard tier external target for the space reclamation process and reupload it after compaction. This leads to increased network bandwidth usage. Cohesity recommends having sufficient network bandwidth between the cluster and the external target; for example, hosting the external target within the same data center or network as the Cohesity cluster.

- Do not apply an AWS Object Lifecycle Policy with the Expiration action to an AWS S3 Compatible External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Access Key ID	Access key ID used to access the storage.
Secret Access Key	Secret access key address to access the storage.
Endpoint	The Endpoint mount path. Use the IP address or Fully Qualified Domain Name (FQDN). **Note:** Do not include 'http://' or 'https://' in the mount path.
Port	The S3 Compatible port that Cohesity communicates with.
Region	Geographic area to use for S3 Compatible.
Secure Connection (HTTPS)	Toggle **Secure Connection** on or off.
AWS Signature Version	Select the AWS signing process version for adding authentication information to requests.

To register other S3 compatible targets like Pure FlashBlade S3 bucket, IBM ICOS S3 bucket and Cohesity S3 bucket, see CloudArchive Deployment & CloudRetrieve Recovery for S3-Compatible Storage.

S3 Compatible - Tape Based Storage Class

- The archival format is *Always Full* with compression enabled by default.
- Cohesity does not support deduplication and incremental archive.
- Cohesity does not support file-level recovery.
- The S3 cache must have sufficient storage to hold the largest Protection Group while the cache is filled with the restored data.

- The recovery workflow retains the restored data from the tape on the S3 cache for seven days. If you want to reduce the retention period of restored data on the S3 cache, contact Cohesity Support.
- If multiple recoveries are triggered within a span of seven days, the cache size must be large enough to hold all the recovered snapshots.
- By default, the number of parallel streams per cluster is set to 10. If you want to configure these values, contact Cohesity Support.
- By default, the number of restore tasks per Tape Storage is set to 1.
- Cohesity recommends not to add more objects in a Protection Group as the recovery workflow could hydrate the complete archive while restoring a single object.
- To protect IBM Storage Protect S3 Compatible, you must contact IBM support.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Access Key ID	Access key ID used to access the storage.
Secret Access Key	Secret access key address to access the storage.
Endpoint	The Endpoint mount path. Use the IP address or Fully Qualified Domain Name (FQDN). **Note:** Do not include 'http://' or 'https://' in the mount path.
Port	The S3 Compatible port that Cohesity communicates with.
Region	Geographic area to use for S3 Compatible.
Secure Connection (HTTPS)	Toggle **Secure Connection** on or off.
AWS Signature Version	Select the AWS signing process version for adding authentication information to requests.

11:11 Systems S3-IA Standard

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Bucket Owner Account Id	The AWS account ID of the Amazon storage you archive the data.
Access Key ID	Access key ID used to access Amazon storage.
Secret Access Key	Secret access key address to access Amazon storage.

> **Note:** A CloudFormation template download is not required for 11:11 Systems.

Micrologic, Rackspace, Expedient, or IBM Cloud S3 Compatible - Standard Storage Class

- Standard Storage Class is the only supported storage class.
- Archive Object Lock is not supported.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Access Key ID	Access key ID used to access the storage.
Secret Access Key	Secret access key address to access the storage.

Field Name	Description
Endpoint	The Endpoint mount path. Use the IP address or Fully Qualified Domain Name (FQDN). **Note:** Do not include 'http://' or 'https://' in the mount path.
Port	The S3 Compatible port that Cohesity communicates with.
Secure Connection (HTTPS)	Toggle **Secure Connection** on or off.
AWS Signature Version	Select the AWS signing process version for adding authentication information to requests.

6. (*Applies only for AWS S3 Targets*) To archive your data over a private network, enable Use Private Endpoint URL and provide the following details:
   - **Amazon S3 Endpoint URL** – The Regional S3 DNS name generated when creating the S3 interface endpoint. This interface endpoint must be created in the same VPC as the cluster and associated with the appropriate subnets and security groups.
   - **AWS Lambda Private Endpoint URL (Optional)** – The *Regional Lambda DNS name* generated when creating the Lambda interface endpoint. This endpoint must also be created in the same VPC as the cluster and associated with the correct subnets and security groups.
7. Provide a name in the **External Target Name** field.
8. In the **Archival Format** field, **Incremental Forever** is auto-selected for external targets that support it. For other targets, **Incremental with Periodic Full** is selected by default. This field cannot be modified. For more information, see Archival Formats.

   > **Note:** This field is not applicable for Tape Based external target.

9. Enable **Archive Object Lock** to lock the archives in the targets to prevent modification of data on the external target.

   For AWS S3 targets with incremental forever archival, you can select one of the object lock modes:

- Governance mode:

  Provides protected retention with the ability for privileged users to override if needed. This option is enabled by default.

- Compliance mode:

  Enforces strict WORM protection. User cannot modify or delete locked objects until retention expires.

  If you select Compliance mode, you must acknowledge the warning message by typing **Yes** in the confirmation text box to proceed with registering the external target with Object Lock in Compliance mode.

For more information, see Archive Object Lock.

10. (*Applies only for AWS S3 targets with incremental forever archival*) Download and run the CloudFormation Template by clicking the **CloudFormation Template** link available in the product UI. For more information, see Execute the CloudFormation Template.

    > **Note:**
    > - A CloudFormation Template is required for configuring garbage collection with Amazon S3. However, for the Glacier (Legacy), S3 Glacier Deep Archive, and S3 Glacier Flexible Retrieval (FR) storage classes, a CloudFormation Template is not required. In these cases, data is written directly to the respective S3 storage class, and garbage collection is automatically managed through S3 lifecycle policies and bucket expiration rules.
    > - This option is not available if you have enabled **Archive Object Lock**.

11. (*Applies only for Azure target (hot blob, cold blob, or cool blob) with incremental forever archival*) Enter the **Function App Name** and **Function App Deployment Key**. For information on obtaining the Function App name and Function App Deployment key, see the Azure Function App.

    > **Note:** This option is not available if you have enabled **Archive Object Lock**.

12. (**Applies only to S3-compatible and NAS external targets using Incremental Forever archival**) For Incremental Forever archival, the Cohesity cluster may download portions of data from S3-compatible and NAS external targets during each garbage collection process and reupload it after compaction. Choose a **Space**

Reclamation option based on your optimization preference:

- **Storage Optimized**: Selecting this option maintains a regular garbage collection frequency to optimize storage consumption on the archive. However, this results in higher network bandwidth usage. To minimize additional network usage, Cohesity recommends using S3-compatible and NAS targets within the same data center or network as the Cohesity cluster.
- **Network Optimized**: Selecting this option reduces garbage collection frequency to minimize network bandwidth usage. However, this results in higher storage consumption on the cluster due to less frequent garbage collection.

13. By default, the **Encryption** option is enabled with **Internal KMS** as the **Key Management Service (KMS) Type**.You can also register an external KMS by selecting one of the following as the Key Management Service (KMS) Type:
    - AWS
    - KMIP Compliant

    For information on configuring external KMS, see External Key Management Service.

    If you have selected **Internal KMS** or **KMIP Compliant** as the **Key Management Service (KMS) Type** for encryption, then you can enable additional security for the external target. By enabling additional security, you can manage the key manually instead of Cohesity managing the encryption key for data that is transferred out. If you choose to manage the key manually, then after registering the external target with the Cohesity cluster, you must download the key from the external target's details page and store it in a secure location outside of the Cohesity cluster. With additional security, only the Cohesity clusters where you have explicitly provided the encryption key can download data from that external target. However, the source Cohesity cluster (the cluster that archived the data) does not require the encryption key of the external target.

    Cohesity recommends you keep the **Encryption** option as enabled in almost all cases, except when the data is already encrypted.

    > **Note:** If **Encryption** is enabled at the cluster level, you must also enable encryption while registering the external target for CloudArchive Direct. Archiving directly to an external target may fail if encryption is enabled at the cluster level but not at the external target level.

14. By default, the **Compression** option is enabled to send data to the external target in a compressed format. Disable this option to send data to the external target in uncompressed format.

15. By default, the **Source Side Deduplication** option is enabled to deduplicate data before sending it to the External Target. This improves cloud performance and reduces the amount of data transferred from the Cohesity cluster. This option is not applicable for:
    - **Tape-based external targets** – Source-side deduplication is not supported.
    - **Archival with Incremental Forever** – Deduplication is already enabled by default in the backend.
16. Enable **Bandwidth Throttling** to limit the maximum data transfer rate *from and to this External Target* during a time window.
    a. **Traffic:** Specify the direction of traffic to limit, by specifying a maximum bandwidth:
        - **Upload:** Limit the data transfer rate from the Cohesity cluster to this External Target.
        - **Download:** Limit the data transfer rate from this External Target to the Cohesity cluster.
    b. **Days**: Specify the days of the week to apply the window.
    c. **Time**: Specify all day or a time period for each day.
    d. **Throttle To**: Specify the maximum data transfer rate.Optionally change the rate type.
    e. Click **Add**.
    f. Repeat these steps to add more upload throttling windows.

    If overlapping windows are specified, the latest window that was created takes precedence. For example, if you first create a window from Monday 9 AM - 10 PM with a data rate of 50 Mbps and then a window from 5 PM to 12 PM with a data rate of 100 Mbps, during 5 PM -12 PM the data rate is limited to 100 Mbps.

    If both cluster-level override throttling window and an external-target override throttling window is set during the same time period, the lowest **Throttle To** number specified is the maximum allowed data transfer rate. For example, if a cluster-level override throttling window of 8-11 AM is set to 250 Mbps and an External-level override throttling window of 9-10 AM is set to 200 Mbps, during 9-10 AM the maximum allowed data transfer rate is 200 Mbps. Specifying transfer rate limits at the External-target level is described at the end of this procedure. For information on setting up cluster-level bandwidth throttling, see Register an External Target for Archival.
17. Click **Register**.

    The Cohesity cluster attempts to connect to the Target. A successful connection displays the External Targets page with the new Target included in the list.

> **Note:** When adding an AWS S3-compatible external target, you may encounter the following error:
> `"an internal component may be down"`
> Ensure the external target is reachable. Verify by pinging the target from the Cohesity CLI.

Archival Formats

Cohesity supports the following formats for archiving the data to an external target:

- **Incremental with Periodic Full**: Performs an initial full archive and then incremental archival of data with a periodic full every 90 days from the Cohesity cluster to an external target.
- **Incremental Forever**: Performs a first full archival and then incremental forever archival of data from the Cohesity cluster to an external target.

> **Important:** In version 7.2.2 and later, Cohesity enforces Incremental Forever as the default archival method for registering external target. If the target does not support Incremental Forever, Cohesity will automatically fall back to Incremental with Periodic Full. Note that if the target supports Incremental Forever, the option to switch to Incremental with Periodic Full is not available is not available.

Incremental with Periodic Full

An archived snapshot is referred to as an archive.

The first archive of a Protection Group is always a **Reference Full Archive.**

A *Reference Full* forms the basis for Incremental Archives.

With incremental archives, the next archive after a reference full is an **Incremental Archive**. It designates the Reference Full Archive as its reference.

The Incremental Archive contains only data that is different between its reference's snapshot and the snapshot being archived. Change blocks are computed and then changed data is compared to the reference, and the necessary data is sent to the external target. If deduplication is enabled and data did not change, only metadata is transferred.

To reconstruct the snapshot of this incremental archive, both the incremental and the reference must be processed.

Then, the next Incremental Archive constitutes an increment on top of the previous incremental archive, i.e. it designates the previous Incremental Archive as its reference.

To reconstruct its snapshot, both Incremental Archives and the Reference Full Archive must be processed.

In this way, incremental archives form an "archival chain" that starts at the archive in question and ends with the Reference Full.

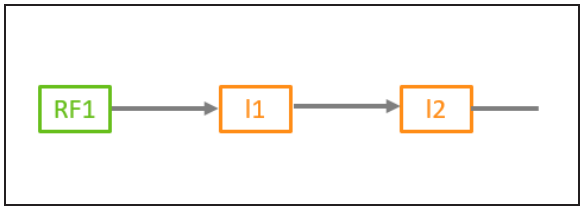

In order to limit the number of archives that must be processed to restore a snapshot, the chain is "broken" every 90 days. At that threshold, the next archive will not be an Incremental Archive, but rather a Full Archive. A regular, non-reference Full Archive is full in the sense that it covers all previous increments since the Reference Full Archive. However, it itself still refers to the Reference Full archive just like the first Incremental Archive.

The next Incremental Archive after a Full Archive will use the Full Archive as a reference.

The threshold is not configurable through the web interface, and is always 90 days regardless of how many archivals have been performed.

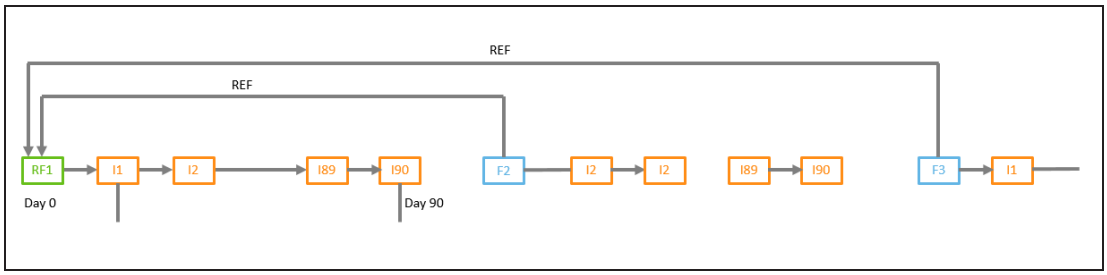

During incremental archives, if the Reference Full Archive becomes too inefficient, mainly if its utilization as a reference drops below 50%, the Reference Full Archive is retired.

Retiring a Reference Full Archive means that the next time a Full Archive needs to be created, it is created as a Reference Full Archive, rather than a regular, non-reference Full Archive. The next Incremental then uses that new Reference Full Archive as its reference.

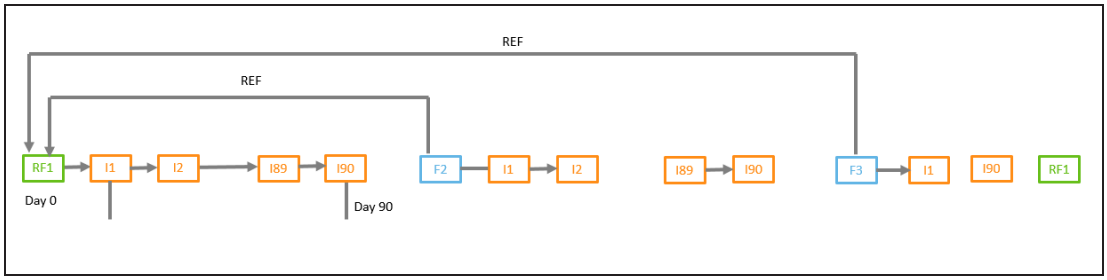

> **Note:** The web interface does not distinguish between Reference Full Archive and non-reference Full Archive.

Each archive, Whether Reference Full, Full, or Incremental, is retained as long as it is referred to by another archive, or as long as its own policy specifies.

For example, for a Protection Job with a 30-day retention policy on archives, on the 90th day, an Incremental Archive will be created, and set to expire on day 120. On day 119, it still refers to an entire archive chain going all the way to the Full Archive from day 1, and therefore even a 119-day old archive is not deleted from the external target.

Furthermore, if that Full Archive refers to a Reference Full Archive that is 6 months old, that Reference Full Archive is also not deleted.

The following table details an example archive workflow for 10 x 100 GiB VMs with a 3:1 deduplication ratio. Each row shows the archive type and amount of data transferred to the External Target based on a change rate.

Change Rate	Cumulative Change Rate	Timeline	Dedup On Incr On	Dedup Off Incr On
N/A	N/A	Day 0	Reference Full - 333 GiB	Reference Full - 1 TiB
1%	1%	Day 7	Incremental - 333 MiB	Incremental - 1 GiB
1%	1%	Day 14	Incremental - 333 MiB	Incremental - 1 GiB
1%	20%	Day 90	Full - 66 GiB	Full - 200 GiB
1%	21%	Day 97	Incremental - 333 MiB	Incremental - 1 GiB
1%	40%	Day 180	Full - 133 GiB	Full - 400 GiB
1%	41%	Day 187	Incremental - 333 MiB	Incremental - 1 GiB
1%	51%	Day 250	Incremental - 333 MiB	Incremental - 1 GiB
1%	60%	Day 270	Reference Full - 333 GiB	Reference Full - 1 TiB
1%	61%	Day 277	Incremental - 333 MiB	Incremental - 1 GiB
1%	80%	Day 360	Full - 66 GiB	Full - 200 GiB

> **Note:** If the **Object Key Pattern** of the S3 View is **Object ID**, you cannot archive the View using the **Incremental with Periodic Full** format.

For the list of external targets and the Protection Groups that supports incremental with periodic full, see Supported Workflows and External Targets.

Incremental Forever

Enable this option to perform a first full archival and then incremental forever archival of data from the Cohesity cluster to an external target. Incremental forever archival eliminates the need for periodic full archive and supports global deduplication, thereby improving the storage utilization of the target.

This option is enabled by default while registering a new external target. For an existing registered external target, you must edit the registered target and enable this option. The migration from **Incremental Archival** to **Incremental Forever Archival** takes place in the next reference archive.

> **Note:**
> - Ensure that the external target does not have WORM/Object Lock with a default locking period configured.

For the list of external targets and the Protection Groups that supports incremental forever archival, see Supported Workflows and External Targets.

Cloud Archive: Incremental with Periodic Full Vs Incremental Forever

The following table lists the differences between Cloud Archive incremental with periodic full archival and Cloud Archive Incremental forever archival.

	Cloud Archive using Incremental with Periodic Full	Cloud Archive with Incremental Forever Archival
Description	Performs an initial full archive and then incremental archive with a periodic full every 90 days. When 90 days elapse, the next archive is full with incremental performed for another 90 days.	Performs an initial full archive and then incremental forever archive. The data is archived from the Cohesity cluster to an external target.
Supported Source/workloads and Targets	See Supported Workflows and External Targets.	See Supported Workflows and External Targets.

	Cloud Archive using Incremental with Periodic Full	Cloud Archive with Incremental Forever Archival
File-level recovery	FLR is supported from the following targets for VMWare and NAS sources with provisioned capacity units: • AWS Glacier(Legacy) Standard • AWS Glacier(Legacy) Gov • AWS Glacier(Legacy) C2S File-level recovery back to the source *is not supported* from the following targets: • QStar Tape • Azure Archive Blob • AWS Glacier-Standard • AWS S3-Glacier-Deep-Archive	Supported
File download	File download is supported from the following target with provisioned capacity units: • Glacier-Standard • Glacier-C2S • S3-Gov-Glacier-Deep-Archive File download *is not supported* from the following targets: • QStar Tape • Azure Archive Blob • AWS Glacier-Standard • AWS S3-Glacier-Deep-Archive	Supported
Cloud Provider Life Cycle Management (LCM)	Entity recovery is supported with Cloud provider LCM from AWS S3 even if Cloud provider LCM is set on the S3 bucket. No recoveries (entity recovery or FLR) are supported from Azure if LCM is configured on the Azure container.	Not supported
CloudRetrive	Supported	Supported

	Cloud Archive using Incremental with Periodic Full	Cloud Archive with Incremental Forever Archival
Deduplication	Deduplication is performed at the Protection Group level thereby utilizing more storage space on the external target	Deduplication is performed at the external target level thereby utilizing less storage space on the external target.
Space reclamation	Does not require any external compute. Space reclamation for older snapshots is performed after a periodic full is done.	Requires external compute (lambda functions) for garbage collection. Continuous space reclamation using external compute is performed as soon as data expires.
WORM support*	Yes	Yes (Early Access feature)

*Cohesity does not support WORM for CloudArchive Direct.

For the list of targets that support WORM, see Supported Workflows and External Targets.

For information on configuring WORM, see Create or Edit a Standard Policy.

Migrate External Targets with Incremental with Periodic full to Incremental Forever

If you want to migrate an existing external target registered with Incremental with Periodic Full archival format in version 7.2.1 or earlier to Incremental Forever, perform the following steps:

1. Confirm the external target supports incremental forever archival format. For the list of supported external targets, see Supported workflow and External Targets.
2. From the Cohesity UI, navigate to the registered external target (**Infrastructure** > **External Target**).
3. Click the Actions menu (⋮) next to the target and select **Edit**.
4. Enable incremental forever as the archival format.
5. If you are enabling Incremental Forever on:
    a. Azure hot tier, you must create Azure Function App. For more information, see Function App.
    b. AWS S3 targets, you must run the **CloudFormation Template** on the AWS console using your AWS account. For more information, see CloudFormation template.
6. Click **Save**.

CloudFormation Template

If you are enabling incremental forever archival for AWS S3 targets, then you must run the **CloudFormation Template** on the AWS console using your AWS account. This will create

a Lambda function, *<S3BucketName>* for efficient garbage collection and an IAM role associated with it. You can download the template by clicking the **CloudFormation Template** link available in the product UI, just below the **Archive Object Lock** option, while registering AWS S3 as an external archival target.

> **Note:**
> - You can use the same or different IAM user to run the CloudFormation Template in your AWS account and to register the AWS target in the Cohesity cluster.
> - The CloudFormation template creates an empty Lambda function. This empty Lambda function will be populated with the Python code once the external target is registered with the Cohesity cluster.

When running the CloudFormation Template, ensure that:

- The IAM user running this template has the following permissions to create the necessary IAM roles and policies.
  - `s3:CreateBucket`
  - `s3:ListBucket`
  - `s3:PutObject`
  - `s3:GetObject`
  - `iam:CreateRole`
  - `iam:GetRole`
  - `iam:PassRole`
  - `iam:CreatePolicy`
  - `iam:PutRolePolicy`
  - `iam:GetRolePolicy`
  - `lambda:CreateFunction`
  - `lambda:GetFunction`
  - `cloudformation:CreateUploadBucket`
  - `cloudformation:CreateStack`
  - `cloudformation:ListStacks`
  - `cloudformation:GetTemplateSummary`
- For the list of permissions required for an IAM user to register an AWS S3 target for archival with incremental format, see Cloud Archive and CloudArchive Direct (Incremental Forever) - S3.

- The storage bucket name you specify while running this CloudFormation template should be the same name you used for AWS external target registration.

To run the CloudFormation template, perform the following steps:

1. In the **AWS Management Console**, navigate to **Services** and search for CloudFormation.
2. Click **Create stack** to create a stack under AWS CloudFormation using the downloaded CloudFormation Template.
3. In the **Prerequisite - Prepare template** page, select the **Template is ready** option and then select **Upload a template** file.
4. Click **Choose File** to browse and upload the JSON CloudFormation template that you have downloaded. Click **Next**.

   The Specify stack details page appears.
5. Provide a **Stack name**.
   1. Enter the Bucket name in the **BucketName** field. The bucket name should match the name you have provided while registering the external target.

      If you have configured **Amazon PrivateLink Interface Endpoints** with the Cohesity cluster to access the AWS S3 external targets for archival, then you must also provide values for the following parameters:

      - **NetworkSecurityGroup**: Enter the security group ID to associate with the lambda function.
      - **VPCSubnets**: Enter one or more subnet IDs in the VPC to use for the lambda function.
        The following screenshot highlights the parameters you need to configure for Amazon PrivateLink Interface Endpoints:

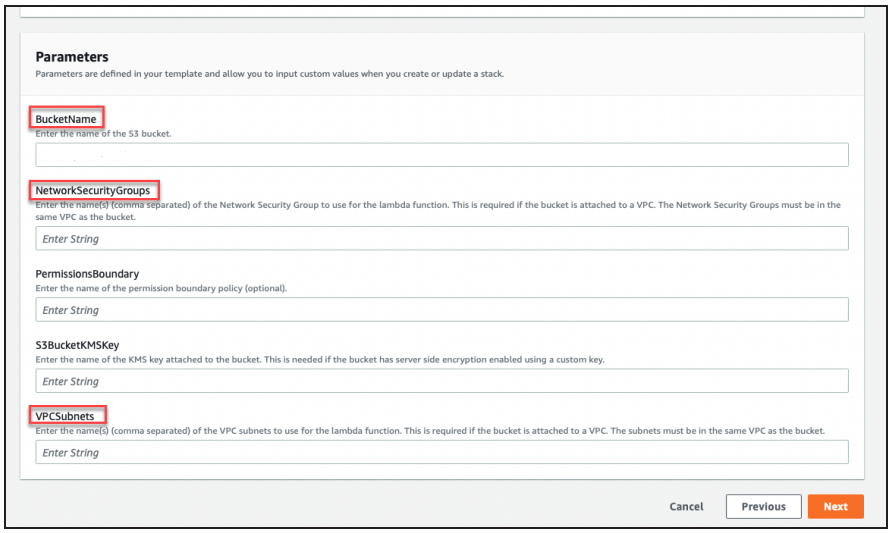

6. Click **Next** > **Next**.

You can skip the **Configure stack options** page.

7. Review the template and select the checkbox to acknowledge the creation of IAM resources, and click **Create Stack**.

On success the IAM resources are created and you can view the details under Resources tab of that stack. You can refresh the page to view the Lambda function, *<S3BucketName>* created for efficient garbage collection and an IAM role associated with it. Note that the Lambda function will be populated with the Python code only after the external target is registered with the Cohesity cluster.

Update Lambda Function for Using Amazon PrivateLink Interface Endpoints

If you have configured **Amazon PrivateLink Interface Endpoints** with the Cohesity cluster to access the AWS S3 external targets for incremental forever archival, then you must update the Lambda function, *<S3BucketName>* as shown below:

Existing code in the lambda function	Update to

`s3 = boto3.client(S3)`	`s3 = boto3.client('s3',endpoint_url='https://bucket. <vpce-id>.s3-<region>.vpce.amazonaws.com')`
	For example:
	`s3 = boto3.client('s3',endpoint_url='https://bucket.vpce-04bc667ade60dcbda-kwt8a5xd.s3.us-west-2.vpce.amazonaws.com')`

Example of the updated lambda function:

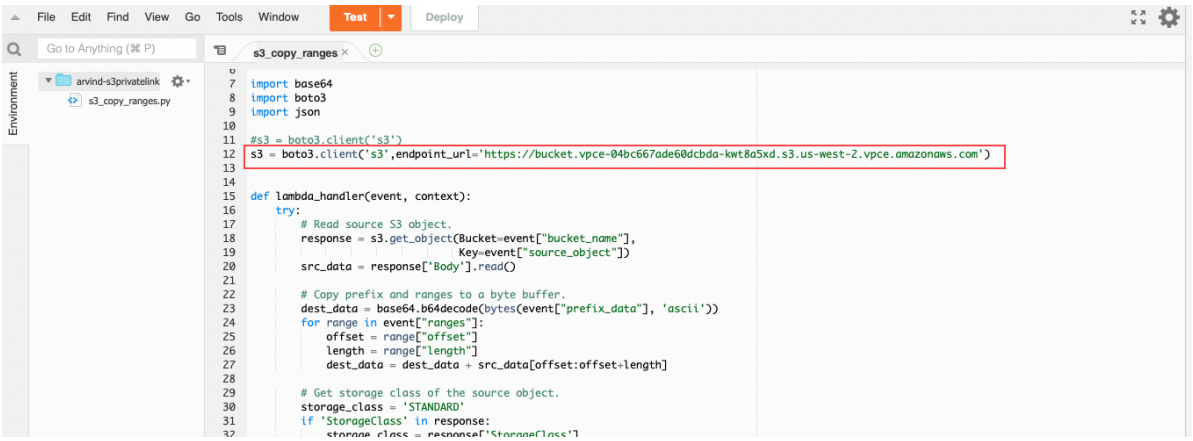

Azure Function App

To archive data to an Azure target (hot blob, cold blob, or cool blob) using the Incremental Forever archival format, you must create an Azure Function App in the Azure portal using your Azure account.

The Azure Function App enables cost-efficient garbage collection. You can create the Function App in one of the following ways:

- Create the Function App directly from the Azure portal (recommended).
- Download the Azure Resource Template from the Register External Target page and upload it to the Azure portal.

> **Note:**
> - Cohesity Cloud Edition (legacy) and Cohesity Cloud Edition clusters do not require an Azure Function App for Incremental Forever archival if the cluster and storage target are in the same region. Therefore, you can proceed with Azure external target registration without specifying a Function App Name or Function App Deployment Key—simply leave these fields empty or enter placeholder values. Validation for these fields is disabled for Cohesity Cloud Edition and Cohesity Cloud Edition (legacy) clusters.
> - The Linux Consumption plan has moved to legacy status and is not getting any new features or language versions. The option to host function apps on Linux in a Consumption plan is retiring on September 30, 2028, after which it can no longer be deployed through the portal. Additionally, function apps still running the end-of-life v3 runtime on Linux in a Consumption plan stop running after September 30, 2026, unless they are upgraded to version 4 (v4). Python version 3.9 is no longer supported by Azure.

Option 1: Create Azure Function from the Azure Portal

Prerequisites

- **Subscription** and **Resource Group**: The values you provide must match the Azure storage account you are registering as an external target.
- **Runtime Stack**: Python.
- **Region**: The same region as the Azure storage account you are registering as an external target.
- **Plan Type**: Select Functions Premium or App Service.
- **Storage Account**: Should be the Azure storage account you registered as the external target.

All resources must be in the same Azure region.

To create the Azure Function App:

1. Log in to the Azure portal and select **Function App**.
2. Click **+Create**.
3. Select either **Functions Premium** or **App Service** as the hosting option and click **Select**.

> **Note:** Consumption or Flex Consumption hosting plans are not supported.

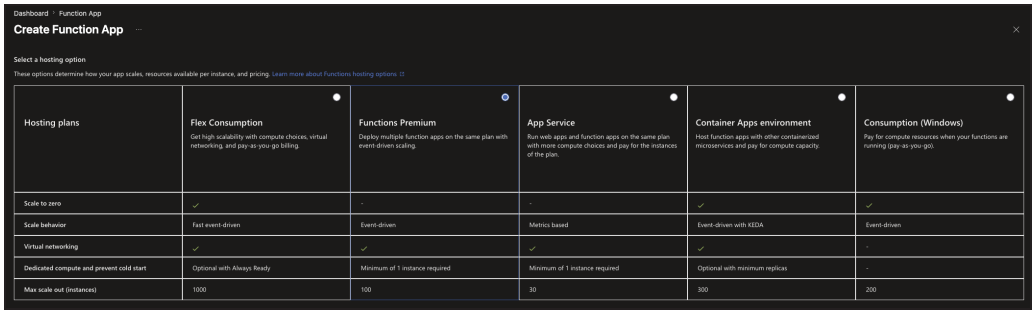

The Create Function App (Consumption) screen appears.

4. On the **Basics** Tab, enter or select the following values, and then click **Next: Storage >**:
   - **Subscription and Resource Group**: Ensure the values you provide are the same as the Azure you register as an external target.
   - **Function App Name**: Provide a name for the Function App and disable the **Try a secure unique default hostname (preview)** option.
   - **Do you want to deploy code or container image**: Select **Code**.
   - **Operating System**: Select **Linux**.
   - **Runtime Stack**: Select **Python**.
   - **Version**: Select **3.12** or the latest Python version.
   - **Region**: Ensure the Region you select for the Function App is the same as the Azure you register as the external target.
   - **Linux Plan**: Select the Azure Service Plan that you created.
   - **Pricing plan**: Select **Elastic Premium EP1 (210 total ACU, 3.5 GB memory, 1 vCPU)**.
   - **Zone redundancy**: Select **Disabled**.

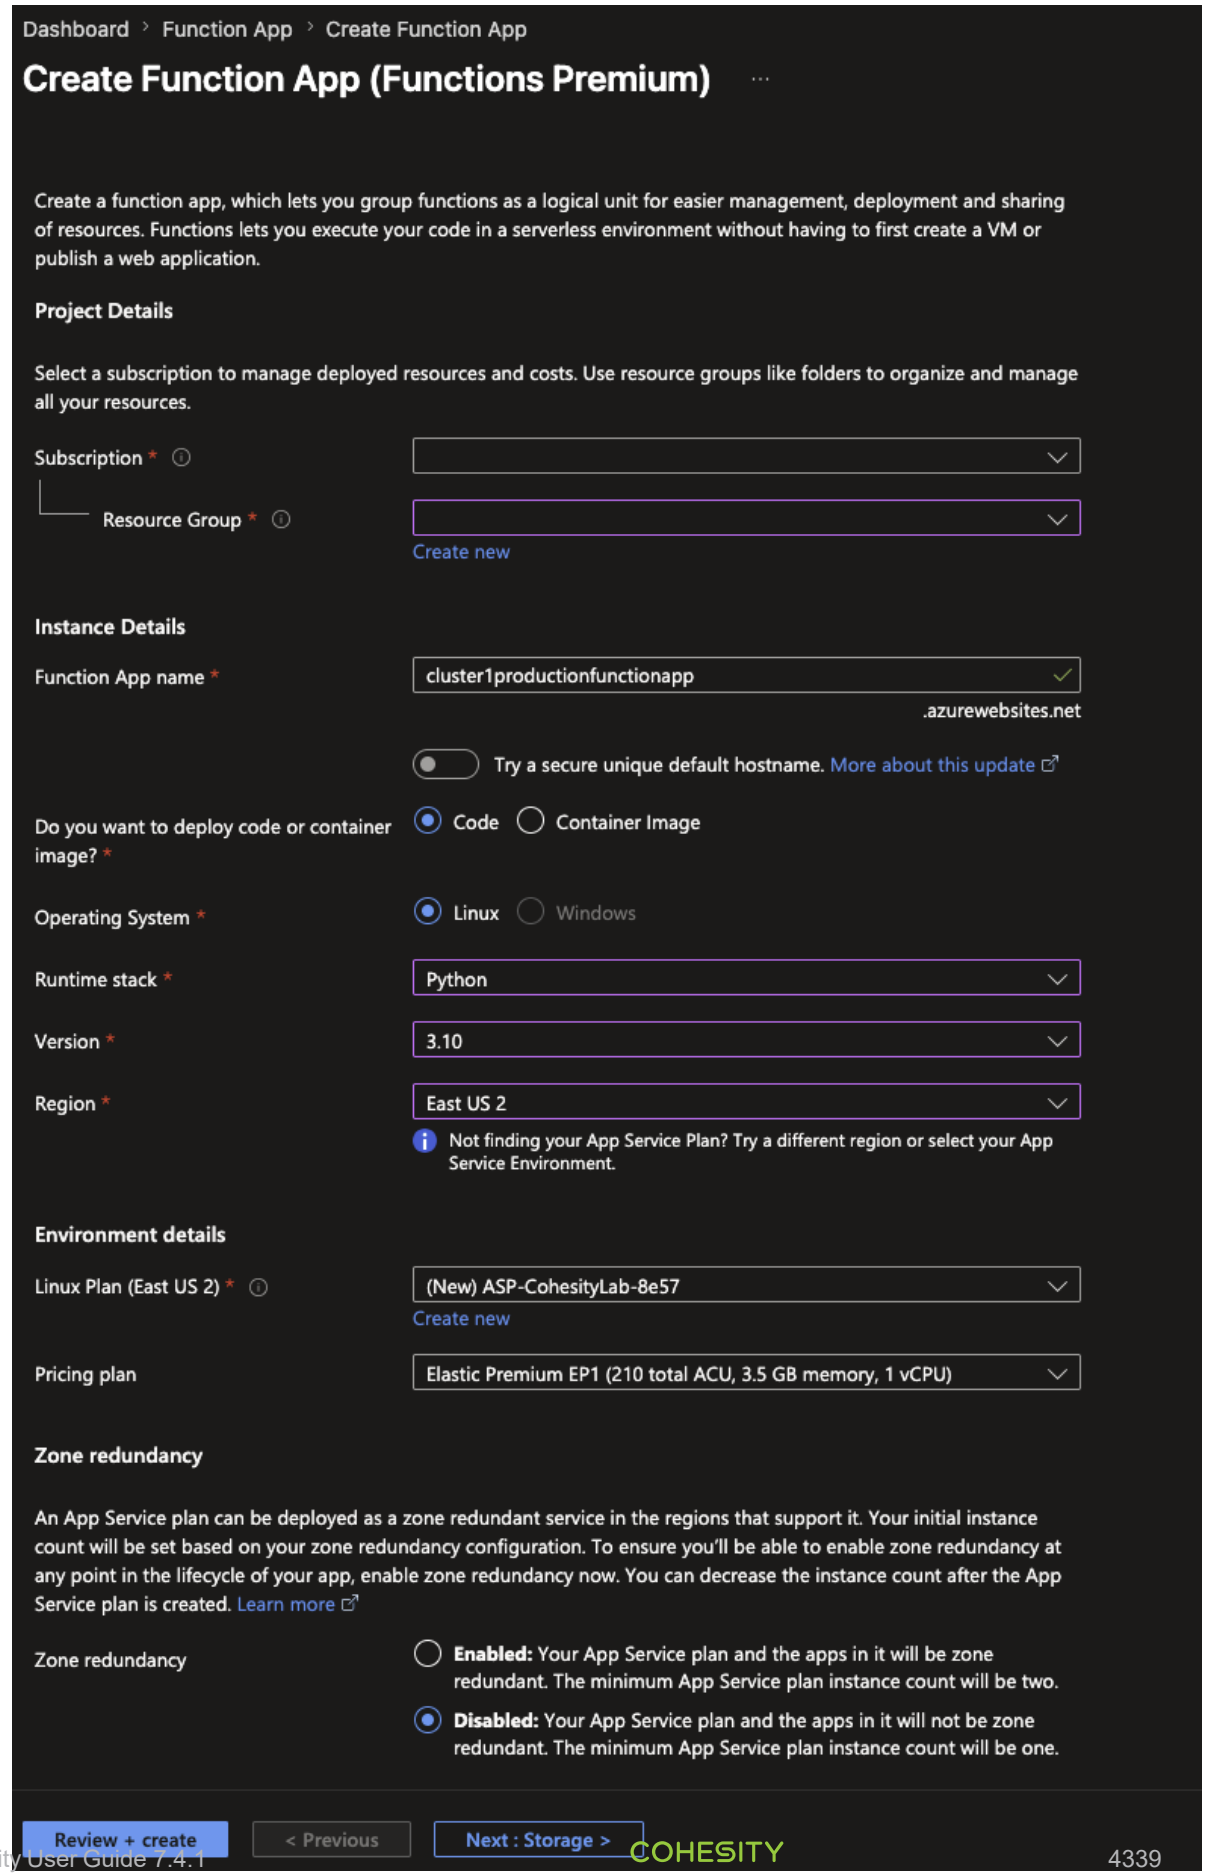
Dashboard › Function App › Create Function App
Create Function App (Functions Premium)
Create a function app, which lets you group functions as a logical unit for easier management, deployment and sharing of resources. Functions lets you execute your code in a serverless environment without having to first create a VM or publish a web application.
Project Details
Select a subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.
Subscription *
Resource Group *
Create new
Instance Details
Function App name *
cluster1productionfunctionapp
.azurewebsites.net
Try a secure unique default hostname. More about this update
Do you want to deploy code or container image? *
Code
Container Image
Operating System *
Linux
Windows
Runtime stack *
Python
Version *
3.10
Region *
East US 2
Not finding your App Service Plan? Try a different region or select your App Service Environment.
Environment details
Linux Plan (East US 2) *
(New) ASP-CohesityLab-8e57
Create new
Pricing plan
Elastic Premium EP1 (210 total ACU, 3.5 GB memory, 1 vCPU)
Zone redundancy
An App Service plan can be deployed as a zone redundant service in the regions that support it. Your initial instance count will be set based on your zone redundancy configuration. To ensure you'll be able to enable zone redundancy at any point in the lifecycle of your app, enable zone redundancy now. You can decrease the instance count after the App Service plan is created. Learn more
Zone redundancy
Enabled: Your App Service plan and the apps in it will be zone redundant. The minimum App Service plan instance count will be two.
Disabled: Your App Service plan and the apps in it will not be zone redundant. The minimum App Service plan instance count will be one.
Review + create
< Previous
Next : Storage >

5. On the **Storage** tab, select the Azure account that you registered as the external target. You can retain the default values for all the other fields on the **Storage** tab.

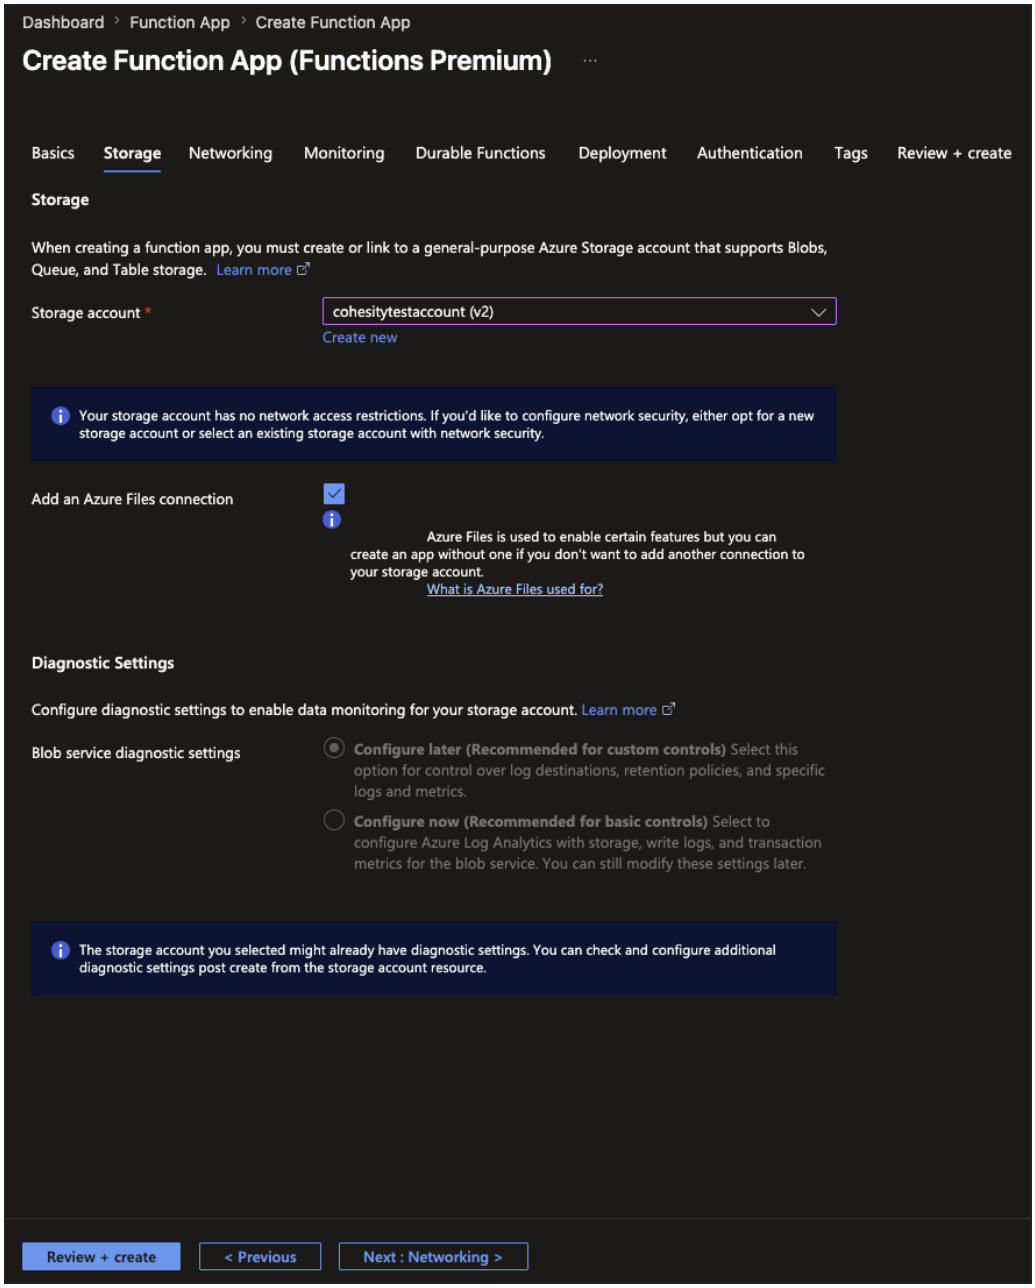

6. Click **Next : Networking >**.
7. On the **Networking** tab, enter or select the following values, and then click **Next: Monitoring >**:
    - **Enable public access**: Select **On**.
    - **Enable virtual network integration**: Select **Off**.

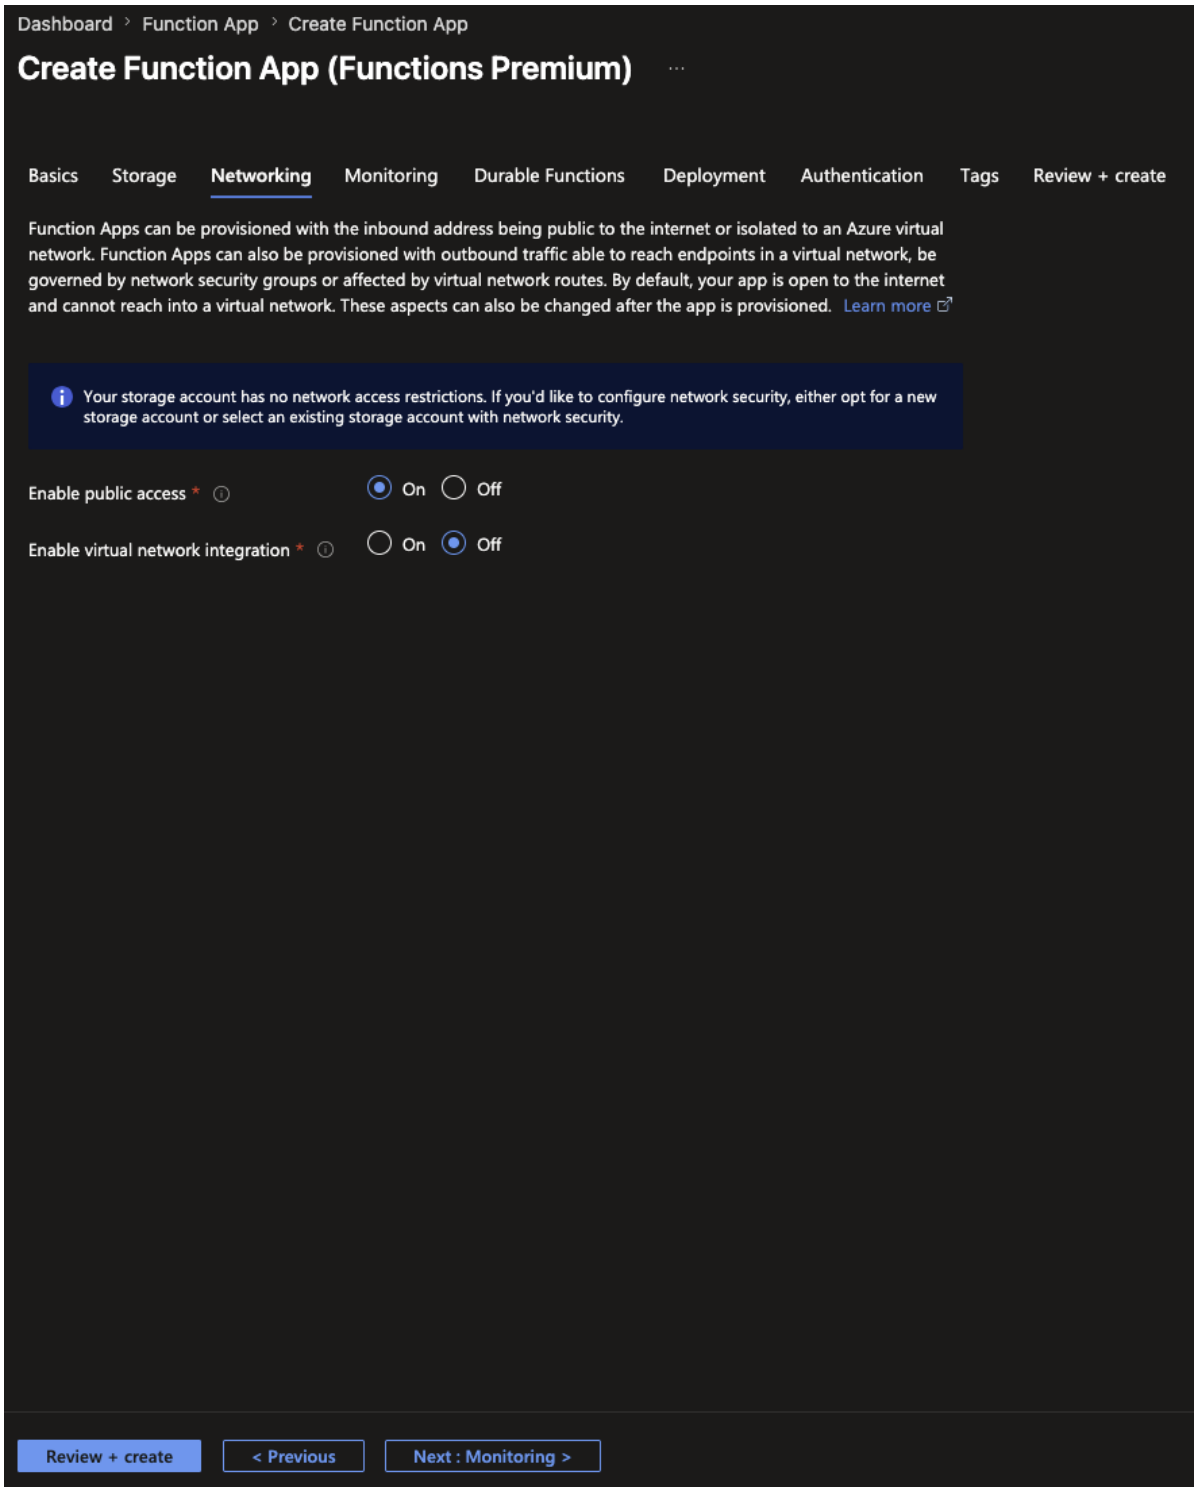
Dashboard › Function App › Create Function App
Create Function App (Functions Premium)
Basics Storage Networking Monitoring Durable Functions Deployment Authentication Tags Review + create
Function Apps can be provisioned with the inbound address being public to the internet or isolated to an Azure virtual network. Function Apps can also be provisioned with outbound traffic able to reach endpoints in a virtual network, be governed by network security groups or affected by virtual network routes. By default, your app is open to the internet and cannot reach into a virtual network. These aspects can also be changed after the app is provisioned. Learn more
Your storage account has no network access restrictions. If you'd like to configure network security, either opt for a new storage account or select an existing storage account with network security.
Enable public access * On Off
Enable virtual network integration * On Off
Review + create
< Previous
Next : Monitoring >

> **Warning:** Enabling Public access will expose the function app to the internet. If you are looking to reduce the access through private links, see Archive Data Over Azure Private Endpoints.

8. (*Optional*) On the **Monitoring** tab, set **Enable Application Insights** to **No**, and then click **Next: Durable Function >**.
9. (*Optional*) On the **Durable Functions** tab, select the desired option, and then click **Next: Deployment >**.
10. On the **Deployment** tab, enter or select the following values, and then click **Next: Authenitcation >**:
    - **Continuous deployment**: Select **Disable**.
    - **Basic authentication**: Select **Enable**.

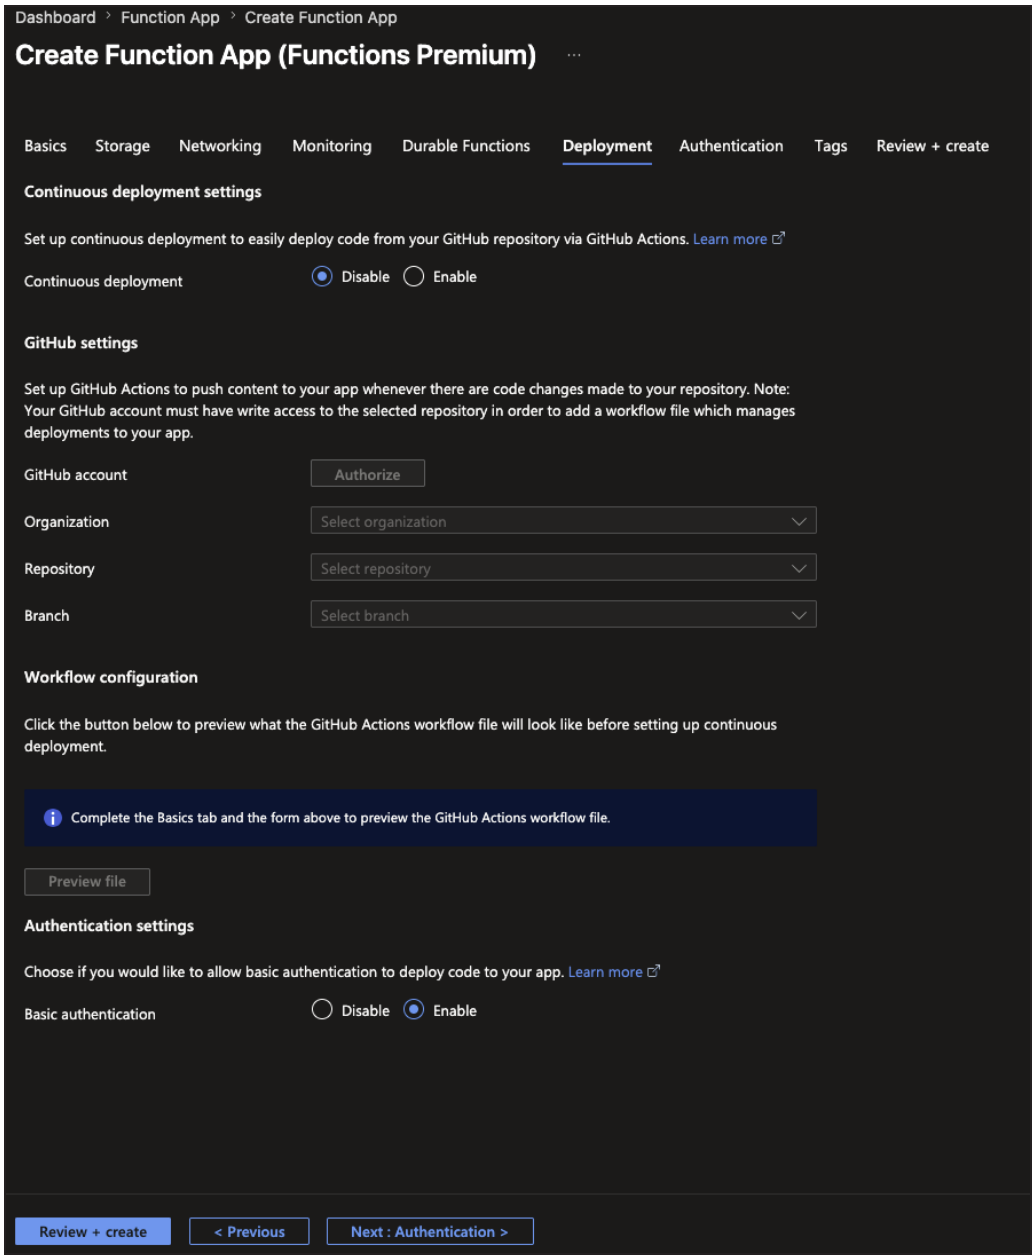

11. On the **Authentication** tab, enter or select the following values, and then click **Next: Tags >**:
    - **Host storage (AzureWebJobsStorage)**: Select **Secrets** as **Authentication type**.

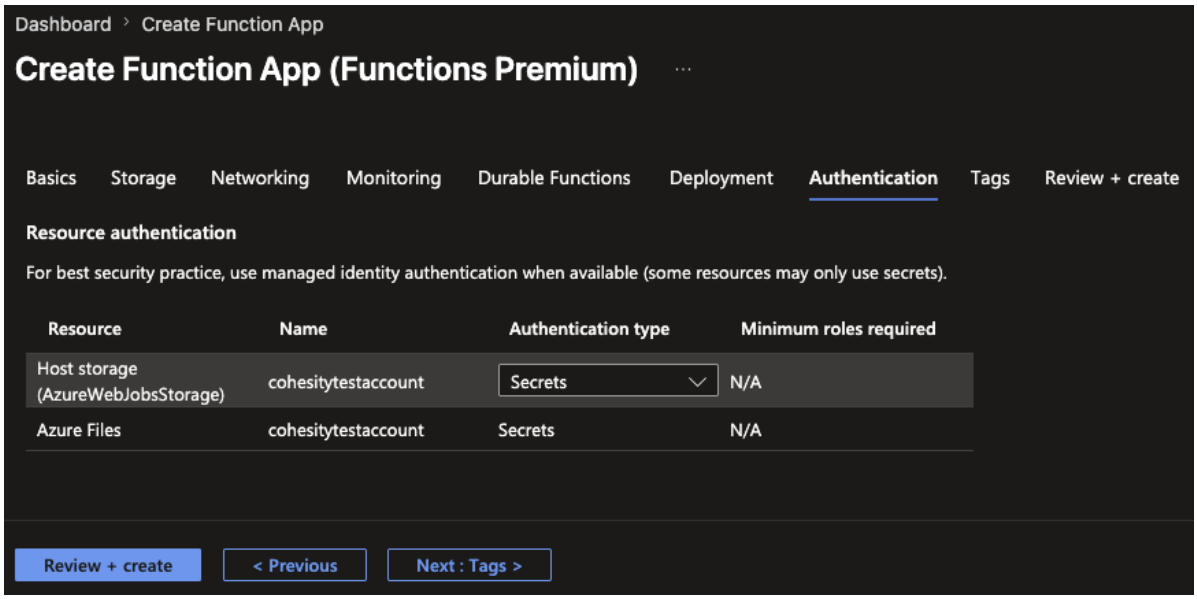

12. On the **Tags** tab, retain the default values and then click **Next: Review + create >**.
13. On the **Review + create** tab, review the configuration, and then click **Create**.

    The Function App deployment is initiated. Once deployed successfully, you can obtain the Function App keys required for registering the Azure as an external target with the Cohesity cluster. For more information, see Obtain the Function App Key of the Function App.

Option 2: Create Azure Function App using Azure Resource template

You can download the template from here. The downloaded template zip file (**ART.zip**) contains two files: **template.json** and **parameters.json**.

Creating the Azure Function App using the Azure Resource template involves the following steps:

- Create Template Spec
- Deploy Template Spec

Create Template Spec

To create the Template Spec:

1. Log into Azure Portal.
2. Navigate to the resource group where you want to create the Function App.

   To simplify resource tracking and troubleshooting, it is recommended to create the Function App in the same resource group as the storage account used for the external target configuration.
3. Click **+ Create**.

4. In the search field, search for **Template Specs** and select **Template Specs** from the result list.
5. From the **Template Specs** tile, click the **Create** drop-down, and then select **Template Specs**.

   The **Create new template spec** page appears.
6. On the **Basic** tab, provide the following information:
   - **Name**: Name of the Template Spec.
   - **Subscription** and **Resource Group**: These fields are auto-populated.
   - **Location**: Region where the Template Spec is created.
   - **Description**: (*Optional*) Provide a description.
   - **Version**: (*Optional*) Enter a version for the Template Spec (example, 1.0).

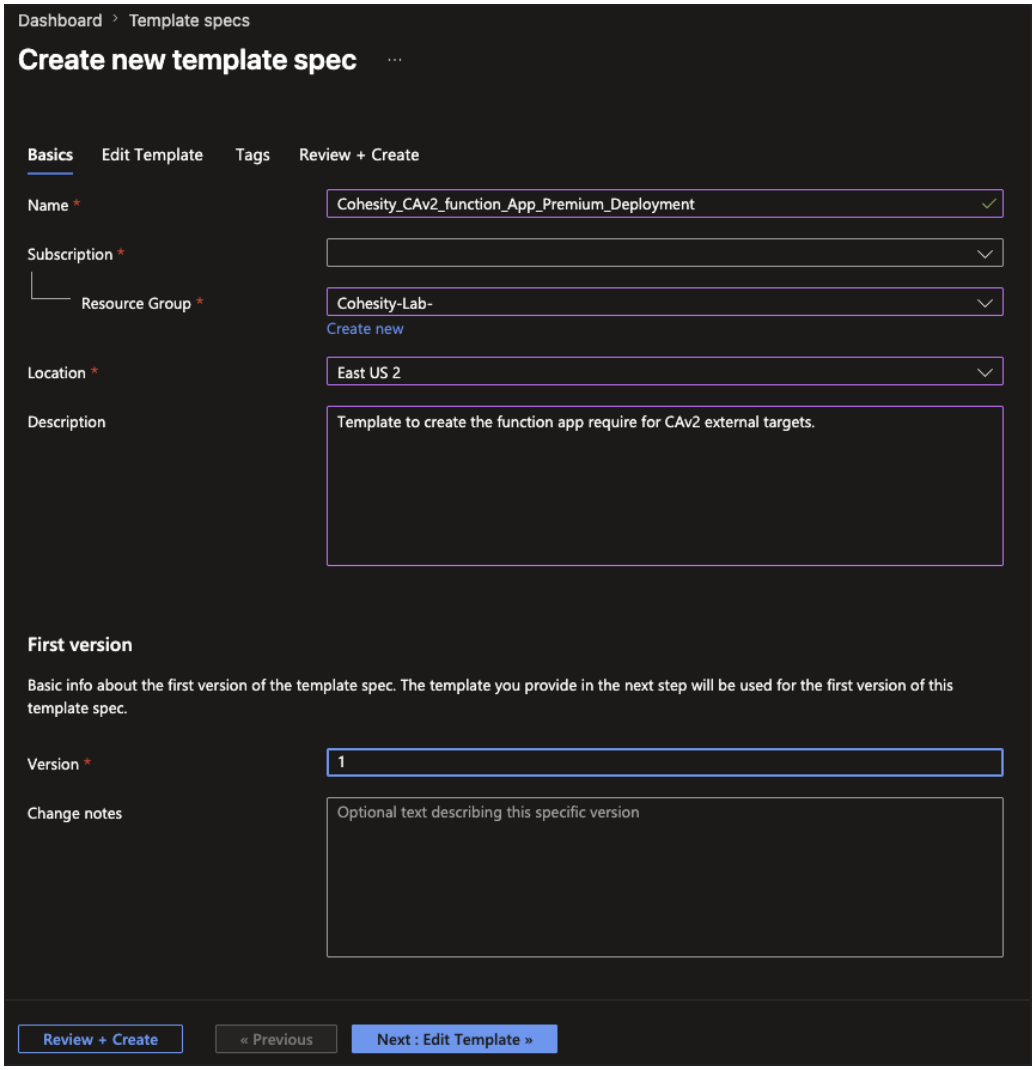

> **Note:** You can select a different Resource Group, Subscription, and location. However, when deploying the function app, the storage account and the app must be in the same region and resource group.

7. Click **Next : Edit Template**.
8. Unzip the downloaded template folder (**ART.zip**) and copy the content from the file **template.json**.
9. In the **Template Content** section of the **Edit Template** tab, replace the existing content with the content from the **template.json** file.

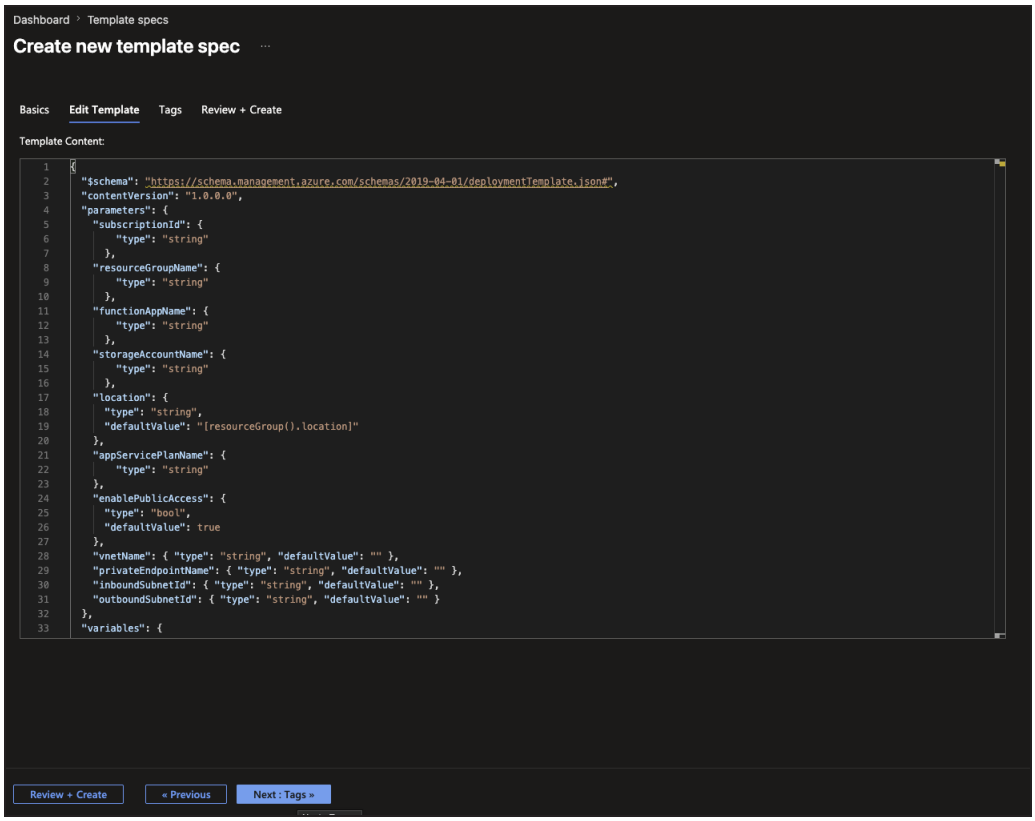

10. Click **Next: Tags**.
11. (*Optional*) On the **Tags** tab, add tags for the Template Spec.
12. Click **Next: Review + Create** and review the summary page.
13. Click **Create**.

    Azure initiates the Template Spec creation.
14. Navigate to the Resource Group that contains the Template Spec.
15. Click **Refresh** on the Resource Group menu to view the newly created Template Spec.

Deploy Template Spec

Once the Template Spec is created, it must be deployed.

To deploy the Template Spec:

1. Navigate to the Resource Group that contains the Template Spec.
2. Click on the Template Spec you created.
3. Click **Deploy**.

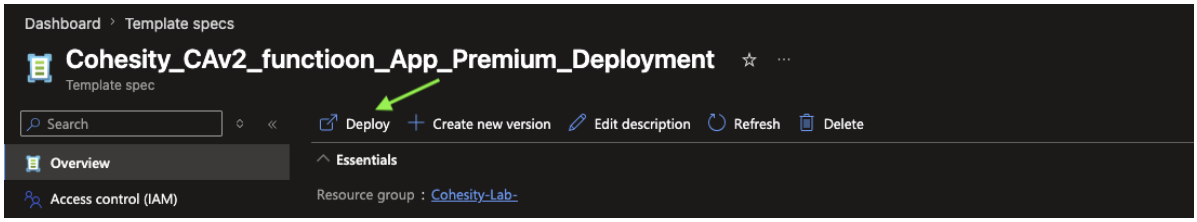

4. Verify the **Project details** section is correct.

   The **Subscription** and **Resource Group** fields are auto-populated.

5. Under **Instance details**, verify the following:
   - **Subscription Id**: Verify the correct Azure subscription is selected. This field is auto-populated.
   - **Region**: Verify the correct Azure region is selected. This field is auto-populated.
   - **Function App Name**: Enter the name for the Function App (The name will be assigned as part of the Function URL).
   - **Storage Account Name**: Enter the name of the storage account that will store the archive data.
   - **App Service Plan Name**: Name of the Function App service. Not related to the Function URL.
   - **Enable Public Access**: Select **true**.

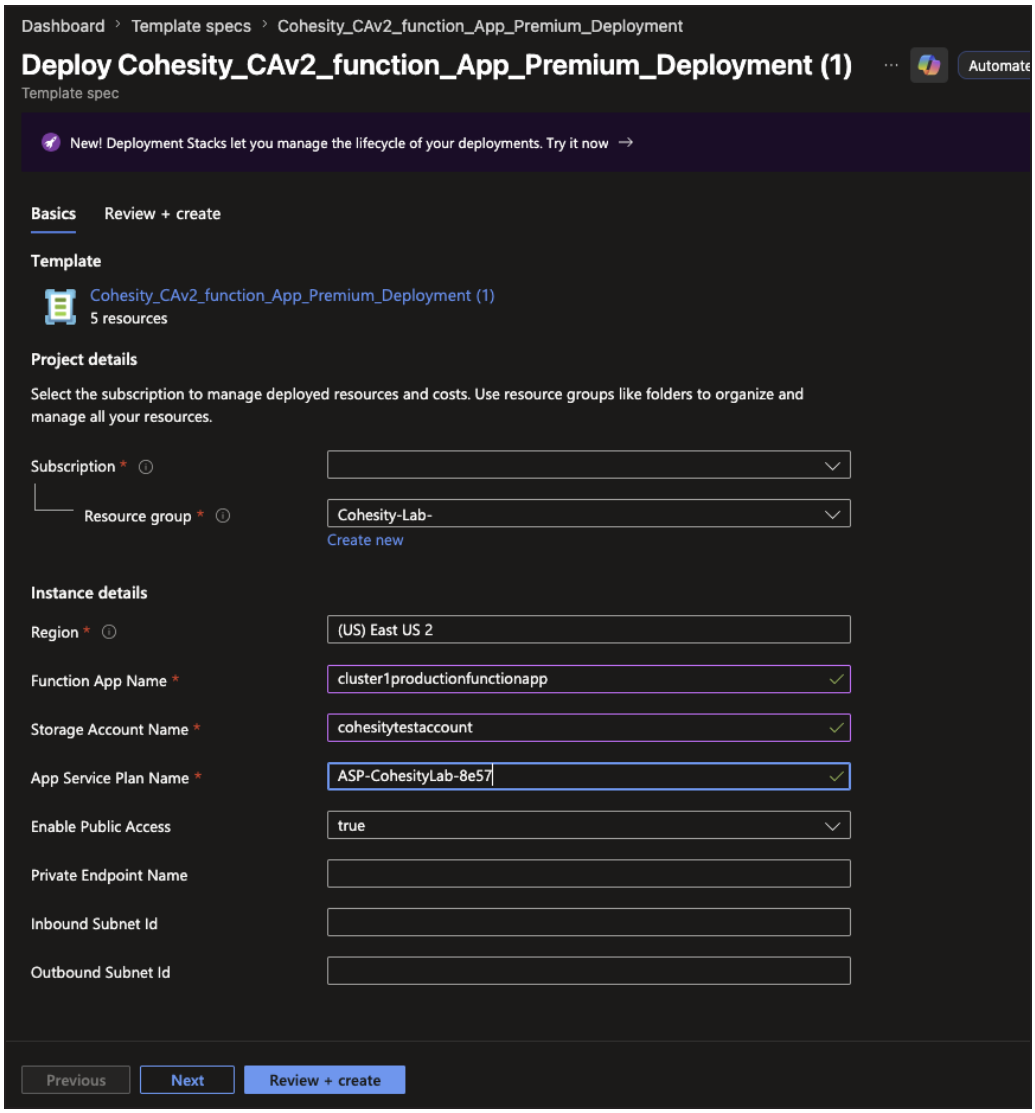

To deploy the function on private links, see Archive Data Over Azure Private Endpoints.

6. Click **Review + Create**.
7. Click **Create**.

   Wait for the resource deployment to complete.
8. When the deployment completes, click **Go to resource group**.

   The Resource Group now contains two new resources: an App Service Plan and a Function App.

   The Function App deployment is initiated. Once deployed successfully, you can obtain the Function App keys required for registering the Azure as an external target with the Cohesity cluster. For more information, see Obtain the Function App Key.

Configure Azure Function App for Incremental Forever Archival

Perform the following steps to configure environment variables for your Azure Function App:

1. On the Function App page, select Environment variables from the left navigation under Settings.
2. In App settings, click **Add** to add the required environment variables, then click **Apply**:
    1. **SCM_DO_BUILD_DURING_DEPLOYMENT**: Set to **true**.
    2. **ENABLE_ORYX_BUILD**: Set to true.
    3. (*Optional*) **WEBSITE_DNS_SERVER**: Set this variable to use Azure DNS for DNS resolution in the Azure Function.
3. Click **Confirm** to save the changes.

Obtain the Function App Key of the Function App

Go to the Function App deployed on the Azure portal.

1. Select **Overview**.
2. Verify the **Status** field displays **Running**.
3. Click **Get publish profile** to download the profile.

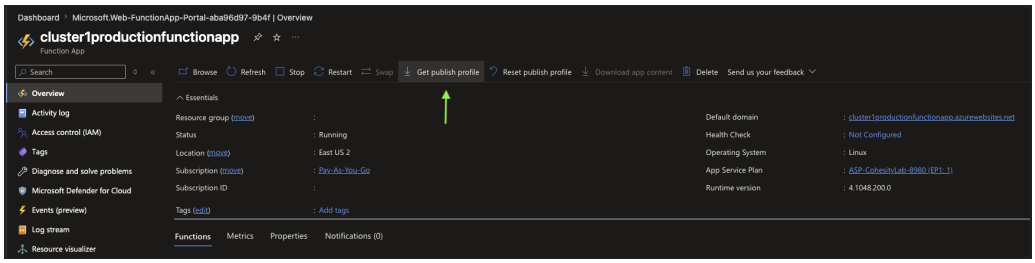

4. Open the downloaded file in a text editor and copy the value of **profileName** and **userPWD**.

    The **profileName** is the **Function App Name**, and **userPWD** the **Function App Deployment Key** you need to provide while registering the Azure as an external target.

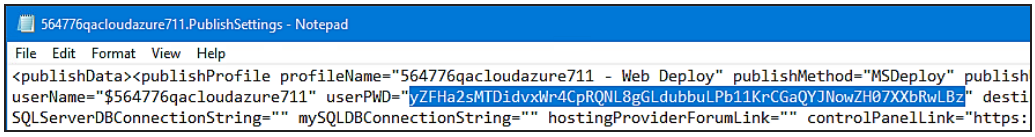

Considerations

Cohesity recommends using the Azure Functions Consumption plan for archival workflows, as it is the minimum required and the only officially qualified plan for Cohesity archival operations. Other Azure Function plans may work; however, they are not officially qualified by Cohesity.

Azure Functions support virtual network (VNet) integration, depending on the Azure Function App plan and its configuration. When using VNet integration or other networking features, ensure that the Azure Function App and the associated storage account allow the required connectivity for archival operations.

For information on plan-specific networking capabilities and configuration requirements, refer to the official Azure documentation.

Register an External Target for Cloud Tier

To register an external target for Cloud Tier, perform the following steps:

1. In the Cohesity Dashboard, select **Infrastructure > External Target**.
2. Click **Add External Target** located at the top right of the page.
3. Provide a name and add a description in the **New Target** and **Description** fields.
4. Select **Tiering** as the External Target's **Purpose**.
5. From the **Type** drop-down list, select the external target you want to register with the Cohesity cluster for Cloud Tiering and enter the information based on the target type you select..

   AWS Glacier C2S

   You must contract with a Commercial Cloud Services (C2S) agency to provide a private instantiation of the AWS commercial cloud.

Field Name	Description
Base URL	URL of the Commercial Cloud Services (C2S) server that contains the Amazon storage provided by the C2S agency.
Agency	The name the agency that is providing the Commercial Cloud Services (C2S) server.
Mission	The account name for accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Role	The IAM role to use when accessing the Amazon storage which is provided by the C2S agency.
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.

Field Name	Description
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. > **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Client Certificate Password	The password to read the Client Private Key when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
CA Trusted Certificate	Upload the Certificate Authority Trusted Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Certificate	Upload the Client Certificate file to access the Commercial Cloud Services (C2S) server which is provided by the C2S agency.
Client Private Key	Upload the Client Private Key file to use when accessing the Commercial Cloud Services (C2S) server which is provided by the C2S agency.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for AWS.

AWS S3 Gov

For IAM user, if you do not already have an access key to use with Cohesity, create a new one in the AWS Management Console (root account credentials required). In AWS, go to **IAM** (AWS Identity and Access Management) and create a new access key ID and secret access key ID. Use this access key ID and secret access key to complete the corresponding fields below.

> **Warning:** Do not apply an AWS Object Lifecycle Policy with the Expiration action to an AWS S3 External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Region	Geographic area to use for Amazon storage. For example, if your region name is US West (N. California), the AWS Region value should be “us-west-1”. **Note:** The AWS Region value is case-sensitive. It must be in lowercase and cannot contain uppercase letters. To know your AWS Region based on the Region name, see AWS Service Endpoints.
Access Key ID	Access key ID used to access Amazon storage.
Secret Access Key	Secret access key address to access Amazon storage.
(For IAM Role Only) **Account ID**	The account ID where the CE is running
(For IAM Role Only) **IAM Role**	The iAM Role for external target registration. Example: CohesityArchiveRole.

Note: IAM Role for AWS S3 Gov is not supported.

For most regions, the Cohesity uses `s3-<region>.amazonaws.com` to construct the S3 endpoints instead of `s3.<region>.amazonaws.com`.

Example

`s3-us-east-2.amazonaws.com`, where `s3-us-east-2.amazonaws.com` is the CNAME for `s3.us-east-2.amazonaws.com`.

However, the Cohesity uses `s3.<region>.amazonaws.com` for regions listed below:

- China
- us-east-1
- S3 acceleration

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for AWS.

You can archive data to Azure Data Box by registering Data Box as NAS External Target on Cohesity cluster. For more information, see How to use Cohesity Microsoft Azure Data Box Integration.

Azure Hot Blob Standard or Gov

After creating a storage account in Azure, if you have not done so already, go to Blob service and create a new container to use with Cohesity. Provide the container name, storage account name and either of the keys generated by Azure to the Cohesity Dashboard.

> **Note:** Cohesity recommends you to use the **Standard** (optimized for high capacity and high throughput) performance tier and not the Premium performance tier when creating a new Azure Blob storage account and container to register with the Cohesity cluster as an external target.

Field Name	Description
Container Name	Name of the existing container where you will store objects. This cannot be modified after you register the Target.
Storage Account Name	Account name used to access Azure storage.
Storage Access Key	Storage access key to access Azure storage.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for Azure.

GCP Multi-Regional or GCP Regional

In the Google Cloud Platform Console, create a new bucket to use with Cohesity if you have not done so already. To create credentials to use with Cohesity, from the Dashboard go to **IAM & Admin > IAM** and create a service account. Give the service account the Role of Project Owner and select Google Nearline or Coldline. After you create the service account, download and open the credential (JSON) file and use the contents (copy the text in the fields between the double quotes) to complete the fields below.

> **Warning:** Do not apply a Google Cloud Object Lifecycle Policy with the Delete action to any GCP External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Project ID	Project ID of the existing project where you will store objects. This cannot be modified after you register the Target.
Client Email Address	Email address used to access Google Cloud Storage.
Client Private Key	Private key used to access Google Cloud Storage.

For more information, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for GCP.

You can archive data to Google Transfer Appliance by registering Google Transfer Appliance as NAS External Target on Cohesity cluster.

Oracle Object Storage

A tenancy must be created by Oracle for your organization. In addition, an administer must set up the Tenancy and possibly create a compartment to hold your resources or a new user if required. Determine the best Oracle storage tier to use for your organization: Oracle Archival Storage (Archive Storage tier) or Oracle Object Storage (Object Storage tier). In the Oracle Cloud console, create a new bucket to use with Cohesity if you have not done so already. The Storage Tier specified for the bucket must match the specified Cohesity External Target Type (Oracle Archive or Oracle Object).

The user with which the bucket is registered must be associated with the following Oracle policy statements:

- Allow group <*group-name*> to manage buckets in tenancy
- Allow group <*group-name*> to manage objects in tenancy

If you do not already have an access key to use with Cohesity, create a secret key. Use the Access Key ID and Secret Access Key to complete the corresponding fields below.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target. In the Oracle Cloud console, create a new bucket to use with Cohesity if you have not done so already. The Storage Tier specified for the bucket must match the specified Cohesity External Target Type (Oracle Archive or Oracle Object).
Region	Geographic area to use for Oracle storage. For more information, see Data Regions.
Access Key ID	Access key ID used to access the storage. For more information about create and copying the key, see create a secret key.
Secret Access Key	Secret access key address to access the storage. For more information about create and copying the key, see create a secret access key.
Tenancy	This is the Tenancy name that is displayed under TENANCY in the top menu bar of the Oracle Cloud Console. This is not the Oracle Cloud Identifier (OCID) for your Tenancy.

S3 Compatible

Cleversafe is a supported AWS S3 Compatible External Target.

When registering a Cleversafe External Target, complete the settings with the following information:

- Bucket Name - Use the vault name from Cleversafe.
- Access Key ID and Secret Access Key - As provided by Cleversafe.
- Endpoint - Use the endpoint of the Cleversafe appliance.
- AWS Signature Version - Select ver 4.

Do not apply an AWS Object Lifecycle Policy with the Expiration action to an AWS S3 Compatible External Target used for archiving data. Applying a policy with this action may result in data loss.

Field Name	Description
Bucket Name	Name of the existing bucket where you will store objects. This cannot be modified after you register the Target.
Access Key ID	Access key ID used to access the storage.

Field Name	Description
Secret Access Key	Secret access key address to access the storage.
Endpoint	The Endpoint mount path. Use the IP address or Fully Qualified Domain Name (FQDN), then a colon (:), followed by the port number. **Note:** Do not include 'http://' or 'https://' in the mount path.
Secure Connection (HTTPS)	Toggle **Secure Connection** on or off.
AWS Signature Version	Select the AWS signing process version for adding authentication information to requests.

To register other S3 compatible targets like Pure FlashBlade S3 bucket, IBM ICOS S3 bucket and Cohesity S3 bucket, see CloudArchive Deployment & CloudRetrieve Recovery for S3-Compatible Storage.

> **Note:** For Cloud Tiering, encryption and compression are always enabled and you cannot disable these options.

6. Enable **Bandwidth Throttling** to limit the maximum data transfer rate *from and to this External Target* during a time window.
    a. **Change Time Zone**: Optionally change the time zone that applies to all the override windows specified.
    b. **Traffic:** Specify the direction of traffic to limit, by specifying a maximum bandwidth:
        - **Upload:** Limit the data transfer rate from the Cohesity cluster to this External Target.
        - **Download:** Limit the data transfer rate from this External Target to the Cohesity cluster.
    c. **Days**: Specify the days of the week to apply the window.
    d. **Time**: Specify all day or a time period for each day.
    e. **Throttle To**: Specify the maximum data transfer rate.Optionally change the rate type.

f. Click **Add**.

g. Repeat these steps to add more upload throttling windows.

If overlapping windows are specified, the latest window that was created takes precedence. For example, if you first create a window from Monday 9 AM - 10 PM with a data rate of 50 Mbps and then a window from 5 PM to 12 PM with a data rate of 100 Mbps, during 5 PM -12 PM the data rate is limited to 100 Mbps.

If both cluster-level override throttling window and an external-target override throttling window is set during the same time period, the lowest **Throttle To** number specified is the maximum allowed data transfer rate. For example, if a cluster-level override throttling window of 8-11 AM is set to 250 Mbps and an External-level override throttling window of 9-10 AM is set to 200 Mbps, during 9-10 AM the maximum allowed data transfer rate is 200 Mbps. Specifying transfer rate limits at the External-target level is described at the end of this procedure. For information on setting up cluster-level bandwidth throttling, see Set Bandwidth throttling at Cluster-level for Tiering to External Targets.

7. Click **Register**.

The Cohesity cluster attempts to connect to the target. A successful connection displays the external targets page with the new target included in the list.

Set Bandwidth throttling at Cluster-level for Tiering to External Targets

You can set bandwidth throttling at the cluster level to limit the data transfer rate to and from all the external targets registered for Tiering on this Cohesity cluster.

1. In the Cohesity Dashboard, navigate to **Inventory** > **External Target**.
2. Select the **Settings** tab.
3. Set the values of the fields displayed in the **Tiering** section:

Field Name	Description
Upload	Set the following options to limit the maximum data transfer rate from the Cohesity cluster to all the External Targets registered on this cluster: 1. **Throttle**: Enable this option and in the Throttle to field, specify the maximum data transfer rate. 2. **Throttling Overrides**: Enable this option and specify a time window during which the maximum data transfer rate you specified for in the above step must be overwritten: a. **Change Time Zone**: Optionally change the time zone that applies to all the override windows specified. b. **Days**: Specify the days of the week to apply the window. c. **Time**: Specify all day or a time period for each day. d. **Throttle To**: Specify the maximum data transfer rate. e. Click **Add**. f. Repeat these steps to add more upload throttling windows. **Note:** If overlapping windows are specified, the latest window that was created takes precedence. For example, if you first create a window from Monday 9 AM - 10 PM with a data rate of 50 Mbps and then a window from 5 PM to 12 PM with a data rate of 100 Mbps, during 5 PM -12 PM the data rate is limited to 100 Mbps.

Field Name	Description
Download	**Throttle**: Enable this option and in the Throttle to field, specify the maximum data transfer rate. 1. **Throttling Overrides**: Enable this option and specify a time window during which the maximum data transfer rate you specified for in the above step must be overwritten: a. **Change Time Zone**: Optionally change the time zone that applies to all the override windows specified. b. **Days**: Specify the days of the week to apply the window. c. **Time**: Specify all day or a time period for each day. d. **Throttle To**: Specify the maximum data transfer rate. e. Click **Add**. f. Repeat these steps to add more upload throttling windows. **Note:** If overlapping windows are specified, the latest window that was created takes precedence. For example, if you first create a window from Monday 9 AM - 10 PM with a data rate of 50 Mbps and then a window from 5 PM to 12 PM with a data rate of 100 Mbps, during 5 PM -12 PM the data rate is limited to 100 Mbps.

4. Click **Save**.

Things to remember

- **When specifying both cluster-level throttling and cluster-level Override throttling windows**: If both cluster-level throttling and a cluster-level override throttling window are set, then the rate specified in the cluster-level override throttling window takes precedence and is used during the specified time window. For example, if the cluster-level throttling is set to 100 Mbps and the cluster-level override throttling window of 9-10 AM is set to 200 Mbps, the maximum allowed data transfer rate during 9-10 AM is 200 Mbps but at any other time, the maximum allowed

data transfer is 100 Mbps.

- **When specifying both cluster-level Override throttling window and External-target Override throttling window**: If both cluster-level override throttling window and an External-target override throttling window is set during the same time period, the lowest Throttle To number specified is the maximum allowed data transfer rate. For example, if a cluster-level override throttling window of 8-11 AM is set to 250 Mbps and an External-level override throttling window of 9-10AM is set to 200 Mbps, during 9-10 AM the maximum allowed data transfer rate is 200 Mbps. Specifying transfer rate limits at the External-target level is described at the end of this procedure.

Add a New External Target for CloudTier

You can add a new external target for CloudTier while the existing external target registered for CloudTier is being removed.

Steps:

1. Verify whether the CloudTier is configured at the Cohesity cluster level or at the Storage Domain level.
    a. If the CloudTier is configured at the Storage Domain level, verify if the Storage Domain has physical storage available to uptier the data back to the Storage Domain with the physical quota being set. If the storage space is less, you must extend the physical quota to uptier the data completely.
    b. If the CloudTier is set at the cluster level, no action is required, proceed with step 2.
2. Remove the external target's disk if you want to remove the existing target from the entire Cohesity cluster. Skip this step if you want to remove the external target only from a set of Storage Domains as other Storage Domains might have data in this external target.
3. Contact Cohesity Support to remove the existing cloud_spill_vec depending on from where you want to remove the existing external target (from a set of Storage Domains or from the entire Cohesity cluster).
4. Wait until uptiering is complete. You can use the **Advanced Diagnostics** statistics (**System** > **Advanced Diagnostics**) on the Cohesity UI to confirm the status.

5. Register a new external target for CloudTier and designate the external target with Storage Domain(s).

   For more information, see Register an External Target for Cloud Tier and Create or Edit Storage Domains.

6. Once the disk of the previous external target is removed, unregister the previous external target. For more information, see Unregister an External Target.

Update Storage Key for Azure External Target

If your organization updates the Azure Storage Key on the Azure Storage Account used for Incremental Forever Archive, you must update the new key information at two places:

1. On the Cohesity cluster: Update the **Storage Access Key** field for the Azure target.
2. On the Azure Function App: Update the **AzureWebJobsStorage** environment variable with the new Azure Storage connection string.

Update Azure Target on the Cohesity Cluster

Perform the following steps on the Cohesity cluster:

1. Navigate to **Infrastructure > External Targets**.
2. Locate the Azure Target configured with Incremental Forever archival.
3. Next to the target, click the pencil icon to edit it.
4. In the **Storage Access Key** field, replace the old key with the new Azure Storage Access Key.
5. Enter the **Function App Name** and **Function App Deployment Key**.
6. Click **Save** to apply the changes.

Update the Azure Function App

Perform the following steps to update the Azure Function App with the New Azure Storage Connection String:

1. Get the New Azure Storage Connection String.
    - In the Azure Portal, navigate to the **Azure Storage Account > Security + Networking > Access Keys** page.
    - On the **Access keys** page, click **Show** next to the Connection string and copy the string.
2. Update the **AzureWebJobsStorage** Environment Variable.
    - In the Azure Portal, navigate to **Azure Function App > Settings > Environment Variables**.
    - Locate the **AzureWebJobsStorage** variable and replace its value with the new Azure Storage connection string.
    - Click **Apply** to save the changes.

AWS Security Token Service Support

Cohesity can secure AWS S3 External Targets with AWS Security Token Service (STS).

Supported S3 AWS Targets

Cohesity supports the following AWS S3 External Targets with AWS STS:

- S3 Intelligent-Tiering
- S3 One Zone-IA
- S3-Standard
- S3-GOV
- S3-C2S
- S3-IA-Standard
- S3-IA-C2S

Configure AWS STS Vault

You can configure the AWS STS vault by using the Cohesity CLI command.

Prerequisites

Before configuring the AWS STS vault, deploy a credentials service by ensuring the following:

- The HTTP endpoint is reachable by the Cohesity cluster.
- The HTTP endpoint accepts POST API requests.
- The HTTP endpoint validates the request body of a POST request and responds with a valid AWS STS token.

The response from the endpoint must generate the JSON output that matches the following syntax:

```
{
    "Version": 1,
    "AccessKeyId": "an AWS access key",
    "SecretAccessKey": "your AWS secret access key",
    "SessionToken": "the AWS session token for temporary credentials",
    "Expiration": "ISO8601 timestamp for when the credentials expire"
}
```

To configure the AWS STS vault:

```
iris_cli vault create name=<name> vault-type=Amazon usage-type=archival com-
pression-enabled=<true|false> encryption-enabled=<true|false> bucket-name-
e=<name> aws-region=<region> tier-type=Amazon-S3-Standard incremental-
enabled=true aws-credential-endpoint=<credential endpoint URL> aws-cre-
dential-blob=<path to local file>
```

To edit the AWS STS vault:

```
iris_cli vault edit id=<vault ID> aws-credential-endpoint=<credential end-
point URL> aws-credential-blob=<path to local file>
```

The following table provides the description of the parameters used for creating the AWS STS vault:

Parameter	Type	Description
`name`	string	Name of the vault.
`vault-type`	string	Type of the vault (Amazon).
`usage-type`	string	Usage purpose of the vault (Archival).
`compression-enabled`	bool	Flag to enable or disable the compression.
`encryption-enabled`	bool	Flag to enable or disable the encryption.
`bucket-name`	string	Name of the bucket or container inside the target vault.

Parameter	Type	Description
`aws-region`	string	AWS region (e.g. us-east-1, us-west-1 etc.) for AWS service account.
`tier-type`	string	Tier of the vault. For the list of supported tier types, see Supported S3 AWS Targets.
`incremental-enabled`	bool	Support incremental archive for archival vaults.
`aws-credential-endpoint`	URL	The endpoint of the credentials service.
`aws-credential-blob`	path	The absolute path of a local file containing the authentication string to send to the credential endpoint.

Kerberos Authentication for NAS External Target

Cohesity supports Kerberos authentication for registration of NAS volumes with NFS versions 4 and above as external targets with the Cohesity cluster.

The following are the use cases supported:

- Using one MIT Kerberos security.
- Using two MIT Kerberos security.
- Cluster in AD using MIT Kerberos security.
- NetApp registered as NAS type of external target and accessed using NFSv4 protocol with Kerberos security.
  The following is the prerequisite procedure for this use case:
  - Login to NetApp.
  - Navigate to **Storage > Storage VMs**. The **Storage VMs** page is displayed.
  - Click the required Storage VM and click the **Settings** tab.
  - Click **Kerberos** and click **Add** to add a Kerberos Realm. The **Add Kerberos Realm** page is displayed.
  - Enter the Name, KDC IP Address, KDC Port, and KDC Vendor.

    When the Kerberos realm is registered, NetApp creates a Service Principal in the KDC.
  - Click **Add** to add a network interface to realm. Enter the Kerberos Interface, Service Principal Name (from Kerberos) , Admin Username, and Admin Password. Click **Save**.

- Click **Save**. Now, the NetApp and KDC can communicate with each other.
- In the Cohesity cluster, Add Kerberos Provider to Cohesity cluster.

  The Cohesity cluster communicates with the KDC and creates two service principles in the KDC.

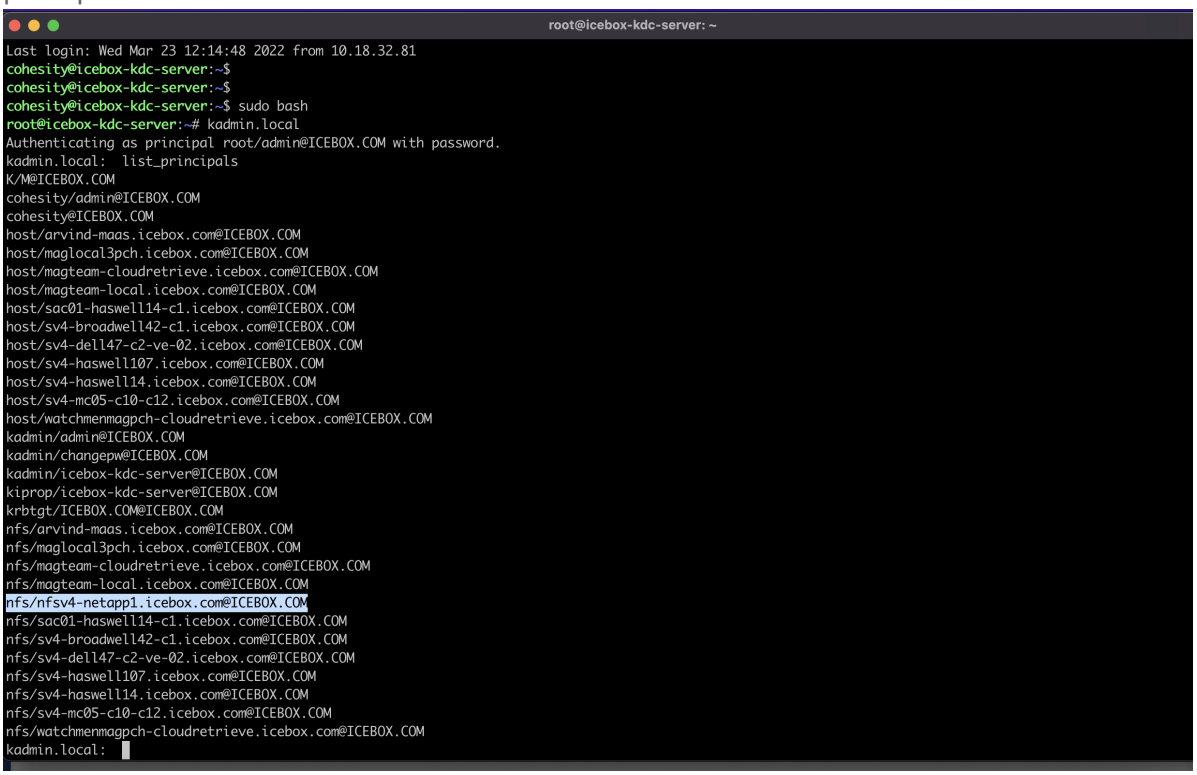

- In NetApp, navigate to **Storage > Storage VMs**.
- Click the required Storage VM. In the **Kerberos to UNIX** section, click **Add**.

- Enter the service principal patterns and map them to *root*.

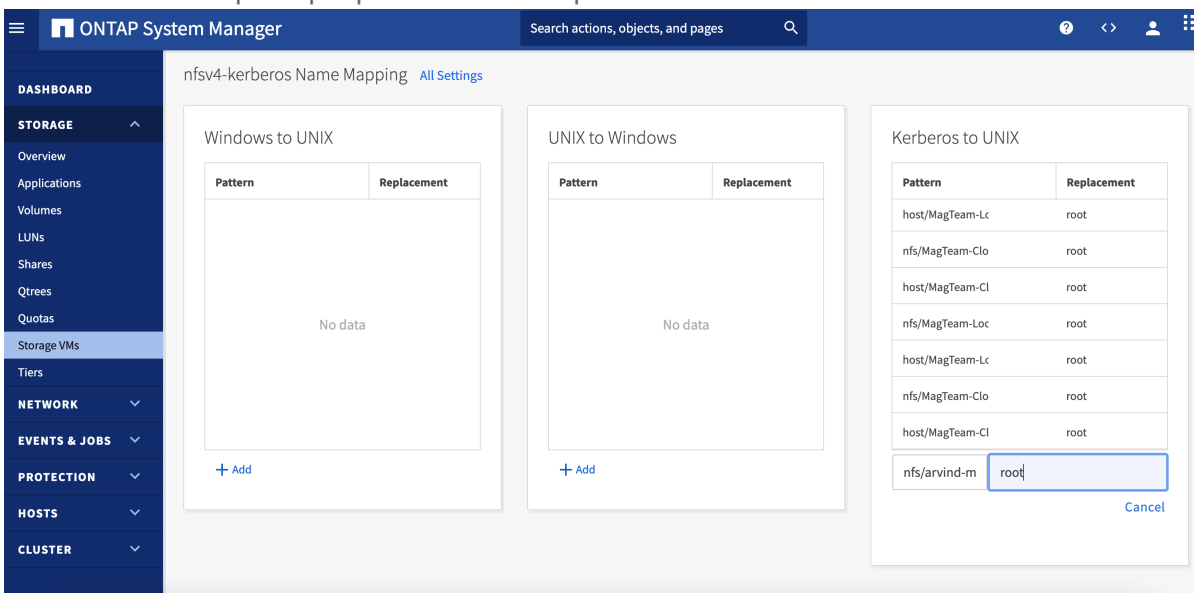

The Kerberos to UNIX mapping will be added on NetApp.

Supported NAS External Targets for Kerberos Authentication

Cohesity supports all NAS volumes of NFSv4 and above versions for Kerberos authentication.

Kerberos Authentication Configuration Steps

Before you register an NFS volume as an external target with the Cohesity cluster using Kerberos authentication, you must perform the following steps:

1. Add Administrators to the Kerberos Database.
2. Add Administrators to the ACLs File.
3. Add Kerberos Provider to Cohesity cluster.

Add Administrators to the Kerberos Database

On the Kerberos database that you maintain, you must add an administrative principal (admin principal) with the required ACLs. You can also use an existing principal that has the required ACLs. This admin principal will be used by the Cohesity cluster during Kerberos provider registration.

> **Note:** The steps for adding admin principal to Kerberos may vary depending on the Kerberos environment.

The following is an example of adding admin principal on a standard Linux environment:

You can use kadmin.local on the Kerberos admin server to create the administrative principal `cohesity/admin`

```
shell% /usr/local/sbin/kadmin.local
kadmin.local: addprinc cohesity/admin@ENG.COHESITY.COM
//ENG.COHESITY.COM is the Realm name//
NOTICE: no policy specified for "cohesity/admin@ENG.COHESITY.COM";
assigning "default".
Enter password for principal cohesity/admin@ENG.COHESITY.COM:  <= Enter a
password.
Re-enter password for principal admin/admin@ENG.COHESITY.COM:  <= re-enter
the password.
Principal "cohesity/admin@ENG.COHESITY.COM" created.
kadmin.local:
```

Add Administrators to the ACLs File

The admin principal you created on the Kerberos database must be added to the Access Control List (ACL) file.

The principal used during Kerberos provider registration must have the following minimum permissions:

- All permissions except extract password for NFS and host principals
- All permissions including extract password for Cohesity principal

> **Note:** The steps for adding admin principal to the ACL file may vary depending on the Kerberos environment.

The following is an example of adding an admin principal to the ACL file on the Linux server:

1. Edit the access control list for Kerberos administration.

   ```
   $ sudo cat /etc/krb5kdc/kadm5.acl
   ```

2. Add the following:

   ```
   cohesity/admin@ENG.COHESITY.COM xe cohesity@ENG.COHESITY.COM
   cohesity/admin@ENG.COHESITY.COM x nfs/*@ENG.COHESITY.COM
   cohesity/admin@ENG.COHESITY.COM x host/*@ENG.COHESITY.COM
   ```

   Here,

   - `cohesity/admin@ENG.COHESITY.COM` is the admin principal that the Cohesity cluster uses to make changes to the Kerberos configuration.
   - `cohesity@ENG.COHESITY.COM`, `nfs/*@ENG.COHESITY.COM`, `host/*@ENG.COHESITY.COM` are the user principals.

- The user principal, `cohesity@ENG.COHESITY.COM` is shared among multiple clusters. The admin principal (`cohesity/admin@ENG.COHESITY.COM`) has read/write privileges (`x`) with respect to this user principal. The admin principal also has privileges to extract password (`e`) for this user principal. The reason to grant the admin principal the permission to extract the password for this user principal is that this user principal is shared among multiple Cohesity clusters in the same Kerberos realm.

  For example, in a multiple clusters environment, the first Cohesity cluster that registers with the Kerberos realm will create a user principal with a random password. This password is saved in its keytab. However, when the second cluster registers with the Kerberos realm, it does not know the password for the user principal that the first cluster created. Therefore, this password needs to be extracted and saved in the second cluster's keytab.
- `nfs/*@ENG.COHESITY.COM`, `host/*@ENG.COHESITY.COM` principals are not shared among clusters. These principles have more restrictive privileges but have wildcard associated with it. The names (`nfs/*` and `host/*`) of these principals are derived from the host alias provided during Kerberos provider registration with the Cohesity cluster and are assumed to be unique for each Cohesity cluster.

Add Kerberos Provider to Cohesity cluster

You need to add the Kerberos provider to the Cohesity cluster to authenticate the Cohesity user to access NAS volumes in a Kerberos environment. The Cohesity user is used for mounting the NFS volume onto the Cohesity cluster.

Prerequisites

- Realm name

  > **Note:** Realm is the domain over which a Kerberos provider can authenticate a user, host, or service. Realm name is the unique identifier of the Realm.

- KDC server name and credentials
- Admin server

To add Kerberos Provider to the Cohesity cluster, complete the following steps:

1. Log in to the Cohesity Dashboard, navigate to **Settings** > **Access Management** > **Kerberos**.
2. In the Kerberos tab, click **Add Kerberos Provider**.
3. In the Add Kerberos Provider screen, complete the following steps:

1. Type **Realm Name**.
2. Type **Host Alias**.

   Host alias is Cohesity cluster names. Default host alias is your Cohesity cluster name. You can add more than one host alias name. The host alias name will be associated with the admin principals that Cohesity creates in the Kerberos provider.
3. (Optional) Toggle the **Use existing Host Aliases**.
4. Type **KDC Servers** names.

   You can add multiple KDC Server hostnames. Use a comma to separate the names.
5. Add **Admin Server** (KDC).

   You can select one of the added KDC servers as an admin server.
6. Type **Admin Principal** you created on the Kerberos database.
7. Type **Admin password**.
8. Click **Add**.

Once the Kerberos provider is added to the Cohesity cluster,

- The Cohesity cluster creates certain principals with the Kerberos provider depending on the provided host aliases:
  - `host/<host_alias>.<realm_name_in_lowercase>@<realm_name>`
  - `nfs/<host_alias>.<realm_name_in_lowercase>@<realm_name>`
- If you have provided multiple host aliases, multiple principals will be created.
- Cohesity cluster additionally creates a Cohesity "user" principal, `cohesity@<realm name>`, which is used by the NFS client when mounting a Kerberos-secured NFS volume.
- The authentication of the principals created is restricted to AES-128/256 encryption types. These encryption types apply to the encryption applied to tickets and Kerberos client/server messages.
- The principals are created with a randomly generated password using the `"-randkey"` `kadmin` option.
- On each Cohesity node, a keytab is created containing the randomly generated passwords of each of the created principals.
  - This is a standard way to store Kerberos client authentication details.
  - Access to the keytab files is protected from non-Cohesity users by access control.

- The keytab contents are stored in the Cohesity cluster configuration and shared among nodes.

Register NAS Volume with Cohesity cluster using Kerberos Authentication

Prerequisites

Before you register an external target with Kerberos authentication, ensure the following:

- You must enable Kerberos security for the NFS volume you register.
- The volume must be accessible through NFS 4.0, 4.1, or 4.2.
- You must provide principals created by Cohesity permission to access the Kerberos-secured NFS volume.

> **Note:** The steps for configuring the prerequisites may be different for each NAS source.

To register an external target for Cloud Archive and CloudArchive Direct, or CloudTier using Kerberos authentication, perform the following steps:

1. In the Cohesity Dashboard, select **Infrastructure** > **External Target**.
2. Click **Add External Target** located at the top right of the page.

   The Register External Target screen appears.
3. Select the **Storage Type** as **NAS**.
4. Select the **NFS version** as **4.0**, **4.1**, or **4.2.**
5. Select one of the following as the **Security Type**:
   - **KRB5**: Provides Kerberos authentication.
   - **KRB5I**: Provides Kerberos authentication and integrity checking.
   - **KRB5P**: Provides Kerberos authentication, integrity checking (KRB5I), and privacy (encryption).
6. Type **Kerberos Realm Name**.
7. In the **NAS Host** field, type the hostname or IP address of the NAS host.
8. In the **Mount Path** field, type the mount path of the NAS host.

   > **Note:** Only NFS is currently supported.
9. Click **Register**.

The NAS volume is now registered with the Cohesity cluster using Kerberos authentication.

QStar Tape Configuration

You can archive backups to tape for long term data retention. Cohesity uses QStar Archive Storage Manager (ASM) from QStar Technologies as the media manager for your existing tape libraries. Cohesity provides the QStar software, license and technical support, but all media management is handled by the QStar software.

For supported tape libraries, requirements, installation and configuration information, see

Archive Data to Tape with QStar & Cohesity

Unregister an External Target

Before you unregister an external target registered for archival, you must make sure that this external target is not used by any of the policies.

To unregister an external target registered for archival from the Cohesity cluster:

1. In the Cohesity Dashboard, navigate to **Infrastructure > External Target**.
2. Click the **Delete** icon.

   If the external target is used by any of the policies, then a prompt is displayed to remove the external target from the policies before unregistering.
3. In the **Unregister External Target** prompt, type **Unregister** for confirmation and optionally, toggle on **Delete Data**.

   By enabling the **Delete Data** option, all the data available at the target will be deleted before unregistering.

   If this is not enabled, then *only* the metadata will be deleted.
4. Click **Unregister**.

   Depending on the number of archives in that external target, the status moves to *Unregistering* along with the percentage displayed. On success, the external target will not be displayed on the External Targets page.

> **Note:** From the Cohesity UI, you cannot unregister an external target registered for CloudTier.

If you want to remove the external target registered for CloudTier from the cloud source, then you must:

1. Ensure to have enough storage space on the Cohesity cluster to uptier the data from the external target to the Cohesity cluster.
2. Gracefully remove the cloud disk.
3. Contact Cohesity Support to remove the external target from the Cohesity cluster.

Chapter 25: Disaster Recovery

Cohesity provides disaster recovery through replication and CloudRetrieve. To learn more, see the following topics:

About Replication 4373

Replication and Remote Access Setup 4391

About CloudRetrieve 4403

About Failover Ready 4408

Disaster Recovery Using CloudEdition 4409

About Replication

Organizations can achieve enterprise-level resiliency with site-to-site replication between Cohesity clusters. Replication can be configured at multiple granularity levels for maximum flexibility, including job-level or Cohesity View level (similar to a storage volume). All data on a single Cohesity cluster can be replicated to one or more clusters through the use of Protection Groups. The Protection Group specifies the objects (like databases, physical servers, VMs, Cohesity Views etc) to be backed up and replicated. Full and partial fail-over functions are fully automated using the Cohesity Orchestration workflow engine for simple, seamless recovery.

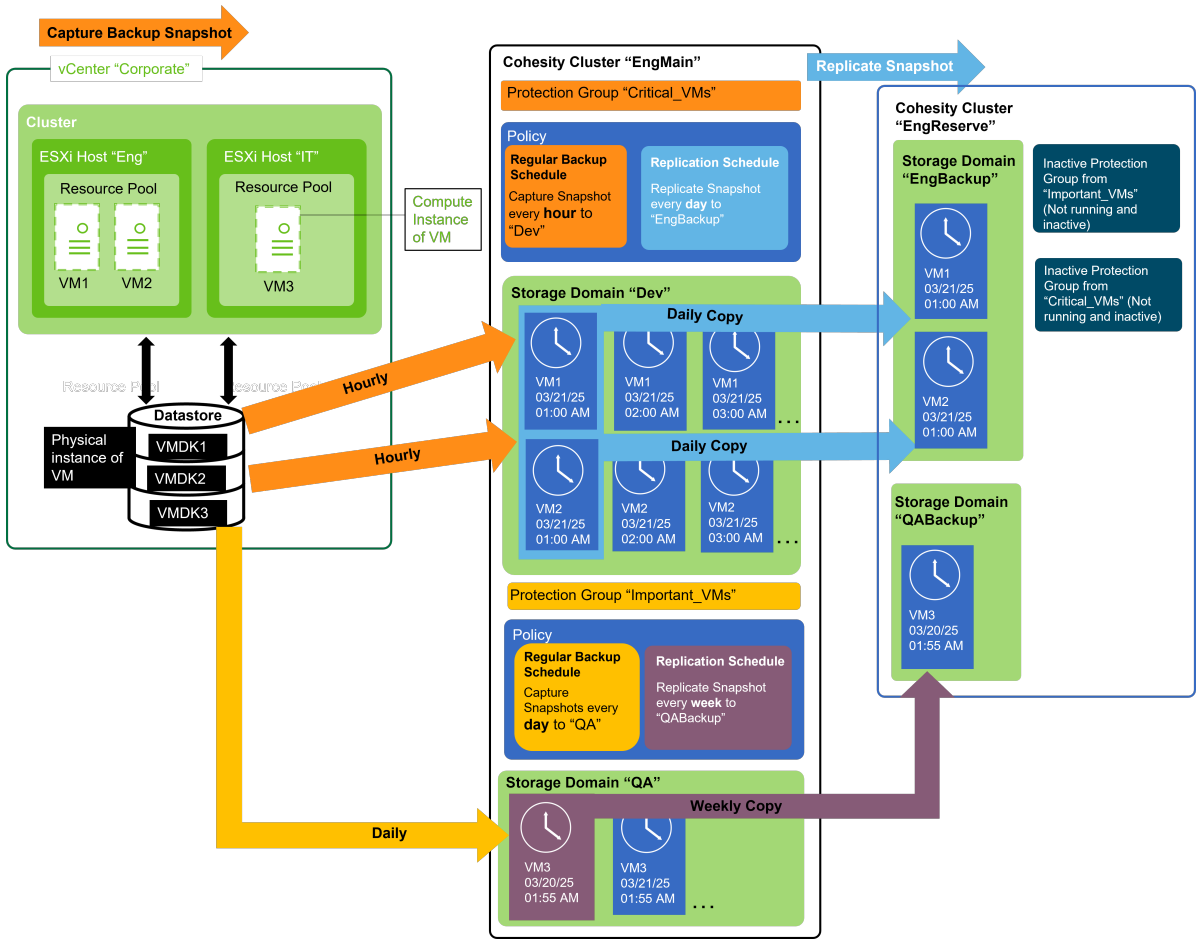

You specify a protection schedule in a Policy and that protection schedule defines how often a Protection Group using that Policy runs. In the example above, a Protection Group using a policy with an hourly protection schedule is capturing Snapshots of VM1 and VM2 and storing those Snapshots in the `Dev` Storage Domain. Another Protection Group using a policy with a daily protection schedule is capturing Snapshots of VM3 and storing those Snapshots in the `QA` Storage Domain.

You can optionally specify a replication schedule in a Policy. The replication schedule defines how often Snapshots are copied from a Storage Domain in the capturing Cohesity cluster and saved on a Storage Domain in the remote Cohesity cluster by a Protection Group using this Policy. In the example above, the Protection Group is making daily copies of the Snapshots in the `Dev` Storage Domain of the `EngMain` Cohesity cluster and storing those replicated Snapshots in the `EngBackup` Storage Domain of the remote Cohesity cluster called `EngReserve`. Another Protection Group with a different Policy is making weekly copies of the Snapshots in the `QA` Storage Domain of the `EngMain` Cohesity cluster and storing those replicated Snapshots in the `QABackup` Storage Domain of the remote Cohesity cluster called `EngReserve`.

The replication schedule for copying Snapshots cannot be more often than the protection schedule for capturing Snapshots. For example if Snapshots are captured hourly, the

Cohesity cluster cannot replicate them every 30 minutes but it can replicate them hourly or any value over one hour.

> **Note:** When replicating large files, if the replication process is slow for more than a day, contact Cohesity Support. For reference, see the KB Article.

> **Note:** After Snapshots of Objects (such as VMs) are created by a Protection Group with replication schedules, copy tasks are immediately created to replicate the Snapshots to the remote Cluster. The Cohesity cluster does not wait until the Snapshots of all the Objects protected by the Protection Group are captured, before it starts replicating.

Using the Replicated VM Snapshots

After VM Snapshots are replicated (copied) to a remote Cohesity cluster you can use the Snapshots to complete the following tasks:

- Recover VMs from the replicated Snapshots to a new location.
- Clone VMs from the replicated Snapshots.
- Recover files from the replicated Snapshots.
- Run the Analytics Workbench on these replicated Snapshots.
- Download VMX files from the replicated Snapshots.

The procedures for initiating these tasks are the same as initiating these tasks from Snapshots created by a Backup task of a Protection Group.

You can also use the Snapshots and Inactive Protection Group that were replicated to the remote Cohesity cluster for disaster recovery. To learn more, see Active Protection Group Replicating to Remote Cohesity cluster.

Using the Replicated View Snapshots

After View Snapshots are replicated (copied) to a remote Cohesity cluster you can use the snapshots to complete the following task:

- Clone a View From a Snapshot.

Using the Replicated SQL Server Snapshots

After SQL Server snapshots are replicated (copied) to a remote Cohesity cluster you can use the snapshots to complete the following task:

- Recover Microsoft SQL Server.

> **Warning:** For replicated MS SQL Server VM Snapshots, you can only restore MS SQL Server to a new location. (You cannot restore MS SQL Server to its original location.)

Object-based Replication

Object-based replication offers granular control by allowing replication at the individual object level. Instead of replicating all the runs in a protection group, you can replicate only the selected objects from a backup run to the target.

Follow these steps to perform object-based replication:

1. In the Cohesity cluster, select **Data Protection** > **Protection**.
2. On the **Protection** page, select the required Protection Group.
3. On the backup run details page, click the Actions menu (⋮) for the run and click **Edit Run**.
4. In the **Edit Run** dialog, click **Replication**.
5. From the **Replicate** dropdown, select the subset of objects in the backup run and click **Continue**.
6. From the **Replication** target dropdown, select the required target.
7. Toggle on **Modify Retention** if required.
    1. **Retention**: From the dropdown, select one of the following:
        1. **Extend**
        2. **Reduce**
    2. **Retain for**: Specify how many days, weeks, months, or years backups are retained on the cluster before they are deleted.
8. Click the save icon.

To add a new target to the backup run, click the plus icon.

About Replication Setup

Replication automatically makes copies of snapshots captured by a Protection Group on one Cohesity cluster and copies them to a second Cohesity cluster.

To learn more about Replication, see About Replication. Before the Cohesity cluster can replicate Snapshots, you must register the Cohesity clusters with each other. For more information, see Replication and Remote Access Setup.

Register Cohesity clusters for Replication

You must register the Cohesity clusters with each other and create at least one Storage Domain pairing between the Cohesity clusters. You can optionally specify a replication schedule in a Policy.

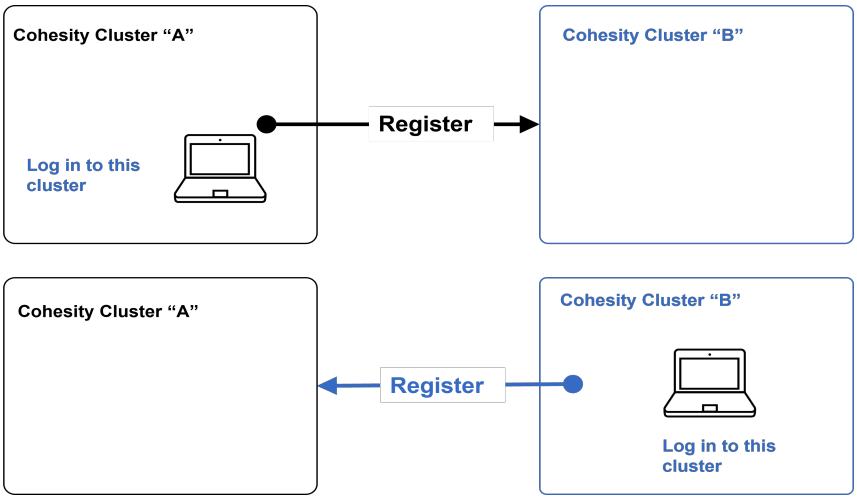

For more information on Replication Compatibility, see Replication Compatibility.

Pair Storage Domains for Replication

In addition, you must pair a Storage Domain on the capturing and replicating Cohesity cluster with a Storage Domain on the remote Cohesity cluster, as shown in the following figure. You cannot pair multiple Storage Domains on one Cohesity cluster to the same Storage Domain on the other Cohesity cluster. This is also shown in the following figure.

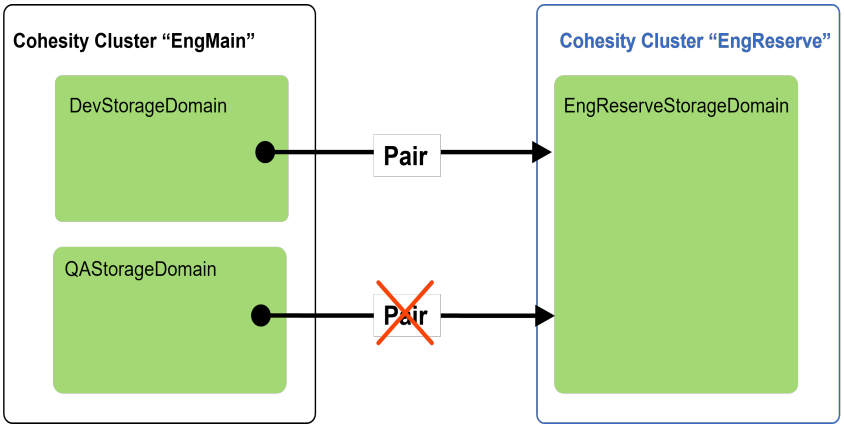

You can, however, pair multiple Storage Domains to a single Storage Domain if the source Storage Domains are on different Cohesity clusters as shown in the following figure.

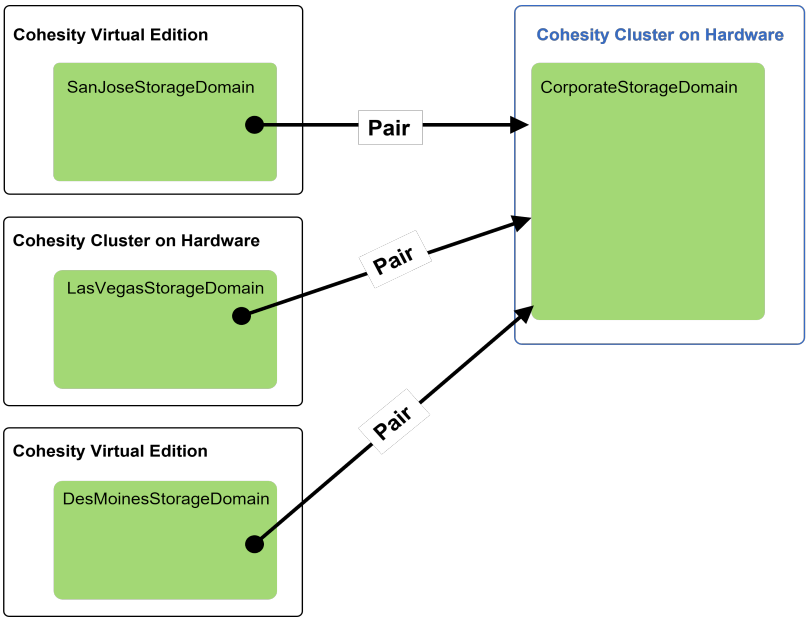

Cascaded Replication and Archival policy

Cohesity now supports Cascaded Replication. You can use this feature to not only replicate your Protection Groups from the source to the target cluster but also from the target cluster to another cluster. Also, protection objects can be archived from any of these targets.

Cohesity supports cascaded replication through scripts only. You have to create a cascaded replication policy that is non-cyclic. For example, replicating data from A to B to C is supported but Cohesity doesn't support the replication of data from A to B to C to A. For more information, see Cascaded Replication, Archival, and Vaulting Policy.

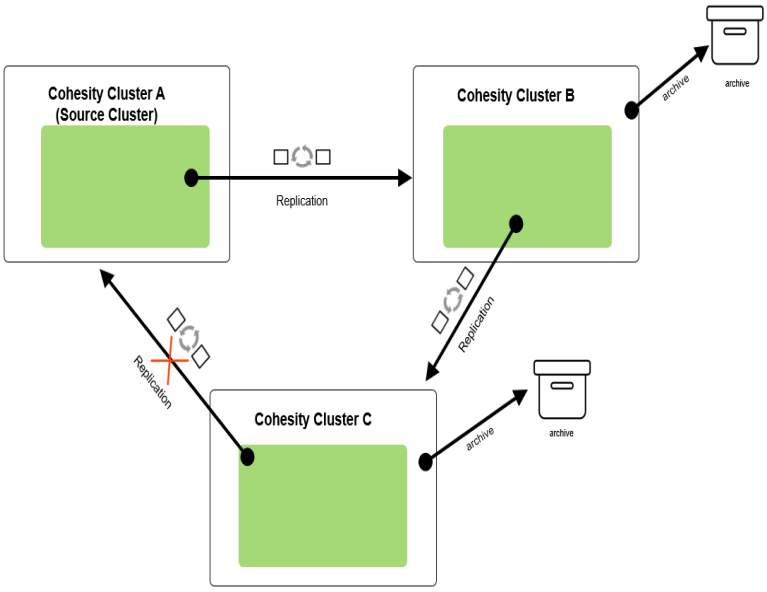

About Replication and Disaster Recovery

When a Protection Group with a replication schedule starts replicating snapshots from a primary Cohesity cluster to a remote Cohesity cluster, a Protection Group in Failover Ready state is created on the remote Cohesity cluster. If the original capturing Cohesity cluster and/or the original Source are down because there is a planned maintenance outage or a disaster destroys them, you can failover the Protection Group in Failover Ready state and deactivate the original task on the capturing Cohesity cluster. As part of the Failover, you are prompted to optionally recover the VMs from the replicated snapshots stored on the remote Cohesity cluster.

> **Warning:** When you create a Protection Group on a Cohesity cluster with a replication schedule after the first snapshot is replicated, a Failover Ready copy of the Protection Group is created on the remote Cohesity cluster that is receiving the snapshot. For more information on Failover Ready, see About Failover Ready. You can failover the Protection Group of VMs. You cannot failover the Protection Group of Views or SQL Servers. You can, however, clone replicated Views and restore SQL Servers. For more information, see Failover a Protection Group.

The disaster recovery procedure consists of the following actions:

1. Active Protection Group Replicating to Remote Cohesity cluster. Protection Group on a Cohesity cluster is actively capturing Snapshots of VMs and replicating them to a remote Cohesity cluster. After the first Snapshot is replicated, an Inactive copy of the Job is also created on the remote Cohesity cluster.
2. Failover from Replicated Copies. Disaster strikes and you must recover from the replicated copies on the remote Cohesity cluster that received the Snapshots. This Failover activates the Job on the receiving (remote) Cohesity cluster. However, the Job is not running.
3. Deactivate Original VM Protection Group. If possible, deactivate the original Job on the Cohesity cluster that originally captured the Snapshots.
4. Restore Activated Protection Group. You can restore the state of the activated Job to running, by assigning a new Policy to the Job and selecting Objects (VMs) to back up.

Active Protection Group Replicating to Remote Cohesity cluster

A Protection Group on a Cohesity cluster is actively capturing Snapshots and replicating them to other remote Cohesity clusters. The Protection Group specifies Objects (such as VMs) to be backed up, replicated or archived. The Policy used by the Protection Group determines the protection schedule, the replication schedule and the archive schedule. The protection schedule determines when Snapshots of the Objects are captured. The replication schedule determines when the Snapshots are replicated to other remote Cohesity clusters. The backup Snapshots of the specified Objects are captured and saved on

a Storage Domain in the capturing Cohesity cluster as represented by the "Back Up" orange arrow in the figure below. The specified Snapshots are replicated from the Storage Domain in the capturing Cohesity cluster ("EngMain" below) to the Storage Domain in the remote Cohesity cluster ("EngReserve" below) as represented by the "Replicate" purple arrow in the figure below. After the first Snapshot is replicated, an Inactive copy of the Protection Group is created on the remote Cohesity cluster.

The example below shows the deactivation of a Protection Group protecting VMs. Deactivating a Protection Group protecting SQL Servers is similar but not all details are the same. For example, the transaction logs are not described or pictured.

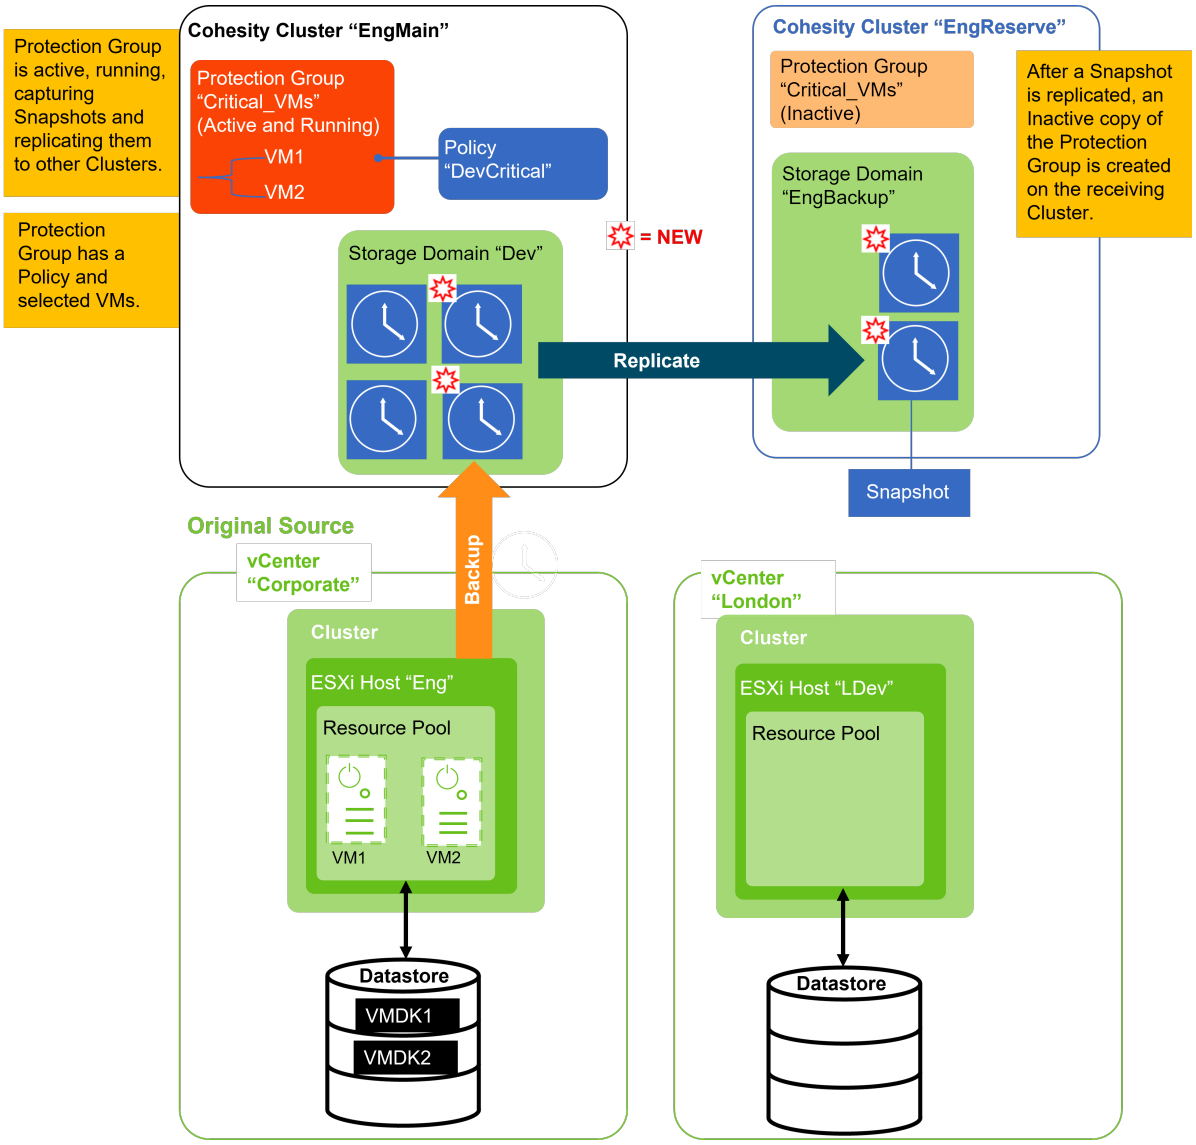

Next learn about Failover from Replicated Copies.

Failover from Replicated Copies

A disaster strikes the Source or there is a planned maintenance outage of the Source. The Cohesity cluster or the Source may be down or not available. From the Cohesity Dashboard, start the Failover procedure. For detailed instructions, see Failover a Protection Group. The example below shows the Failover of a task protecting VMs. Failover of a task protecting SQL Servers has transaction logs that are not pictured.

> **Note:** When you create a Protection Group on a Cohesity cluster with a replication schedule after the first snapshot is replicated, a Failover Ready copy of the Protection Group is created on the remote Cohesity cluster that is receiving the snapshot. For more information on Failover Ready, see About Failover Ready. You can failover the Protection Group of VMs. You cannot failover the Protection Group of Views or SQL Servers. You can, however, clone replicated Views and restore SQL Servers. For more information, see Failover a Protection Group.

During Failover, the following actions occur:

1. **Disaster at Source** - On the primary Cohesity cluster, the backup job fails to capture the snapshots of the source.
2. **Replication Fails** - As the source is down or unreachable, the replication job fails to replicate the snapshots onto the remote Cohesity cluster. The Failover Ready for VMs on the remote Cohesity cluster is activated and causes the rejection of the incoming replicated snapshots created by the associated job on the primary Cohesity cluster.
3. **Failover VMs** - You are prompted to optionally recover the VMs from the replicated snapshots on the remote Cohesity cluster to a new source.

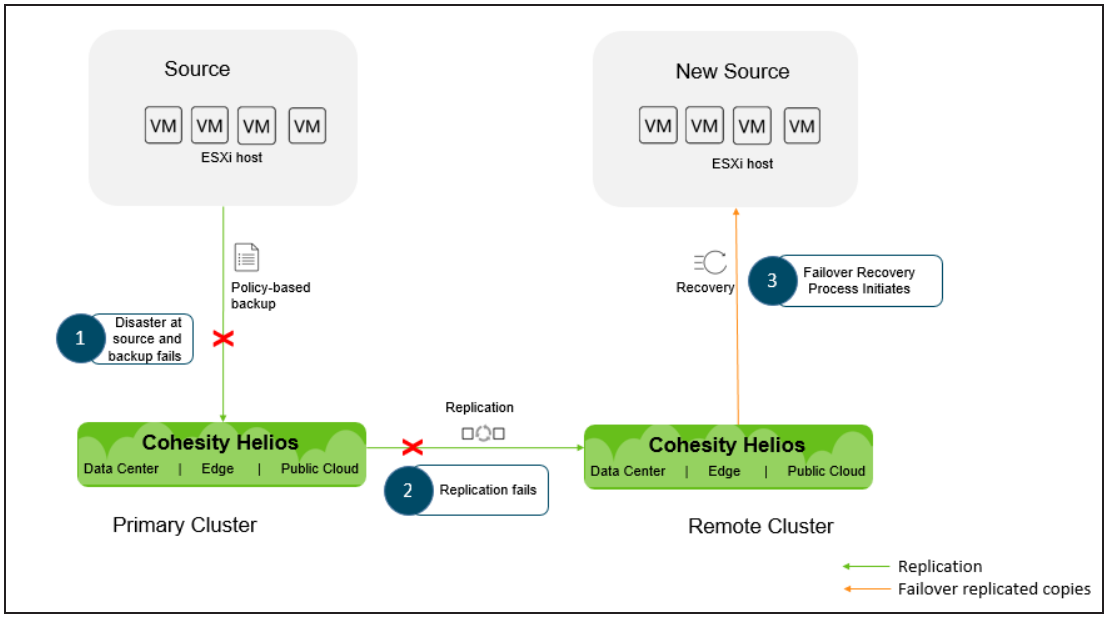

Deactivate Original VM Protection Group

Only Protection Groups that are protecting VMs can be deactivated. Protection Groups that are protecting a view cannot be deactivated.

> **Note:** If possible, deactivate the original Protection Group on the original capturing Cohesity cluster. If the original Cohesity cluster that captured the Snapshots is down or unreachable, you may be unable to deactivate the Protection Group.

The example below shows the deactivation of a Protection Group protecting VMs. Deactivating a Protection Group protecting SQL Servers is similar but not all details are the same. For example, the transaction logs are not described or pictured.

During deactivation, the following actions occur:

1. Snapshots are no longer captured by the original Protection Group on the original capturing Cohesity cluster. However, the existing snapshots captured by the Protection Group remain in the storage domain. They are not deleted.
2. Snapshots created by the Protection Group are no longer replicated to the receiving (remote) Cohesity cluster. (If you have already initiated the failover procedure, the replication has already stopped.)
3. You can optionally choose that the VMs on the original source be powered off. The VMs and the VM files, however, are not deleted. (The Cohesity cluster may be unable to power off the VMs if the original source is down or unreachable.)

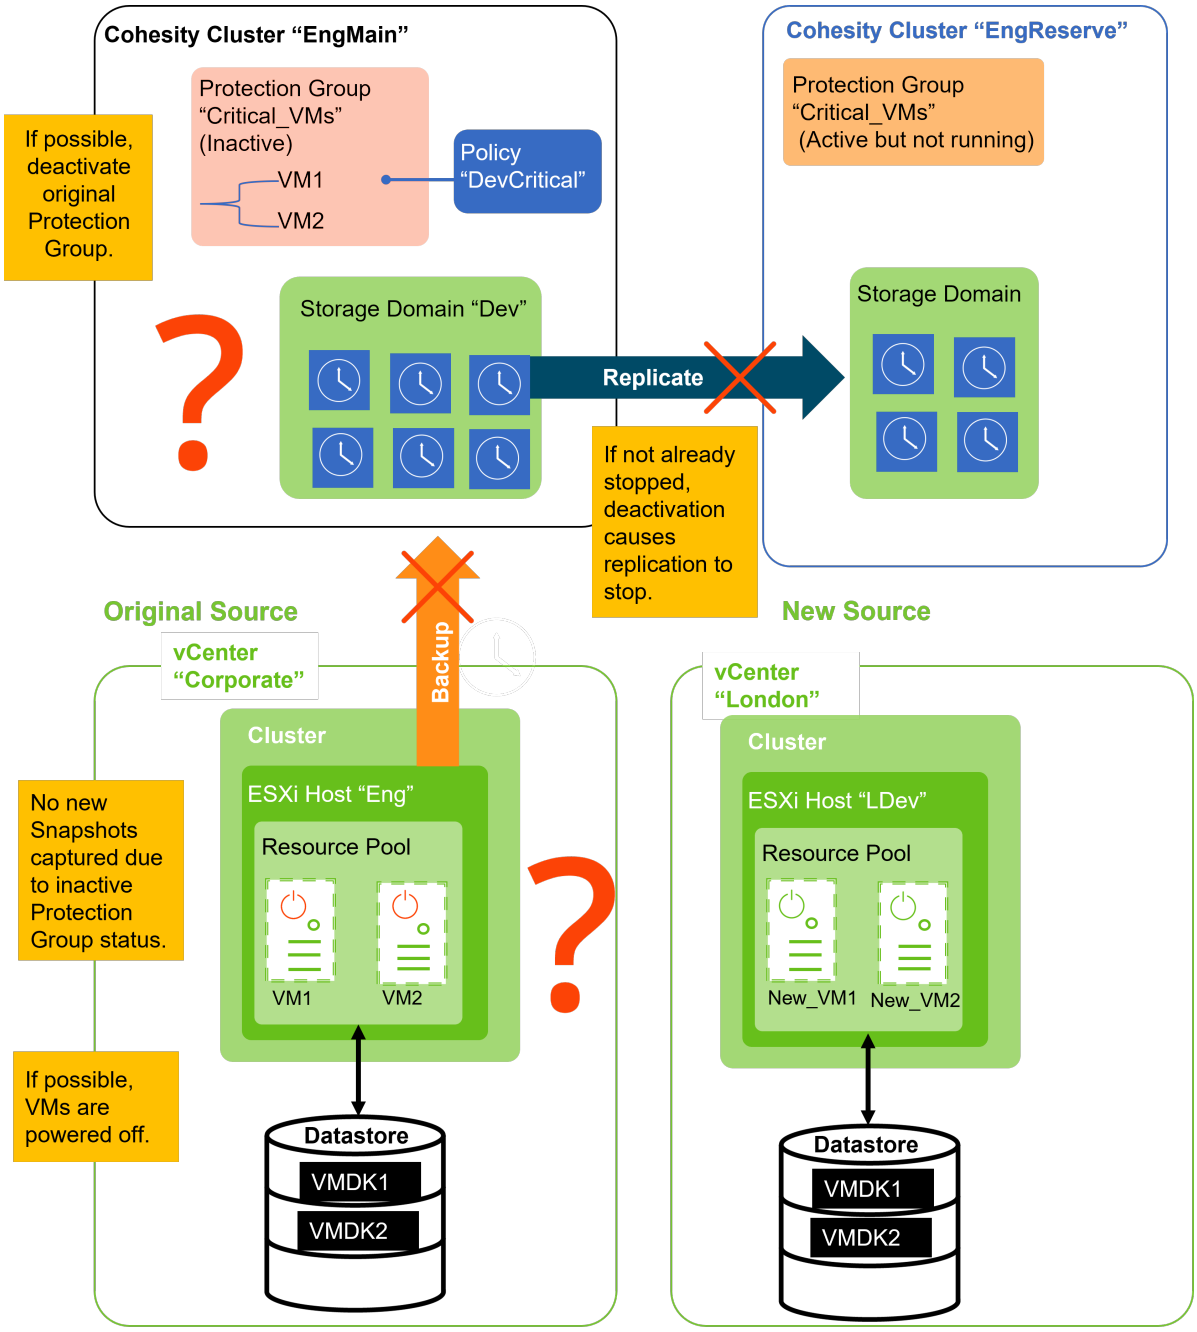

To deactivate the original Protection Group:

1. Log in to the Cohesity cluster where the original Protection Group was running.
2. Navigate to **Data Protection > Protection**.
3. Find the original Protection Group in the list, click on the actions menu ⋮ and select **Deactivate**.
4. Click **Deactivate and Power off VMs**, or click **Deactivate Only** if you do not want

to power off the VMs.

5. Optionally delete the VMs and VM files using the VMware vSphere client.

Restore Activated Protection Group

You can restore the activated Protection Group to a running state by assigning a new policy to the Protection Group and selecting Objects (VMs) to back up. The original receiving (remote) Cohesity cluster becomes the new capturing Cohesity cluster. If the new policy defines a replication schedule, the new capturing Cohesity cluster starts replicating the snapshots to a new receiving (remote) cluster as shown in the following figure. If the original Source is still up, you can also optionally clean up the old VM files and VMs. For detailed instructions, see Restore VM Protection Group to Running.

The example below shows restoring an activated Protection Group protecting VMs. Restoring an activated Protection Group protecting SQL Servers is similar but not all details are the same, for example the transaction logs are not described or pictured.

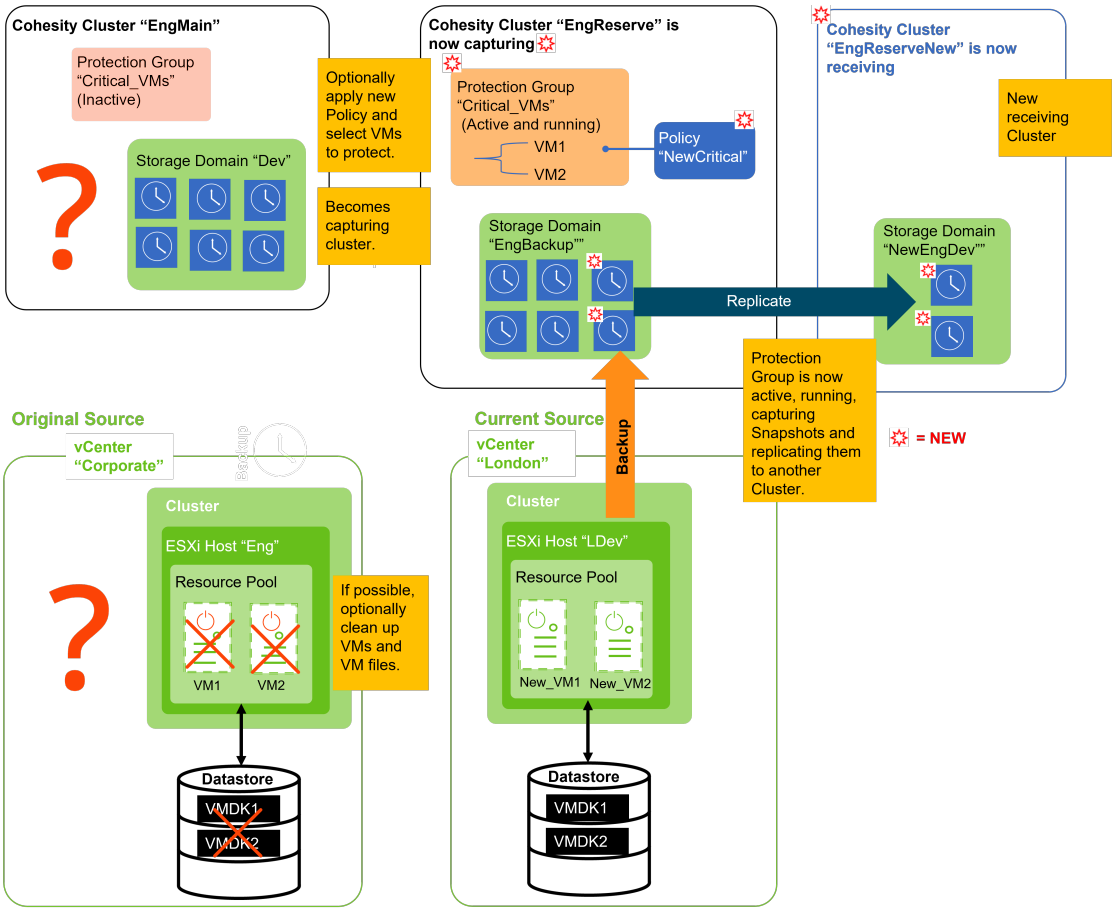

Cascaded Replication, Archival, and Vaulting Policy

By default, Cohesity supports site-to-site replication of the Protection Groups, which includes objects such as databases, physical servers, VMs, and Cohesity Views. Cohesity now supports Cascaded Replication. You can use this feature to not only replicate your Protection Groups from the source to the target cluster, but also from the target cluster to another cluster. Also, protection objects can also be archived from any of these targets.

Prerequisites

Make sure the following prerequisites are met before creating a cascaded replication policy:

- Install Python on the Virtual machine on the cluster with the following modules:
  - `argparse`
  - `ast`
  - `copy`
  - `requests`
  - `logging`
  - `kson`
  - `sys`
  - `os`
  - `csv`
  - `time`
  - `re`
  - `datetime`
  - `subprocess`
  - `urllib3`

    > **Note:** Python 3.10 must have all the modules except for `requests`.

    Install the unavailable modules using the following command:

    ```
    pip install <modulename>
    ```

    Example pip command:

    ```
    pip install request
    ```

    For more information, see How to install pip on windows.
- Storage Domains of the source cluster and the target cluster must be paired.
- The Cohesity cluster version must be 6.8.1 and above.
- Download the **manage_cascaded_policy.py** and sample inputs file here.

Before you Begin

Cohesity provides the following files to create a cascaded replication policy:

- **manage_cascaded_policy.py**: The `manage_cascaded_policy.py` file contains the code and is used as the source file to create a cascaded replication policy. For more information, see Create a Cascaded Replication policy.

  > **Warning:** Do not edit the `manage_cascaded_policy.py` file.

- **user_friendly_policy.json**: The `user_friendly_policy.json` file contains the policy name, IP address, replication topology and other information required for the cascaded replication and archival.

  Topology is the connection between two nodes in a network and it represents the replication happening from the source cluster to the target cluster.

  In topology, where replication must be done from cluster A > B > C. While providing information for cluster A, mention the immediate cluster B and not C.

> **Note:** If you are creating a new `user_friendly_policy.json` file, pass the file path as `--user_friendly_json_file_path=<PATH>` while executing the `manage_cascaded_policy.py` script. By default, the `user_friendly_policy.json` file is in the same directory as the `manage_cascaded_policy.py` file.

Use the following format to create or modify the `user_friendly_policy.json` file.

Here is an example of the `user_friendly_policy.json` to create a cascaded policy named 'sample_policy_name' and the topology is A > B > C for replication where C is 'Sample_target_2' for archival.

The IP address should be in the following format:

<xxx.xxx.xxx.xxx>

- **log_in_credentials.json**: The `log_in_credentials.json` file contains the credentials of the IPs provided in the `user_friendly_policy.json` file. Provide the login credentials of the source and target clusters in the format given in the .json file.

Note:

- If you are creating a new `log_in_credentials.json` file, pass the file path as `--log_in_credentials_file_path=<PATH>` while executing the `manage_cascaded_policy.py` script . By default, the `log_in_credentials.json` file is in the same directory as the `manage_cascaded_policy.py` file.
- The login credentials are required only while creating the cascaded policy. The credential file can be deleted once the policy is created.

Supported Parameters

Parameters	Default	Description
`-m, --mode`		Mode of script execution. (generate-only-iris-json/ create-policy-from-user-friendly-json/ create-policy-from-iris-json)
`-t/ --target`		Specify the source cluster to create the policy. This is required while executing the script in the `create-policy-from-iris-json` mode.
`-u/ --username`	'admin'	Specify the generic username of the Cohesity cluster. You can also provide unique usernames in the `log-in_credentials.json` file.
`-p/ --password`	'Cohe$1ty'	Specify the generic password of the Cohesity cluster. You can also provide unique usernames in the `log-in_credentials.json` file.
`-bkp_skd_unit/ --backup_run_schedule_unit`	'Hours'	Period for scheduling protection runs [Hours, Days, Weeks, Months].
`-bkp_skd_freq/ --backup_run_schedule_frequency`	1	Frequency to perform the backup.

Parameters	Default	Description
`-bkp_ret_unit/ --backup_retention_unit`	'days'	Retention period for the protection Runs [Days, Weeks, Months, Years].
`-bkp_ret_freq/ --backup_retention_frequency`	28	Timeframe to retain the protection Runs.
`-rep_ret_unit/ --replication_retention_unit`	'days'	Retention period for the replication runs [Days, Weeks, Months, Years].
`-rep_ret_freq/ --replication_retention_frequency`	14	Timeframe to retain the replication runs.
`-arc_ret_unit/ --archival_retention_unit`	'days'	Retention period for the archival runs [Days, Weeks, Months, Years]
`-arc_ret_freq/ --archival_retention_frequency`	14	Timeframe to retain the archival runs.
`-path/ --iris_json_path`	`cascaded_protection_policy.json`	Specify the path of the cascaded policy creation request payload json file. This is required while executing the script in `create-policy-from-iris-json` mode.
`-auth_file_path/ --log_in_credentials_file_path`	`log_in_credentials.json`	Specify the path of the login credentials json file. The users are expected to provide the credentials in the `log-in_credentials.json` file, or pass the appropriate file path with the same format as the `log-in_credentials.json` file.

Parameters	Default	Description
`-user_json/ --user_friendly_json_file_path`	`user_friendly_policy.json`	Specify the path of the user-friendly json file. The users are expected to make the required changes in the `user_friendly_policy.json` or pass the appropriate file path with the same format as the `user_friendly_policy.json` file.
`-output_path/ --cascaded_policy_path`	`cascaded_protection_policy.json`	Specify the path of the generated `cascaded_protection_policy json` file. By default, the generated json file is saved as `cascaded_protection_policy.json`.

Manage Cascaded Replication Policy

You can create and edit a cascaded replication and archival policy through script only.

Create a Cascaded Replication policy

You can create a cascaded replication policy using any of the following methods.

Method 1

Generate a policy json file that can be used by the Cohesity Platform to create a cascaded protection policy.

1. Open the terminal (python must be accessible from this terminal) on the Virtual machine or the cluster .
2. Run the following command to create the `cascaded_protection_policy.json` file and click **Enter**:

   ```
   python ./manage_cascaded_policy.py --mode=generate-only-iris-json
   ```

   If the `user_friendly_policy.json` file and the `log_in_credentials.json` file are in different folders or the name of the `log_in_credentials.json` file is different from the default name, add the path of the `user_friendly_policy.json` and the `log_in_credentials.json` file along with the python command.

   ```
   --user_friendly_json_file_path=<json file path> --log_in_credentials_file_path=<log in file path>
   ```

   If you want the `cascaded_protection_policy.json` file to be created in a different folder, add the following command along with the python command:

```
--cascaded_policy_path=<path_where_you_want_to_save_cascaded_policy>
```

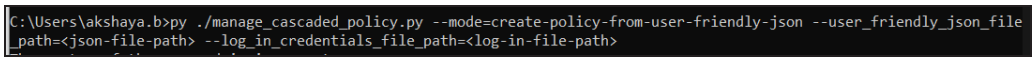

If an error '**Can't open file <path of the manage_cascaded_policy.py file>: No such file or directory**', is displayed, run the following command and click **Enter**.

```
cd <path of the manage_cascaded_policy.py file>
```

On successful creation of the `cascaded_protection_policy.json`, the following message appears.

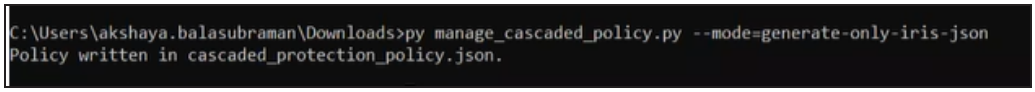

3. To create a policy in the Cohesity cluster using the `cascaded_protection_policy.json` file, run the following command and click **Enter**:

   ```
   python ./manage_cascaded_policy.py --mode=create-policy-from-iris-json
   --target=<source-cluster-ip>
   ```

   where *iris json* is the created cascaded replication policy.

   If the json file containing the cascaded replication policy is in a different folder, add the following along with the above python command.

   ```
   --iris_json_path=<path-to- iris-json>
   ```

Here is an example of the generated `cascaded_protection_policy.json` file.

The Cascaded replication policy is now created in the Cohesity Platform and you can view the policy by navigating to **Data Protection** > **Policies**.

Method 2

1. Open the terminal on the Virtual machine or the cluster and to create a policy using the `user_friendly_policy.json` file, run the following command and click **Enter**:

   ```
   python ./manage_cascaded_policy.py --mode=create-policy-from-user-
   friendly-json
   ```

   If the `user_friendly_policy.json` file and the `log_in_credentials.json` file are in different folders, run the following commands along with the command above.

   ```
   --user_friendly_json_file_path=<json file path> --log_in_credentials_
   file_path=<log in file path>
   ```

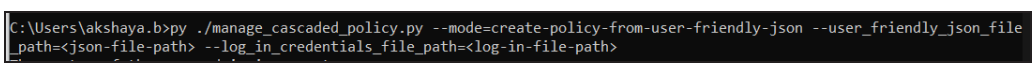

If an error '**Can't open file <path of the manage_cascaded_policy.py file>: No such file or directory**', is displayed, run the following command and click **Enter**.

```
cd <path of the manage_cascaded_policy.py file>
```

On successful creation of the `cascaded_protection_policy.json`, the following message appears.

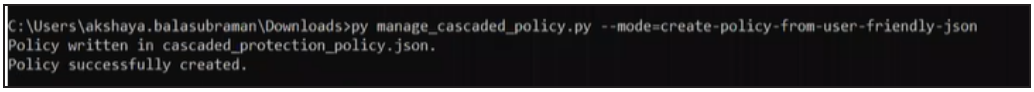

The json file is created in the directory where the `manage_cascaded_policy.py` file resides and the policy is created in the Cohesity Platform.

2. To view the created policy, log in to the Cohesity cluster, navigate to **Data Protection** > **Policies**.

> **Note:** You cannot edit the created cascaded policy from the Cohesity Platform.

Edit the Cascaded Replication Policy

To edit the cascaded replication policy, you require the cascaded replication policy ID.

Use the following URL format to fetch the created cascaded policy id:

```
https://<cluster IP address>/siren/v1/remote?remoteUrl=http%3A%2F%2F<IP address>%3A20000%2Fpolicyz
```

Considerations

- Cohesity supports cascaded replication through scripts only. Ensure to provide accurate data for the replication to be successful.
- Create a cascaded replication policy which is non-cyclic. For example, replicating data from A to B to C is supported but Cohesity doesn't support the replication of data from A to B to C to A.
- CloudRetrieve is not intended to be used to recover data to the source or target Cohesity clusters. CloudRetrieve can be used to retrieve data to an alternate location.

Replication and Remote Access Setup

On the Remote Cohesity clusters page, you can configure replication and remote access to another Cohesity cluster:

- Replication automatically makes copies of snapshots captured by a Protection Group on one Cohesity cluster and copies them to a second Cohesity cluster.You must register the Cohesity clusters with each other and create at least one Storage Domain

pairing between the Cohesity clusters. You can optionally specify a replication schedule in a Policy. A single Cohesity cluster can act as both a local and remote Cohesity cluster. For more information, see About Replication.

- Remote access lets you set up multi-cluster access and manage remote Cohesity clusters from the local Cohesity cluster. The Cohesity clusters must be the same version (6.0 or later). Once you add a remote Cohesity cluster to the local Cohesity cluster, you can switch to the remote Cohesity cluster by selecting its name from the drop-down list on the top left of the Cohesity Dashboard. The drop-down is not displayed until you add at least one remote Cohesity cluster. You can also display the consolidated statistics for all the remote access Cohesity clusters by selecting **All Clusters Dashboard** from the drop-down.

Note that replication and remote access are independent of each other; you can enable either or both for a remote Cohesity cluster. One does not require the other.

From the Remote Cohesity cluster's Dashboard page, you can begin setting up replication and remote access by clicking **Register Remote Cluster**. For detailed instructions, see Create a Connection to the Remote Cohesity Cluster.

To configure the Cohesity cluster's 1G interface for remote replication, see Set Up the Secondary Interface for Replication.

Note the following about replication:

- For information about using replication between Cohesity clusters running different versions, see Replication Compatibility.

  For the latest notes about replication compatibility, prerequisites and limitations, see the *Release Notes*.
- If you have a Protection Group that is capturing and replicating Snapshots multiple times a day, Cohesity recommends configuring the replication schedule to copy Snapshots daily instead of replicating Snapshots after each protection run. If the replication schedule is too frequent, the replication may lag behind the capturing of Snapshots resulting in a backlog of replication tasks.
- If Snapshots of a VM are replicated to a remote Cluster and the VM is renamed in the vCenter Server, the Cohesity Dashboard on the remote Cluster displays the original VM name in the protection run Details page. However, you can search for new VM name while recovering or cloning and the search results display the new VM name. Replication is not affected by this issue.
- In production environments, Cohesity recommends not replicating from one single node Cohesity cluster Virtual Edition to another single node Cohesity cluster Virtual Edition. Cohesity recommends replicating from Cohesity cluster Virtual Editions to Cohesity clusters running directly on hardware.

Prerequisites:

- When registering a remote Cluster, log in to the local Cluster with a user in the default LOCAL domain that has an Admin Role such as the default `Admin` account. Do not use users or groups that were added to the Cohesity cluster for Active Directory principals to register a remote Cluster.
- Ensure that the required firewall ports (443, 11111, 20000, and 23335) are accessible for replication to a remote Cohesity cluster.

  > **Note:** If both the local and remote Cohesity clusters support TLS 1.3, port number 23335 is utilized for TLS communication during replication.

- Verify the IP reachability for both VIP and Node IPs between the Cohesity cluster capturing and replicating Snapshots and the Cohesity clusters receiving the replicated Snapshots.
- If the cluster replication is configured, verify that the network connectivity is functioning properly during the upgrade to ensure the cluster replication relationship is successfully upgraded to use AES-256-GCM for encryption.
- In a replication environment, if either the transmitting or receiving cluster is a multitenancy-enabled, both clusters must be running version 6.1 or later. Storage domain pairing must be set up on the source cluster. One or both clusters can be multitenancy-enabled.

Note the following about multi-cluster access:

- Both the local Cohesity cluster and the remote access Cohesity cluster must be running the same version (6.0 or later) of the Cohesity cluster software.
- On the local Cohesity cluster for which multi-Cluster access is configured, verify that firewall port 443 (TCP outbound) is open to all the remote access Cohesity clusters.
- If a Cohesity cluster is configured for multi-Cluster access and then the Cohesity cluster is re-imaged or re-installed, Cohesity recommends deleting and re-adding the remote access configurations to ensure that multi-Cluster access functions properly.

The columns on the Remote Cohesity clusters page are described below.

Column Name	Description / Action
Cluster Name	The name of the Cohesity cluster that will be accessed remotely and/or used to store replicated (copied) Snapshots. If the Cohesity cluster is auto-registered at the target, then the **Auto Registered** tag will be displayed next to the Cohesity cluster name. For more information, see Create a Connection to the Remote Cohesity Cluster.

Column Name	Description / Action
Remote Access	Whether remote access has been configured for the Cohesity cluster.
Replicate	Whether replication has been configured for the Cohesity cluster. • **Yes** indicates replication is enabled and at least one Storage Domain pairing between the Cohesity clusters is configured. • **No** indicates replication might be enabled, but there are no Storage Domain pairings between the Cohesity clusters.
Actions	To manage a remote Cohesity cluster, select one of the following options: • **Edit** ( )–Change the connection settings for the remote Cohesity cluster. Follow the instructions provided in Create a Connection to the Remote Cohesity Cluster. • **Delete** ( )–Delete the connection between the local and remote Cohesity cluster. When you delete a connection: ◦ If replication was configured, all Storage Domains Pairings in this connection are deleted. In addition, any Replication Tasks created by Protection Groups that are using this connection will fail and Snapshots will not be copied to the Storage Domains on the remote Cohesity cluster. ◦ If remote access was configured, the remote Cohesity cluster cannot be managed from the local Cohesity cluster.

Create a Connection to the Remote Cohesity Cluster

If you are setting up replication on the Cohesity cluster that contains the original snapshots to replicate, you must register the connection to the remote Cohesity cluster that will store the replicated snapshots. The capturing (local) Cohesity cluster captures the snapshots and replicates them to the remote Cohesity cluster.

> **Note:** A remote cluster is reachable by the source cluster only after the registration is complete because the firewall ports on the remote cluster are open after the registration.

To register a connection to a remote Cohesity cluster for replication and remote access:

1. Log in to the Cohesity cluster on which the original snapshots are stored.
2. Navigate to **Infrastructure > Remote Clusters**.

3. On the **Remote Clusters** page, click **Register Remote Cluster**.

   The **Register Remote Cluster** dialog is displayed.

4. **VIP or Node IP Addresses**—Specify one of the following:

   - One or more Virtual IP (VIP) addresses of the remote Cohesity cluster. Cohesity recommends specifying VIPs because of their ability to failover to other nodes in the Cohesity cluster if a node is down or unreachable.
   - One or more node IP addresses of the remote Cohesity cluster.

     If **All Cluster Node IPs are Reachable** is enabled for replication, a single node is specified in this field and the primary Cohesity cluster can establish an initial connection to the specified remote node, and the Cohesity cluster distributes the network traffic to other nodes in the remote Cohesity cluster. Cohesity recommends specifying VIPs. However, if you specify node IP addresses, Cohesity recommends specifying multiple node IP addresses. By specifying multiple node IP addresses, if the network traffic cannot reach a node, the Cohesity cluster can try a different node to establish the initial connection.

     If network traffic is restricted in your environment, specify the VIP or node IP addresses of the Cohesity cluster that are configured to receive network traffic in your environment and turn off the **All Cluster Node IPs are Reachable** toggle.

     > **Note:** To add more IP addresses after the remote connection is configured, edit the remote connection, click **Change**, enter the addresses in the **VIP or Node IP Addresses** field, and click **Connect** to verify connectivity to the addresses.

5. **Username**—Specify a local or Active Directory user.

   Authentication by Active Directory user is an Early Access feature. Contact your Cohesity account team to enable the feature.

   - For remote access, Cohesity recommends you specify an admin user. Ensure you disable the MFA for this user.
   - For replication, Cohesity recommends you specify a local user with only the **Replication** role on the remote Cohesity cluster. Ensure you disable the MFA for this user. For more information see Disable MFA for Local Users.

> **Note:**
> - The local user with only the **Replication** role on the remote Cohesity cluster does not have access to the Cohesity UI and can only perform replication.
> - The audit log records the user specified while setting up the remote Cohesity cluster and not the user logged in to the remote Cohesity cluster. All the actions done on a remote Cohesity cluster are done on behalf of the remote user configured while setting up the remote Cohesity cluster.

6. **Password**—Specify the password of the user specified.

   > **Note:** The password must adhere to the ISO-8859-1 character set; otherwise, the cluster registration will fail.

7. **Domain**—Specify the domain of the user who can authenticate the remote cluster registration. The default domain is LOCAL. For AD users, you can select the respective domain.

   > **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

8. Select **Interface Group**—Select one of the following:
   - Select **Automatically** if you want the Cohesity cluster to select the interface group.
   - Select **Manually** to select the associated interface group from the dropdown list.
9. Click **Continue** to specify the cluster options.

   The cluster connection status is displayed.
10. Turn on the **Remote Access** toggle to manage the remote Cohesity cluster from the Cohesity cluster you are currently using. Remote access is not required for replication. After remote access is enabled, you can select the remote Cohesity cluster name from the dropdown in the top left of the Cohesity cluster.
11. **Replication**—Starting from the 6.8 release, replication is enabled by default. Configure replication settings to the remote Cohesity cluster.
12. **Pair Storage Domains for Replication**—Select the storage domains:

a. Select a local storage domain. This storage domain contains the original snapshots that will be replicated.

> **Note:** You cannot pair multiple storage domains on one Cohesity cluster to the same storage domain on the other Cohesity cluster. You cannot pair a storage domain to two different storage domains on the same remote cluster. You can pair a storage domain to two different storage domains if they are in different remote clusters. For more information, see Pair Storage Domains for Replication.

b. Select a remote storage domain. This storage domain is where the copies of the replicated snapshots are stored.

13. Outbound compression is enabled by default. Cohesity recommends keeping this option enabled so that this Cohesity cluster can compress the replication data before sending it to another Cohesity cluster.
14. **All Cluster Node IPs are Reachable**—Turn on the toggle if traffic can be distributed to all nodes. Turn off if your network environment requires restricting network traffic to the subset of IP addresses listed in the **Node IP Address** field. For example, if your environment has a firewall that is limiting the network traffic to the remote Cohesity cluster and if **All Cluster Node IPs are Reachable** is enabled, the Cohesity cluster automatically distributes the network traffic to other nodes in the remote Cohesity cluster.

> **Note:** If there is Network Address Translation (NAT) between two clusters, ensure that the **All Cluster Node IPs are Reachable** option is disabled. This prevents communication issues due to translated IP addresses that may not be directly routable across the clusters.

15. Turn on the **Encryption** toggle to encrypt the replication data generated by this Cohesity cluster. Click **Generate New Key**. The encryption key is displayed. The Cohesity cluster that is receiving encrypted data must enable encryption and specify the same encryption key as the Cohesity cluster that sent the replication data. Copy the encryption key to the clipboard, so you can paste it while creating the connection from the capturing Cohesity cluster back to the remote Cohesity cluster as described in Create Connection Back to the Capturing Cohesity cluster.
16. Turn on the **Throttle** toggle to set a replication data transfer rate limit from this Cohesity cluster to another Cohesity cluster. Specify the data rate in Megabits per second. For this connection, this value limits the data transfer rate when copying replication data such as snapshots from this Cohesity cluster to another Cohesity

cluster.

17. Turn on the **Quiet Times and Throttle Overrides** toggle to specify one or more replication override windows. These windows override the setting in the **Data Transfer Rate Limit** field during the time period specified for the window.
    - You can optionally change the time zone that applies to all the override windows specified for the connection between the local Cohesity cluster and the remote Cohesity cluster.
    - Select the days of the week when the override window must be applied.
    - Specify a time period in a day when the override window must be applied.
    - Specify a type for the window:
        - **Quiet Time**—A window when data is not replicated from the local Cohesity cluster to a remote cluster. This value overrides the **Data Transfer Rate Limit** value specified.
        - **Throttle**—A window when this data transfer rate overrides the **Data Transfer Rate Limit** value specified. This window sets the maximum data transfer rate between the local Cohesity cluster and the remote Cohesity cluster during the specified time period.
    - If you selected **Throttle**, enter the replication maximum data transfer rate for this window in Megabits per second. For this connection, this value limits the rate that replication data such as snapshots are copied from the local Cohesity cluster to the remote Cohesity cluster during this window.

    > **Note:** If overlapping windows are specified, the latest one that was created takes precedence. For example, if you first create a quiet time window from Monday 9 AM to 10 PM and then a throttle window from 5 PM to 12 PM with a data rate of 100 megabits per second, from 5 PM to 12 PM the data rate is limited to 100 megabits per second.

18. Click **Continue** to view the summary.
19. Verify the details and click **Register** to register the remote Cohesity cluster.

> **Note:** If you configured replication, ensure that you set up replication on the remote Cohesity cluster. For more information, see Create Connection Back to the Capturing Cohesity cluster.

Edit Remote Cohesity Cluster Details

To edit a remote Cohesity cluster connection:

1. Navigate to **Infrastructure > Remote Clusters**.
2. On the **Remote Clusters** page, hover over the remote cluster that you want to edit and click the edit icon.

   The **Edit Remote Cluster** dialog is displayed.
3. Modify the remote cluster details as needed and click **Save**.

Delete the Remote Cohesity Cluster Connection

To delete a remote Cohesity cluster connection:

1. Navigate to **Infrastructure > Remote Clusters**.
2. On the **Remote Clusters** page, hover over the remote cluster that you want to delete and click the delete icon.
3. Click **Delete** to confirm.

   After you delete, all storage domain pairings in this connection are deleted and any replication task using this connection fails.

Create Connection Back to the Capturing Cohesity cluster

On the remote Cohesity cluster that will store the replicated snapshots, you must register the Cohesity cluster that is capturing the original snapshots. The capturing (local) Cohesity cluster captures the snapshots and replicates them to the remote Cohesity cluster.

To register a connection back to the Cohesity cluster that is capturing the original snapshots:

1. Log in to the Cohesity cluster that will store the replicated snapshots.
2. Navigate to **Infrastructure > Remote Clusters**.
3. On the **Remote Clusters** page, click **Register Remote Cluster**.

   The **Register Remote Cluster** dialog is displayed.

   > **Important:** In this procedure, you are logged in to the remote Cohesity cluster and registering the capturing Cohesity cluster. Specify the IP address and user credentials of the local Cohesity cluster.
4. In the **VIP or Node IP Addresses** field, specify one of the following:
   - One or more Virtual IP (VIP) addresses of the local Cohesity cluster. Cohesity recommends specifying VIPs because of their ability to failover to other nodes in the Cohesity cluster if a node is down or unreachable.
   - One or more node IP addresses of the capturing Cohesity cluster. If **All Cluster Node IPs are Reachable** is enabled for replication, a single node is specified

in this field and this Cohesity cluster can establish an initial connection to the specified primary node, and the Cohesity cluster distributes the network traffic to other nodes in the capturing Cohesity cluster. Cohesity recommends specifying VIPs. However, if you specify node IP addresses, Cohesity recommends specifying multiple node IP addresses. By specifying multiple node IP addresses, if the network traffic cannot reach a node, the Cohesity cluster can try a different node to establish the initial connection.

If network traffic is restricted in your environment, specify the VIP or node IP addresses of the Cohesity cluster that are configured to receive network traffic in your environment and disable **All Cluster Node IPs are Reachable**.

> **Note:** To add more IP addresses after the remote connection is configured, edit the remote connection, click **Change**, enter the addresses in the **VIP or Node IP Addresses** field, and click **Connect** to verify connectivity to the addresses.

5. **Username**—Specify a local Cohesity user.
   - For remote access, Cohesity recommends you specify an admin user.
   - For replication, Cohesity recommends you specify a local user with only the Replication role on the remote Cohesity cluster. Ensure you disable the MFA for this user. For replication, Cohesity recommends you specify a local user with only the **Replication** role on the remote Cohesity cluster. Ensure you disable the MFA for this user. For more information, see Disable MFA for Local Users.

   > **Note:**
   > - The local user with only the **Replication** role on the remote Cohesity cluster will not have access to the Cohesity UI and can only perform replication.
   > - The audit log records the user specified while setting up the remote Cohesity cluster and not the user logged in to the remote Cohesity cluster. All the actions done on a remote Cohesity cluster are done on behalf of the remote user configured while setting up the remote Cohesity cluster.

6. **Password**—Specify the password of the user specified.

   > **Note:** The password must adhere to the ISO-8859-1 character set; otherwise, the cluster registration will fail.

7. Click **Continue** to specify the cluster options.

   The cluster connection status is displayed.

8. Turn on the **Remote Access** toggle to manage the capturing Cohesity cluster from the remote Cohesity cluster. Users will be able to select the remote Cohesity cluster name from the drop-down in the top left of the Cohesity Dashboard.
9. **Replication**—Starting from the 6.8 release, replication is enabled by default. Configure replication settings to the remote Cohesity cluster.
10. **Pair Storage Domains for Replication**—The storage domain paring that was configured on the capturing Cohesity cluster is displayed. You can add additional parings as needed, but only one pairing is required. To add more pairings:
    a. Select a local storage domain. This storage domain contains the original snapshots that will be replicated.

    > **Note:** You cannot pair multiple storage domains on one Cohesity cluster to the same storage domain on the other Cohesity cluster. You cannot pair a storage domain to two different storage domains on the same remote cluster. You can pair a storage domain to two different storage domains if they are in different remote clusters. For more information, see Pair Storage Domains for Replication.

    b. Select a remote storage domain. This storage domain is where the copies of the replicated snapshots are stored.
11. Outbound compression is enabled by default. Cohesity recommends keeping this option enabled so that this Cohesity cluster can compress the replication data before sending it to another Cohesity cluster.
12. **All Cluster Node IPs are Reachable**—Turn on the toggle if traffic can be distributed to all nodes. Turn off if your network environment requires restricting network traffic to the subset of IP addresses listed in the **Node IP Address** field. For example, if your environment has a firewall that is limiting the network traffic to the capturing Cohesity cluster and if **All Cluster Node IPs are Reachable** is enabled, the Cohesity cluster automatically distributes the network traffic to other nodes in the capturing Cohesity cluster.
13. **Encryption**—If encryption was enabled on the capturing Cohesity cluster that is sending replication data to this remote Cohesity cluster, complete the following steps:
    a. Turn on the **Encryption** toggle.
    b. Paste the encryption key copied from the capturing Cohesity cluster. The Cohesity cluster that is receiving encrypted data must specify the same encryption key as the capturing Cohesity cluster that sends the replication data.

> **Note:** If you do not have the encryption key in your clipboard, log in to the capturing Cohesity cluster, select **Infrastructure > Remote Clusters**, select **View Details** ( ) and click **Copy Key to Clipboard**.

14. **Throttle:** Turn off if this Cohesity cluster is only receiving replicated data. If this Cohesity cluster is also copying replicated data, turn on the **Throttle** toggle to set a replication data transfer rate limit from this cluster to another cluster. Specify a data rate in Megabits per second. For this connection, this value limits the data transfer rate when copying replication data (such as snapshots) from this Cohesity cluster to another Cohesity cluster.
15. **Quiet Times and Throttle Overrides**—Turn off the toggle if this Cohesity cluster is only receiving replicated data. Turn on the **Quiet Times and Throttle Overrides** toggle to specify one or more replication override windows. These windows override the setting in the **Data Transfer Rate Limit** field during the time period specified for the window.
    - You can optionally change the time zone that applies to all the override windows specified for the connection between the local Cohesity cluster and the remote Cohesity cluster.
    - Select the days of the week when the override window must be applied.
    - Specify a time period in a day when the override window must be applied.
    - Specify a type for the window:
        - **Quiet Time**—A window when data is not replicated from the local Cohesity cluster to a remote cluster. This value overrides the **Data Transfer Rate Limit** value specified.
        - **Throttle**—A window when this data transfer rate overrides the **Data Transfer Rate Limit** value specified. This window sets the maximum data transfer rate between the local Cohesity cluster and the remote Cohesity cluster during the specified time period.
    - If you selected **Throttle**, enter the replication maximum data transfer rate for this window in Megabits per second. For this connection, this value limits the rate that replication data such as snapshots are copied from the local Cohesity cluster to the remote Cohesity cluster during this window.

> **Note:** If overlapping windows are specified, the latest one that was created takes precedence. For example, if you first create a quiet time window from Monday 9 AM to 10 PM and then a throttle window from 5 PM to 12 PM with a data rate of 100 megabits per second, from 5 PM to 12 PM the data rate is limited to 100 megabits per second.

16. Click **Continue** to view the summary.
17. Verify the details and click **Register** to register the capturing Cohesity cluster.

You can now add a replication schedule to a policy. For more information, see Create or Edit a Standard Policy.

About CloudRetrieve

CloudRetrieve provides the ability to download data that was archived from a Cohesity cluster to an alternate (non-original) Cohesity cluster. For example: You have Cohesity cluster A, which archives data to a cloud external target, but you want to download that archived data to Cohesity cluster B.

> **Note:** CloudRetrieve is not intended to be used to recover data to the original Cohesity cluster. To restore data to the original Cohesity cluster, see About Recovering VMs.

Download to an Alternate Cohesity cluster

The following example depicts one source Cohesity cluster archiving data to a cloud external target followed by an alternate Cohesity cluster retrieving the archived data from that external target. The action steps in the following procedure are represented in the figure below with blue numbered time squares:

.

1. Register the external target that contains the data you want to download to the alternate Cohesity cluster.
2. Submit a search for Protection Groups that contain the data you want to download from the cloud external target. The search may take a few hours to complete. If the job to be downloaded is from a particular Cohesity cluster, specify the job name and Cohesity cluster name to narrow the results.
3. If the external target's encryption key file is managed manually, upload the key file to the alternate Cohesity cluster before triggering the search. See Create a CloudRetrieve Search Task for more details.

4. Download a snapshot from any of the Protection Groups in the search results. You can also build a local index of Protection Groups that ran during a specified date range. Downloading can take a few hours to complete. See Create a CloudRetrieve Download Task for more details.
5. Recover or clone any downloaded objects using the Cohesity recovery or clone workflows. See CloudRetrieve Downloaded Protection Groups for more details.

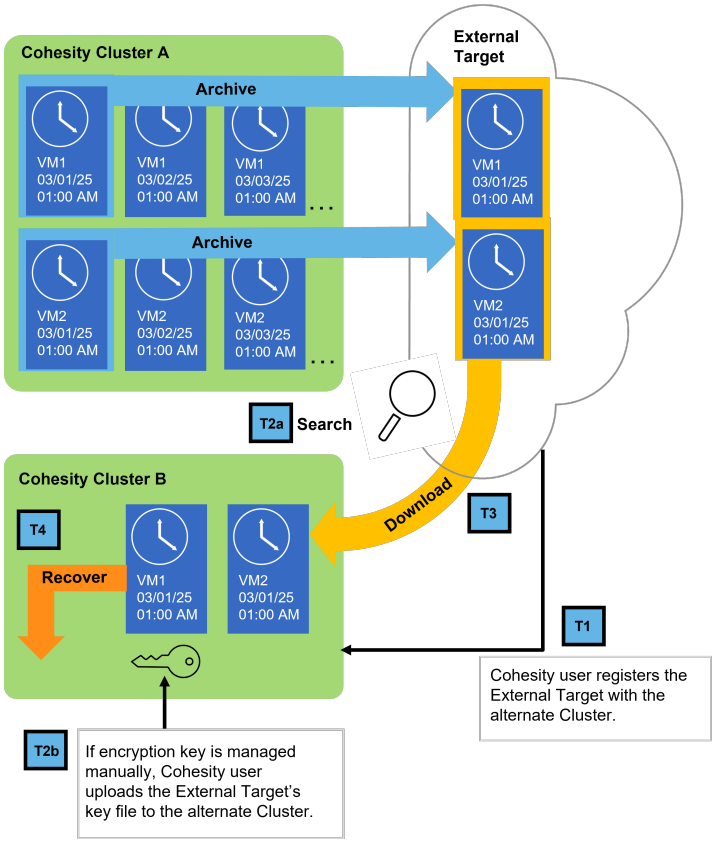

Considerations

- It is recommended that the Cloud Retrieve cluster runs a newer software version than the source cluster. The version of archives performed increases across releases and the older clusters cannot retrieve newer versions.
- Cohesity recommends not registering the same AWS S3 bucket, Azure storage container, or GCP bucket more than once for CloudRetrieval, even if the storage class differs.

CloudRetrieve Searches

The Searches tab lists search tasks and their status. Click a search task name to set up a download task.

Create a CloudRetrieve Search Task

> **Note:** CloudRetrieve is not intended to be used to recover data to the original Cohesity cluster.

In this section you create a task to search for Cohesity clusters and Protection Groups that were archived to the selected external target. To learn more, see About CloudRetrieve.

Before you can submit a search, an external target must be registered. See Register an External Target.

Search Cohesity clusters and Protection Groups

1. Select an external target. You may need to register an external target if the one you want to search is not registered with the Cohesity cluster.
2. In the **Task Name** field, you can optionally change the default name for the task.
3. Change the **Date Range** by clicking the date range link (for example: 06/30/2017 - 07/06/2017), and adjusting the time period. A larger date range results in a longer search time. The default date range is one year.
4. Optionally provide a **Cohesity Cluster Name** to search for. This is the Cohesity cluster where the Protection Groups originated.
5. Optionally provide a **Protection Group Name** to search for.

   > **Important:** Wildcard characters (like '*') are NOT supported.

   > **Note:**
   > - If you enter search terms for both **Cohesity Cluster Name** and **Protection Group Name**, your search must find matches for the Protection Group within clusters that match.
   > - If your search is too narrow, try entering a search term for **Cohesity Cluster Name** or **Protection Group Name**, or leave one or both empty.

6. If the external target is protected by a manually managed key, click **Attach** to upload the encryption key file to the Cohesity cluster.
7. Click **Search**.

   The search can take up to several hours to complete.

> **Note:** The success of a CloudRetrieve search does not guarantee that the search found any matches. It only means that the search operation was completed successfully. If your search results are empty, broaden your search with partial names for the cluster and/or Protection Group, leave them blank, and/or extend the date range.

8. Go to searches on the CloudRetrieve page and click the name of the search task to create a download task. See Create a CloudRetrieve Download Task for more details.

Create a CloudRetrieve Download Task

In this section, you can create a task to download data from an external target. To learn more, see About CloudRetrieve.

Download Data

1. If desired, filter search results by Protection Group type or Cohesity cluster. You can also narrow results by searching for a Protection Group or Cohesity cluster.
2. Select the Protection Groups to download. Selecting the check box adjacent to Protection Group selects all filtered jobs.
3. For each selected Protection Group, you can edit download options by clicking Edit ( ). Or click **Edit Selected** to set download options for all selected Jobs together.

   Select the **Download Job Meta-Data** check box to fetch and index Job meta-data for the specified range. After the Protection Groups are indexed, they are listed as inactive Protection Groups on the Manage Protection Groups page. Jobs protecting VMs can be activated using the failover action. For more information, see Failover a Protection Group.

   Select the **Download Snapshot** check box to download the specified snapshot to the Cohesity cluster. Selecting a specific Snapshot is not available when editing multiple Protection Groups together (you clicked **Edit Selected**); the latest snapshot is always downloaded for these Protection Groups.

   > **Note:** In addition to the selected archive, a download fetches job meta-data for all archives that the currently selected archive depends on. This results in all the dependent protection runs displaying on the Job Details page. Additionally, selecting only a snapshot for download also downloads job meta-data.

4. Select or create a storage domain to contain the downloaded objects.

> **Note:** Selected archived Protection Groups that are linked to Protection Groups on this Cohesity cluster use the local Protection Group's storage domain for their download location. Selected archived Protection Groups without a known local Protection Group are downloaded to the selected storage domain.

5. Click **Download**.

   Downloads may take up to several hours to complete.

   (*Applies only for cloud archive incremental with periodic full*) Provisioned Capacity Units (PCU) are required for all file-level recoveries. Provisioned capacity ensures that your retrieval capacity for expedited retrievals is available when needed. Each unit of capacity guarantees that at least three expedited retrievals can be performed every five minutes and provides up to 150 MB/s of retrieval throughput.

   Glacier Expedited retrieval offers a significantly improved latency (1-5 minutes) over Standard/Bulk retrievals (3-12 hours). As part of this, AWS may fail expedited Glacier requests with a lack-of-resources response:

Option	Description	HTTP Status Code	Client
uri	Returned if there is insufficient capacity to process the expedited request. This error only applies to expedited retrievals and not to standard or bulk retrievals.	503 Service Unavailable	Server

   The Cohesity Bridge warning logs will return the following error when the cloud provider cannot process the expedited request.

```
curl_http_rpc_executor.cc:226] Response for the failed RPC: 9173 is
{"code":"InsufficientCapacityException","message":"There is
insufficient capacity to process this Expedited
request.","type":"Server"}
```

6. Go to Downloaded Jobs on the CloudRetrieve page and choose an action to perform from the actions menu.

CloudRetrieve Downloaded Protection Groups

The Downloaded Job tab lists Protection Groups that were previously downloaded. Mouse over the actions menu and choose an action to perform.

About Failover Ready

The state of a Protection Group protecting VMs can be "Failover Ready" for any of the following reasons:

- When you create a Protection Group on a Cohesity cluster with a replication schedule after the first snapshot is replicated, a Failover Ready copy of the Protection Group is created on the remote Cohesity cluster that is receiving the snapshot.

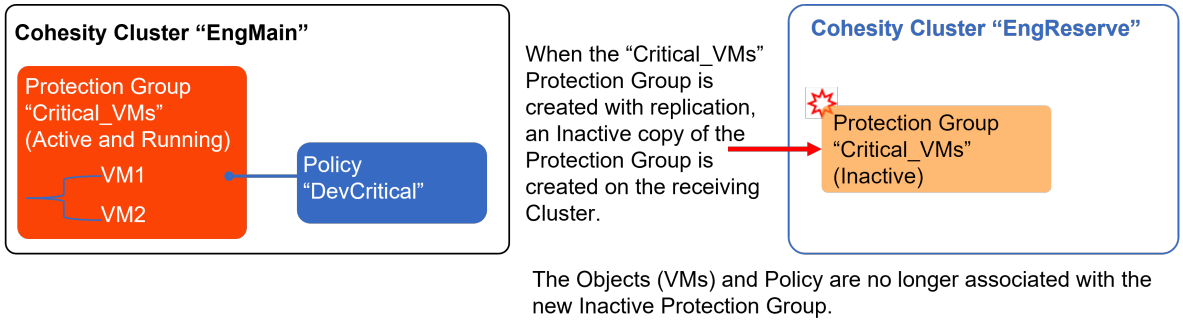

- When an active Protection Group is deactivated ( ), the Protection Group becomes Inactive and the selected VMs and Policy are still associated with the Protection Group.

> **Warning:** Cohesity recommends only deactivating a Protection Group as part of the failover process made possible by replication. Deactivate the original Protection Group running on the Cluster that captured the Snapshots before the failover and restoration of the Protection Group on the remote Cluster receiving the Snapshots. For more information, see Deactivate Original VM Protection Group and Failover from Replicated Copies.

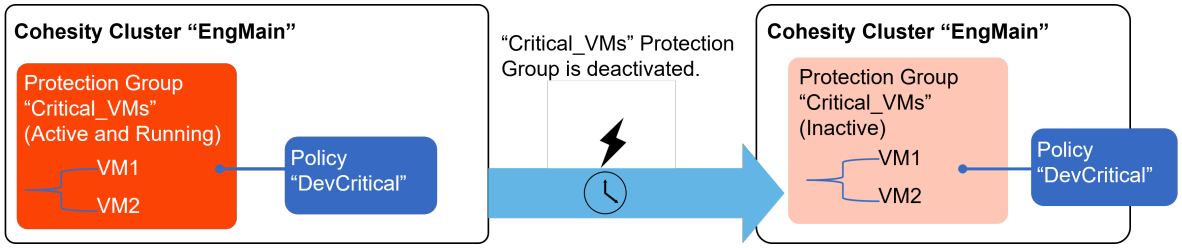

- When the Protection Group meta-data is downloaded during CloudRetrieve, meta-data for the specified range is fetched and indexed. After the Protection Groups in this range are indexed, they are listed as Inactive Protection Groups on the Manage

Protection Groups page.

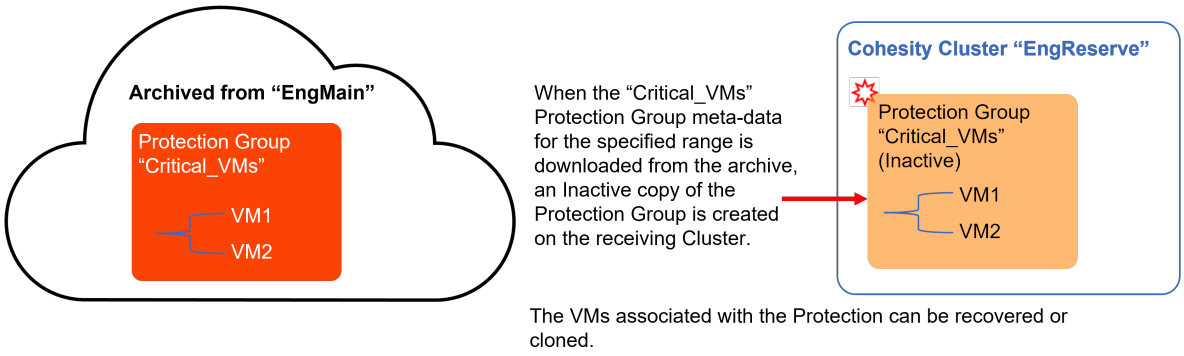

> **Note:** Only Protection Groups that are protecting VMs can be deactivated. Protection Groups that are protecting a view cannot be deactivated.

Disaster Recovery Using CloudEdition

You can use Cohesity's Cohesity Cloud Edition for disaster recovery. For more information, see Disaster Recovery of VMs Using AWS Cohesity Cloud Edition.

Chapter 26: Manage Sources

From the Sources page, you can start the following actions:

- Add (register) a new source that contains objects to protect
- Back up objects in a source
- Download agents to physical or virtual servers.

The Sources page also provides an overview of the sources that have already been registered to the Cohesity cluster. Display modes can show VMs in a hierarchy or a flat list of all VMs.

Sources are organized by type and you can search on the source name. To get details and protection statistics about a source, click a source link in the source column of the table.

To register a new source on the Cohesity cluster, the host settings checker checks the health of different types of hosts. Host registration is not affected by the failure of any or all checks, but it is recommended that users address all error messages according to the instructions below.

Note the following about sources:

- The host settings checker is currently supported only for MS SQL hosts.
- For MS SQL FCI clusters, health checks are run only on the primary node. If a failover causes the primary node to change, health checks will begin to run on the new primary and stop on the previous node.
- If a Virtual Distributed Switch is created *after* the vCenter Cohesity user was created, then ensure that the Cohesity user has privileges to access the vCenter datacenter and its children objects (ESXi hosts, clusters, distributed switches, etc.). Without these privileges, some children ESXi hosts and clusters will not display in the object protection status hierarchy.

Manage Existing Sources Using Actions Menu Options

For a multitenancy Cohesity cluster, you can only modify the objects that belong to your organization.

- **Protect** ( )—Create a Protection Group using this source.
- **Edit** ( )—Configure the settings of a registered source. For detailed instructions, see Register or Edit a Hypervisor Source.
- **Unregister** ( )—Unregister this source. The Cohesity Dashboard prevents you from unregistering a source that is currently being used by a Protection Group.
- **Refresh** ( )—Manually force a refresh of the object hierarchy. The Cohesity cluster tracks all the objects (for example VMs, resource pools, datacenters, folders, and ESXi hosts) found within a source in an object hierarchy. This object hierarchy is automatically refreshed every four hours. If a new object such as a VM is added to the source, the Cohesity Dashboard does not list the new VM until the object hierarchy is refreshed, either manually or automatically after four hours.
- **Register Applications** —For more information, see the Manage Registered Applications Running on Virtual Servers section.
- **Deploy Cohesity Agent**—For more information, see the Manage Cohesity Installed Agents Running on Virtual Servers section.
- **Maintenance**—For more information, see the Maintenance Mode section.

Maintenance Mode

With your Cohesity cluster, you can easily enable Maintenance mode for registered sources. Enabling Maintenance mode for your Cohesity cluster's registered source will skip any scheduled activities or manually initiated operations, such as protection runs, restores, and upgrades. Ongoing backups will be canceled, but ongoing restores will continue without any interruptions.

> **Note:** You can only apply the maintenance mode at the level at which you registered it with the Cohesity cluster. For example, if you registered a vCenter, the entire vCenter will be put into maintenance mode. You won't be able to place individual ESXi hosts or VMs within that vCenter exclusively into maintenance mode.

To place a registered source in maintenance mode:

1. Navigate to **Data Protection** > **Sources**.
2. From the Action Menu (⋮) of the source, select **Maintenance**.

3. You have the option to initiate the maintenance mode immediately or later. To initiate the maintenance mode immediately:
   a. Select **Start Now**. The **Start Maintenance Now** dialog displays the list of Protection Groups available in that source.
   b. Review the list of Protection Groups and enter yes to confirm.
   c. Optionally, you can enter a note.
   d. Click **Confirm**.

   To schedule the maintenance mode:
   a. Select **Schedule**.
   b. On the **Schedule Maintenance** dialog, enter the start date and time and the end date and time.
   c. Optionally, you can enter a note.
   d. Click **Save**.

Register Applications Running on Virtual Servers

If a server is running MS SQL Server and/or MS Exchange, you can register these applications with the Cohesity cluster.

Navigate the hierarchy to the server, mouse over its actions menu and click **Register Applications** (+). See Set Up MS SQL on VMs as a VMware vCenter application with the Installed Agent and Register Exchange on a VM.

Manage Registered Applications Running on Virtual Servers

- **Edit Application Registration** ( )—Configure the settings of a registered application.
- **Unregister All Applications** ( )—Unregister all applications on the server.

Manage Cohesity Installed Agents Running on Virtual Servers

To upgrade the Cohesity agents on VMware, see Upgrade the Agent on VMware.

To upgrade Cohesity agents on Hyper-V, see Upgrade the Agent on SCVMM and Hyper-V Hosts.

To upgrade Cohesity agents on RHV/OLVM, see Upgrade the Linux Agent.

Manage Existing Physical Servers Using the Actions Menu

For a multitenancy Cohesity cluster, you can only modify the objects that belong to your organization.

- **Edit Server** ( )—Configure the settings of a registered server. For detailed instructions, see Register or Edit a Physical Server.
- **Unregister Server** ( )—Unregister this server from the Cohesity cluster.
- **Refresh Server** ( )—Manually force a refresh of the object hierarchy.

  > **Note:** The refresh operation only collects information about the agent version, new objects, and deleted objects. No agent restart is initiated during this process

- **Register as MS SQL Server** ( )—If a physical server is running MS SQL, you can register it as MS SQL Server with the Cohesity cluster. Before registering, see Set Up MS SQL on Virtual or Physical Servers using the Agent or Set Up Microsoft SQL Server FCI.
- **Unregister MS SQL Server** ( )—If MS SQL Server is already registered with the Cohesity cluster, you can unregister it.
- **Upgrade Agent** ( )—Upgrade the Cohesity agent on the server if an upgrade is available. If an agent upgrade is requested while backup jobs or recover tasks are active, the upgrade starts after the jobs and tasks finish. This page also displays any agent upgrades that are in progress and if you must manually re-install an agent.

Health Check for SQL Sources

The health checks are run when a source is refreshed, automatically or manually. The current automatic source refresh interval is 8 hours.

Health Check Name	Error Message	Corrective Action
Cohesity Agent service **Description** Checks whether the cohesity agent service is running.	Error querying service status for Cohesity Agent Service.	1. Ensure that the Cohesity Agent is installed. 2. Provide local administrator rights to the CohesityAgent service.
	The service Cohesity Agent Service is not running. Current state: <current_state>	Ensure that the Cohesity Agent service is running. If yes, try restarting it.
SQL Server VSS Writer **Description** Checks whether the SQL Server VSS Writer service is running.	Error querying service status for SQL Server VSS Writer.	1. Ensure that the Cohesity Agent is installed. 2. Provide local administrator rights to the CohesityAgent service.
	The service SQL Server VSS Writer is not running. Current state: <current_state>	Ensure that the Cohesity Agent service is running. If yes, try restarting it.

Health Check Name	Error Message	Corrective Action
SQL Server (SQL2012), SQL Server (SQL2016), SQL Server (SQL2017) **Description** Checks if all SQL instances are running.	Error querying service status for <display name>.	1. Ensure that the Cohesity Agent is installed. 2. Provide local administrator rights to the Cohesity Agent service. 3. Ensure the SQL instance <display name> is correctly installed.
	Querying registry for SQL instance names failed.	Provide local administrator rights to the CohesityAgent service.
	Inactive SQL instances detected: <display names>	Start the SQL instances if required.

Health Check Name	Error Message	Corrective Action
SQL Server (SQL2012), SQL Server (SQL2016), SQL Server (SQL2017), SQL Server VSS Writer, Cohesity Agent Service **Description** Validates that the required permissions have been granted to the logons for the services being checked.	Querying registry for SQL instance names failed.	Provide local administrator rights to the CohesityAgent service.
	Error querying service status for <display name>	1. Ensure that the Cohesity Agent is installed. 2. Provide local administrator rights to the CohesityAgent service. 3. Ensure the SQL instance <display name> is correctly installed.
	The service <display name> is not using a local admin or service logon.	Provide either admin rights or a default system service account to the service login.
	The logon <logon name> for the service <display name> is not a sysadmin of the SQL instances(s) <SQL display name>.	Assign the logon the sysadmin role in the specified SQL Server instance.

Chapter 27: Manage Protection Groups

From the Protection page, you can start the backup process by clicking **Protect** in the upper right corner. You can also view protection group settings, run a protection group, edit protection group settings, delete a protection group, deactivate a protection group or failover a protection group using the icons listed in the actions column of the table. (See below for details.)

This page also provides the following information:

- Overview of the protection groups that are already created. You can filter these protection groups based on:
  - Groups
  - Group types
  - Policies
  - SLA
  - Status
  - Organizations
- The run status for all the protection groups in the Cohesity cluster that meet the specified filter criteria.

The columns for the table at the bottom of this page are described below.

Column Name	Description / Action

Group Name	Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.
	You can view these details either in the details view or in calendar mode.
	• To view the details in the details view mode, click the icon on the right.
	• To view the details in calendar mode, click the icon displayed on the right corner of the page.
	  • You can view the status of the protection group based on the date you select from the calendar.
	  • If there is no protection group run scheduled on a day, then that day will be greyed out in the calendar.
	  • If there is a protection group run scheduled on a day, then that day will be highlighted in the calendar.
	  • If the scheduled protection group run fails, then the icon will be displayed against the scheduled day.

Organizations	Drill-down on this link to get details of the organization that is associated with the protection group.
Start time	The start time and date of the last run of this protection group.
Duration	The time duration taken to complete the last run of this protection group.
SLA	Reports the SLA status. –The last run of this Protection Group completed within the SLA time period specified for this protection group. –The last run of this Protection Group did not complete within the SLA time period specified for this protection group .

Column Name	Description / Action
Status	For the last run of this protection group, this column reports the current status: • **Succeeded**–The protection group run completed without error. • **Canceled**–A protection group run started and then was canceled in the Cohesity Dashboard The protection group run did not complete. • **Canceling**–In the Cohesity Dashboard, the cancel option was selected for the protection group run and the cluster is in the process of canceling the run but the run is not yet canceled. • **Failed**–The protection group run did not complete due to an error. • 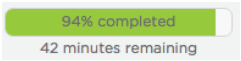 –The protection group run is currently running but is not completed yet. Additional information about some protection run issues may be available here.

Cancel Replication and Archive Copy Tasks of a Protection Run

A Protection Group with replication and/or archive schedules creates copy tasks that copy Snapshots to remote Cohesity cluster or External Targets. You can cancel these copy tasks for a specific protection run using the Cohesity Dashboard. Here are some example scenarios when you may want to cancel copy tasks:

- A protection run takes a lot of time to complete and you want to free the resources used by Runs of this Protection Group:
  1. Pause the protection run.
  2. Cancel the replication and archive tasks of past Runs of this Protection Group.
- You delete a Job that is replicating or archiving Snapshots and not all of the Snapshots captured by this Job have been copied to the remote Cohesity cluster or the External Target, the Cohesity cluster continues to replicate and archive the remaining captured Snapshots even after the Job has been deleted. Before or after deleting this Job, you may want to cancel the replicate and archive tasks of existing protection runs that copy the remaining captured Snapshots.

To cancel the copy tasks of a Protection Group Run:

1. Navigate to **Data Protection > Protection**.
2. From the list of Protection Groups, click on the Protection Group that has replication or archival running to view that Protection Group's details.

> **Tip:** Use the Status filter's Replication Running or Archival Running options to filter the Protection Groups.

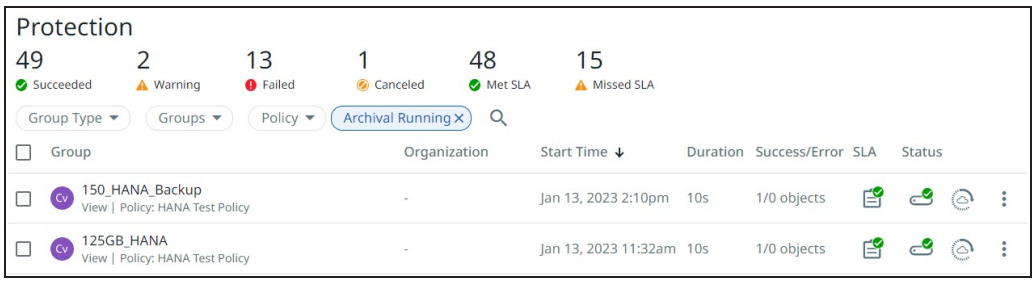

3. The Group Details page lists the runs for that Protection Group. Click the Actions menu (⋮) for a protection run and select **Edit Run**.
4. On the Edit Run page:
   - To cancel a replication copy task, under Replication, select **Cancel**.
   - To cancel an archive copy task, under Archival, select **Cancel**.

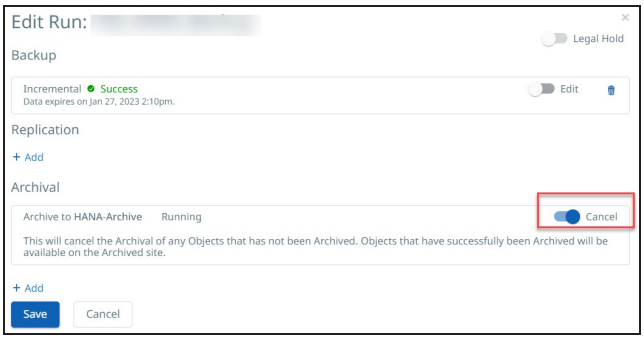

5. Click **Save**.

Delete Protection Groups, Snapshots and Objects

You can delete a Protection Group from the local Cohesity cluster where the Protection Group is running. As part of the Protection Group deletion, you can optionally delete the snapshots captured by the Protection Group. You can also delete snapshots or individual Objects for protection runs (without deleting the Protection Group) as described in Deleting Snapshots or Objects. However, the deletion of snapshots is possible only if they're not protected by the DataLock policy or if the DataLock policy has expired.

> **Note:** DataLock policy can be changed at any time. After you change the policy, the snapshots that were taken earlier will continue to follow the previous policy time and only the new snapshots will follow the time defined in the updated policy. This means if you've set a 15-day DataLock policy for the snapshots and you change it to 7 days, then the old snapshots, which were taken before the policy change, will continue to be protected for 15 days from the day of its creation. However, any new snapshots that are taken after the policy change will be protected for 7 days.

Deleting a Protection Group

You can delete a Protection Group if you:

- Do not want to retain the backups, local and archived snapshots that are generated by the Protection Group.
- Want to reclaim the storage space utilized by the Protection Group.

Cohesity provides the following ways to delete a Protection Group:

- **Delete Protection Group Only**: Allows you to delete the Protection Group. All the backups generated by the Protection Group are preserved till the snapshots expire.
- **Delete Protection Group and Snapshots**: Allows you to delete the Protection Group and all the local and archived snapshots generated by the Protection Group.

> **Note:**
> - After deleting the Protection Group, the reclaimed storage space details will not be displayed immediately. It will take some time to reflect on the Cohesity Dashboard.
> - If a Protection Group is deleted manually while preserving the snapshots, all the unexpired snapshots will be available for recovery. However, when the snapshots expire as per the saved policy or by manual expiry, those expired snapshots will not be available for recovery.

Delete a Protection Group

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, select the Protection Group that you want to delete, click the actions menu (⋮), and select **Delete**.

or

3. Drill down on a Protection Group name in the **Group** column of the list. In the Protection Group Details page, click the actions menu (⋮) and select **Delete**.
4. In the **Delete Protection Group?** dialog box that appears, select any one of the following options:
   - **Delete Protection Group Only**
   - **Delete Protection Group and Snapshots**

   If you select **Delete Protection Group and Snapshots**, any local snapshots captured by this Protection Group and stored on the Cohesity cluster are deleted. In addition, any archived snapshots stored on External Targets that were captured by this Protection Group are also deleted. (To selectively delete the snapshots for specific protection runs without deleting the Protection Group, see Deleting Snapshots or Objects).

   If the current Protection Group is replicating snapshots to a remote cluster or archiving snapshots to an external target, you should probably cancel these replication or archive tasks before deleting the Protection Group. For instructions, see Cancel Replication and Archive Copy Tasks of a Protection Run.

   If the Protection Group on the local cluster is replicating to a remote cluster, the snapshots that were replicated (copied) from the local cluster will also be deleted. Archived snapshots are also deleted.

Deleting Snapshots or Objects

You can delete the Snapshots or individual Objects for specific protection runs of a Protection Group without deleting the entire Protection Group. This feature is useful if a long retention time was specified for Snapshots and the disk space on the Cohesity cluster is getting low. You can clean up Snapshots to reclaim space on the Cohesity cluster.

Delete Snapshots for Protection Runs on the Protection Group Details page

1. Log in to the Cohesity cluster that is storing the snapshots that you want to delete.
2. Select **Data Protection > Protection**.
3. On the **Protection** page, drill down on a Protection Group name in the **Group** column of the list.

   In the Protection Group Details page, a checkbox appears in the **Start Time** column. In addition, check boxes also appear next to the protection runs.
4. Select protection runs by clicking the check boxes in the **Start Time** column. All the snapshots that were captured by the *selected* protection runs will be marked for deletion. (Some limitations apply, see below.) To select all the protection runs displayed on the current page, click the check box in the **Start Time** header. If you

cannot select a check box, the snapshots for the protection run are already marked for deletion or the protection run has not been completed.

> **Note:** If there are multiple pages of protection runs, selecting the check box in the **Start Time** header does not select the protection runs on the other pages. It only selects the protection runs on the current page.

5. Click the **Delete All Snapshots** icon. If possible, all the snapshots that were captured by the *selected* protection runs are marked for deletion. Note the following about the snapshot deletion process:
    - The Cohesity cluster runs the Snapshot deletion tasks asynchronously in the background. Snapshots that are marked for deletion may continue to appear in the Cohesity Dashboard until the deletion tasks are complete. The deletion task takes only a few minutes to complete. Once the snapshot is deleted, it is marked as a "deleted snapshot" in the Cohesity Dashboard. You cannot restore the deleted Snapshot.The Cohesity cluster reclaims the space utilized by the deleted snapshot asynchronously in the background. You can increase the priority of the space reclamation by throttling the system activities running in the background. For more information, see Throttling Background System Activities.
    - If a selected protection run spawns copy tasks that replicate or archive Snapshots and those copy tasks have not been completed before the request for deletion is received, the Snapshots for the protection run are not deleted even after the copy tasks are complete. If other protection runs are selected and those runs do not have these restrictions, the snapshots captured by these protection runs are marked for deletion.
    - If the Protection Group on the local cluster is replicating to a remote cluster, the snapshots that were replicated (copied) from the local cluster will also be deleted. Archived snapshots are also deleted.

Delete Snapshots for a single Protection Run on the Run Details page

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the **Protection** page, drill down on a Protection Group name in the **Group** column of the list. The Protection Group Details page appears.
3. In the **Start Time** column, drill down on a protection run. The Run Details page appears.
4. Click **Delete All Snapshots**. If possible, all the snapshots captured by the selected protection run are marked for deletion. Note the following about the snapshot deletion process:
    - The Cohesity cluster runs the Snapshot deletion tasks asynchronously in the background. Snapshots that are marked for deletion may continue to appear in

the Cohesity Dashboard until the deletion tasks are complete. The deletion task takes only a few minutes to complete. Once the snapshot is deleted, it is marked as a "deleted snapshot" in the Cohesity Dashboard. You cannot restore the deleted Snapshot.The Cohesity cluster reclaims the space utilized by the deleted snapshot asynchronously in the background. You can increase the priority of the space reclamation by throttling the system activities running in the background. For more information, see Throttling Background System Activities.

- If a selected protection run spawns copy tasks that replicate or archive Snapshots and those copy tasks have not been completed before the request for deletion is received, the Snapshots for the protection run are not deleted even after the copy tasks are complete.
- If the Protection Group on the local cluster is replicating to a remote cluster, the snapshots that were replicated (copied) from the local cluster will also be deleted. Archived snapshots are also deleted.

Delete individual Objects within a Snapshot

> **Note:** Deleting a single MS SQL database within a snapshot is not supported.

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. On the Protection page, drill down on a Protection Group name in the **Group** column of the list. The Protection Group Details page appears.
3. In the **Start Time** column, drill down on a protection run. The Run Details page appears.
4. Select Objects by clicking the checkboxes in the Object name column. To select all the Objects displayed on the current page, click the check box in the Object name header.
5. Click the **Delete Snapshots** icon.

   If the Protection Group includes replication, the confirmation includes the following additional information:

   - The Cohesity clusters that the snapshot was replicated to.
   - Replicated snapshots are not deleted through this workflow. To delete them, you must log in to those remote Cohesity clusters and delete the snapshots there.

   > **Note:** You cannot delete individual objects within an archived snapshot.

6. Click **Delete** to confirm.

Protection Group Run Issues

If a protection run fails, additional information may be available below:

- If backing up MS SQL Server fails with the following error: Backup task failed with error: [kGuestOperationsError]: Error while pushing the agent file, it means Cohesity was unable to push its agent file to the VM. Possible causes: VMware Tools went into a bad state, VM was not running in high performance mode.

  Possible fix - Restart VMware Tools and ensure the VM is running in high performance mode.

- Lingering VSS shadows on the VM can cause backups to fail with the following error: "Backup task failed with error: [kDatabaseBackupMismatch]: Incompatible LSN numbers for..." Running `vssadmin list shadows` in elevated mode in the CLI of the VM encountering issues shows the VM's existing shadows.

  Workaround - If no other backup apps are running on that VM, you can delete all VSS shadows using this command:

  ```
  PS C:\Windows\system32> vssadmin delete shadows /All
  ```

  If the previous command does not work, open DiskShadow and use this command:

  ```
  DiskShadow> delete shadows all
  ```

- If you encounter SOAP errors between Cohesity cluster and VMware vSphere, refer the Knowledge Base article, Troubleshooting SOAP errors between Cohesity cluster and VMware vSphere to troubleshoot the commonly seen errors during the SOAP exchange.
- If you encounter any VSS errors in a Cohesity environment, refer the Knowledge Base article, Troubleshooting VSS issues in a Cohesity environment. These error messages generally point to an environment issue and not a Cohesity protection or recovery issue.

Add or Remove a Legal Hold to a Snapshot

Users who are assigned the Data Security role can put a legal hold on existing Snapshots to preserve them for legal purposes. Legal hold Snapshots cannot be deleted until the legal hold is removed. Legal hold can be added to unlocked and DataLocked Snapshots. Cohesity recommends that their data security user enable Legal Hold for business-critical workloads after a ransomware attack is confirmed.

You can add a legal hold to Backup, Replication, and Archival targets in the protection run or individual objects in a protection run:

- If you add a legal hold to a protection run, it applies to all the Snapshot objects that were backed up by that protection run.
- You must manage the legal hold status on the remote replication Cluster manually.

> **Note:** Legal hold prevents Snapshots from being deleted until the legal hold is removed. Using legal hold for long periods of time may result in the cluster running out of space

To add a legal hold:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click a Protection Group name.
3. To add a legal hold for Backup, Replication and Archival targets in a protection run:
   a. Click the Actions Menu (⋮) and select **Edit Run**.
   b. Toggle on **Legal Hold** for the target that requires a legal hold. This option is enabled only if you have the Data Security role. Users assigned with other roles can only view the Legal Hold status.
   c. Click **Save**.
4. To add a legal hold to objects in the protection run:
   a. Drill-down on a protection run start time.
   b. Click the checkbox next to objects and then click the **Add Legal Hold** gavel icon ( ) displayed at the top of the list. The icon is enabled only if you have the Data Security role and you have selected at least one object.

   The legal hold is not added to the replicated objects. Consider adding legal hold manually to objects on the remote replication cluster.

A gavel icon ( ) indicates that a protection run or object has a legal hold on it.

To remove a legal hold:

1. In the Cohesity Dashboard, select **Data Protection > Protection**.
2. Click a Protection Group name.
3. To remove legal hold from Backup, Replication and Archival targets in a protection run:
   a. Click the Actions Menu (⋮) and select **Edit Run**.
   b. Toggle off **Legal Hold** for targets that do not require a legal hold. This option is displayed only if you have the Data Security role.
   c. Click **Save**.
4. To remove legal hold from objects in the protection run:
   a. Drill-down on a protection run start time.

      Legal hold objects have a gavel icon ( ) next to the object name.

b. Click the checkbox next to the objects and then click the **Remove Legal Hold** gavel icon ( ) displayed at the top of the list. The icon is displayed only if you have the Data Security role and have selected at least one object.

The legal hold is not removed from the replicated objects. Consider removing legal hold manually from objects on the remote replication cluster by signing on to the remote cluster and following the procedure above.

Considerations

- Only users assigned with a Data Security role can add and remove Legal Hold. Users with other roles can only view the Legal Hold status.
- Synchronization of Legal Hold happens only from the local cluster to the replication cluster and not from the replication cluster to the local cluster.
- The Legal Hold toggle for the whole local copy is enabled only if all the objects are on Legal Hold.
- The legal hold cannot be applied at the object level for SQL objects; it is only supported at the run level.

Failover a Protection Group

If you are failing over a Protection Group for VMs, follow the steps in Failover Protection Group for VMs.

You cannot failover the Protection Group of views and SQL servers.

You can, however, clone a view from a view snapshot stored on the remote Cohesity cluster. For detailed instructions, see Clone a View From a Snapshot.

You can also restore SQL servers from a SQL server VM snapshot stored on the remote Cohesity cluster. To learn more, see About Restoring Microsoft SQL Server Databases.

> **Warning:** For replicated MS SQL Server VM Snapshots, you can only restore MS SQL Server to a new location. (You cannot restore MS SQL Server to its original location.)

Failover Protection Group for VMs

> **Note:** Failover of Protection Groups is only supported for VMware VMs.

When a Protection Group with a replication schedule starts replicating snapshots from a primary Cohesity cluster to a remote Cohesity cluster, a Protection Group in Failover Ready state is created on the remote Cohesity cluster. If the original capturing Cohesity cluster and/or the original Source are down because there is a planned maintenance outage or a disaster destroys them, you can failover the Protection Group in Failover Ready state and deactivate the original task on the capturing Cohesity cluster. As part of the Failover, you are prompted to optionally recover the VMs from the replicated snapshots stored on the remote Cohesity cluster.

> **Note:** When you create a Protection Group on a Cohesity cluster with a replication schedule after the first snapshot is replicated, a Failover Ready copy of the Protection Group is created on the remote Cohesity cluster that is receiving the snapshot. For more information on Failover Ready, see About Failover Ready. You can failover the Protection Group of VMs. You cannot failover the Protection Group of Views or SQL Servers. You can, however, clone replicated Views and restore SQL Servers. For more information, see Failover a Protection Group.

During Failover, the following actions occur:

1. **Disaster at Source** - On the primary Cohesity cluster, the backup job fails to capture the snapshots of the source.
2. **Replication Fails** - As the source is down or unreachable, the replication job fails to replicate the snapshots onto the remote Cohesity cluster. The Failover Ready for VMs on the remote Cohesity cluster is activated and causes the rejection of the incoming replicated snapshots created by the associated job on the primary Cohesity cluster.
3. **Failover VMs** - You are prompted to optionally recover the VMs from the replicated snapshots on the remote Cohesity cluster to a new source.

The state of a Protection Group may also become "Failover Ready" when an active Protection Group is deactivated. For more information, see About Failover Ready.

For more information about what happens during a Failover, see Failover from Replicated Copies.

For more information on disaster recovery of VMs using the Cohesity Cloud Edition, see Back Up VMs to On-Prem, Replicate and Failover to AWS Cloud.

To Failover a Protection Group on the remote Cluster:

1. Navigate to **Data Protection > Protection**.
2. Find the Protection Group in the Failover Ready state in the list, click on the actions menu ⋮, and select **Failover**.

3. Select a source to place the recovered VMs and click **Failover Job and Continue to Recovery**.

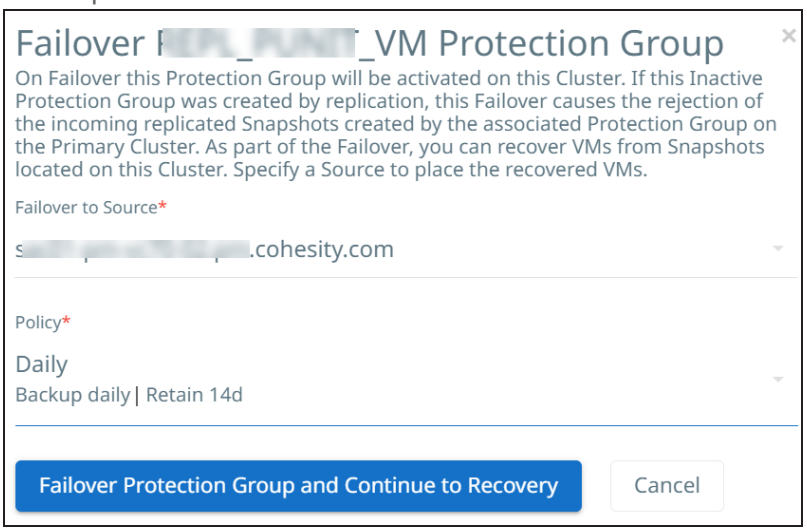

4. Set Recover Options for VM Failover.
5. View VM Failover Confirmation.
6. Deactivate Original VM Protection Group. The Failover of the Protection Group in the Cohesity Dashboard does not automatically deactivate the original task on the original cluster that captured the snapshots. To deactivate, you must follow the instructions in Deactivate Original VM Protection Group
7. Restore VM Protection Group to Running. The Failover of the Protection Group in the Cohesity Dashboard does not automatically restore the task to running on the new cluster. To restore the task to a running state, you must follow the instructions in Restore VM Protection Group to Running.

Set Recover Options for VM Failover

In this section, you select the recover options while recovering VMs as part of the Failover of an Inactive Protection Group.

To set recover options:

1. Toggle **Rename Recovered VMs** on if you want to add Prefix and/or Suffix strings to the VM names of the new VMs created by this task. For example, you could specify Prefix `Test_` and Suffix `_QA` to the VM name `WinOS` and when the VM is recovered, the new VM that is created is called `Test_WinOS_QA`.
2. Select the appropriate **Recovery Location** option
   - **Recover back to original location**—Select this option to recover the VM files (such as the VMDK files) to their original datastores and create new instances of the VMs in the original location in the original source. For more information, see Recover to the Original Location.
   - **Recover to a new location**—Select this option to recover the VM files (such as the VMDK files) to an alternate datastore and create new instances of the

VMs in a new location of a registered Source. For more information, see Recover to a New Location. For details about virtualization or cloud platform-specific options, see the respective section.

- Acropolis
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Storage Container:** Select a storage container to store the recovered VMs.
- AWS
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Region:** Select a region in the project to recover the VMs.
  c. **VPC:** Select an Amazon Virtual Private Cloud (VPC) to store the recovered VMs.
  d. **Subnet:** Select a subnet in the Amazon VPC to store the recovered VMs.
  e. **Key Pair:** Select an Amazon key pair that has permission to log in to the recovered VMs.
  f. **Network Security Group:** Select the associated network security group.
- Azure
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Resource Group, Storage Account:** Select the corresponding Resource Group and Storage Account to which the recovery has to be done.
  c. **Storage Container:** Select the storage container in which page blobs will be created.
  d. **Virtual Network:** Select the virtual network where the VM must be placed.
  e. **Subnets:** Subnets in the virtual network where VMs must be placed.
- GCP
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Project:** Select a Google Cloud Platform (GCP) project to recover the VMs to.
  c. **Region:** Select a region in the project to recover the VMs to.
  d. **Availability Zone:** Select an availability zone in GCP to recover the VMs.

e. **VPC:** Select a Virtual Private Cloud (VPC) to store the recovered VMs.

- Hyper-V
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Compute:** Select the compute host for the recovered VMs.
  c. **Volume:** Select a volume to store the recovered VMs.

  When recovering to a standalone Hyper-V host as a new location, SMB volumes are not supported locations.
- RHV
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Data Center:** Select a data center to store the recovered VMs.
  c. **Cluster:** Select an RHV/OLVM to store the recovered VMs.
  d. **Storage Domain:** Select an RHV/OLVM storage domain to store the recovered VMs.
- VMware
  a. **Source:** Select a registered Source to recover the VMs to.
  b. **Resource Pool:** Select a Resource Pool to store the recovered VMs.

  > **Note:** When protecting a VMware Source, if you select a Resource Pool that is not part of the ESX cluster, all datastores in the vCenter are displayed. Ensure you select a datastore that is part of the selected Resource Pool. If you select the wrong datastore, storage migration fails, but the VM remains attached to the ESX Host and runs from the Cohesity datastore. You can manually migrate the VM to the desired datastore at a later time.

  c. **Datastore:** Select one or more available target datastores. The VM files (such as VMDK files) extracted from the Snapshots are stored in the selected datastores. The number of target datastores you select must be equal to or less than the number of VMs selected for recovery.

  Cohesity also displays each datastore's amount of free space and total capacity. For more information about recovering VMs to multiple datastores, see Recover VMs to Multiple Datastores.

> **Note:** Not applicable when recovering to a standalone ESX host. After recovery, VMs continue to run on the Cohesity datastore because Storage vMotion is not supported.

d. **VM Folder** Optionally select a logical VM folder in the datacenter in which to create the VMs.

> **Note:** Not applicable when recovering to a standalone ESX host.

- vCloud Director
  a. **Source:** Select a registered vCloud Director source to which you plan to recover the VMs or vApp.
  b. **Organization VDC:** Select the target organization VDC.
  c. **vApp:** Optional. Select the target vApp. If not specified, VMs are recovered as standalone VMs inside the target organization VDC.
  d. **Storage Profile:** Select the target Storage Profile. The Cohesity cluster chooses any compliant datastore associated with the Storage Profile to store the VM files. If you are recovering multiple VMs, the Cohesity cluster chooses one or more datastores associated with the Storage Profile.

3. Select the appropriate **Recovery Type** option.
   - **Instant Recovery**: if you select this option, the VM(s) will be usable instantly in the target environment and will be moved to target storage later. The VM can be used in production during Instant Recovery when data is copied to the VM from the Cohesity cluster. In such scenarios, the VM performance may be low as the Storage vMotion is still in progress.

     When recovering VMware VMs, by default, the selected VM is always instantly available after recovery.

     > **Note:** For instant recovery to Standalone ESXi, the VM(s) will be available immediately in the target environment and will have to be moved to the target storage manually.

   - **Copy Recovery**: You can select this recovery type for **VMware** sources only. If you select this option, the recovered VMs will be usable in the target environment only after all the data is copied from Cohesity to the target

storage. Once the restore is complete, the VM will be available to use and the performance of the VM will be normal.

> **Note:** You should have `VirtualMachine.Config.Rename` permission on your user account to use this recovery option.

Considerations when using the Copy Recovery option to recover VMware sources

- This is the only option to recover data from VMware cloud solutions.
- Copy recovery is not supported on Hybrid Extender deployments. Only vCenter and stand-alone ESXi hosts are supported.

4. Specify the networking options.
The network options displayed depend on the Recovery Location option that was selected.
   - **Keep original**—For each recovered VM, keep the original virtual Network Interface Cards (vNICs) and attach them to the original network connections. This option is only supported when VMs are recovered back to their original location.
   - **Start Connected**—For each recovered VM, connect to the original or new network when the VM reboots. If this option is not selected, the VMs are not connected to any network on reboot.
   - **Detach network**—For each recovered VM, the virtual Network Interface Card (vNIC) is removed from the VM.
   - **Attach to a new network**—For each recovered VM, attach the virtual Network Interface Card (vNIC) to a new network. The new network might be a standard port group or a distributed port group.

     When recovering VMware VMs, the selected Network port group will be attached to all the network interfaces of the VMs being recovered.

     When recovering Hyper-V VMs, the selected virtual switch will be attached to all the network interfaces of the VMs being recovered.

Note:

- Recovering a network attached CentOS VM on VMware or Hyper-V does not attach the VM to the network. Attach the VM to the network manually.
- If the VMware Resource Pool is not part of the ESX cluster, all networks available to the Datacenter are displayed. Ensure you select a network available to the selected Resource Pool. If you select the wrong network, VM creation and recovery will fail.

5. **Auto Select:** When enabled, the recovery task automatically selects the correct VLAN.
6. Select **Instant Recovery** if you want to change the Hyper-V recovery mechanism so that VMs become available instantaneously after recovery. Applicable to Hyper-V only.

   Select **Preserve VMs UUID** if you want to preserve the UUID for the recovered VMs. Applicable to Hyper-V only.
7. Select **Leave recovered VMs powered off** if you want the recovered VMs to remain powered off after they are created.

   If you recover VMs from a Storage Domain that has Cloud Tier enabled, Cohesity recommends leaving the VMs powered off.
8. Select **Continue recovery even if errors occur when recovering VMs** if you want the Recover task to continue even if errors occur when recovering the VMs. For example, if one of the VMs cannot be created, the Cluster will still attempt to create the other VMs. If this is unselected, then the Recover task is discontinued only if VM creation fails, so if VM creation succeeds, but VM storage migration fails, the task continues.
9. In the Selected Objects section, click **Show VMs** to view the VMs associated with the Failover Job. By default, the VM Snapshots of the last Run of the Job are recovered.

   Click Edit ( ) to select Snapshots captured by a different protection run.
10. Click **Finish**.

The next step is View VM Failover Confirmation.

View VM Failover Confirmation

In this step, you can view a summary of the options for the current Failover task and a list of objects that will be recovered by the task.

The next step is to Deactivate Original VM Protection Group

Deactivate Original VM Protection Group

Only Protection Groups that are protecting VMs can be deactivated. Protection Groups that are protecting a view cannot be deactivated.

> **Note:** If possible, deactivate the original Protection Group on the original capturing Cohesity cluster. If the original Cohesity cluster that captured the Snapshots is down or unreachable, you may be unable to deactivate the Protection Group.

The example below shows the deactivation of a Protection Group protecting VMs. Deactivating a Protection Group protecting SQL Servers is similar but not all details are the same. For example, the transaction logs are not described or pictured.

During deactivation, the following actions occur:

1. Snapshots are no longer captured by the original Protection Group on the original capturing Cohesity cluster. However, the existing snapshots captured by the Protection Group remain in the storage domain. They are not deleted.
2. Snapshots created by the Protection Group are no longer replicated to the receiving (remote) Cohesity cluster. (If you have already initiated the failover procedure, the replication has already stopped.)
3. You can optionally choose that the VMs on the original source be powered off. The VMs and the VM files, however, are not deleted. (The Cohesity cluster may be unable to power off the VMs if the original source is down or unreachable.)

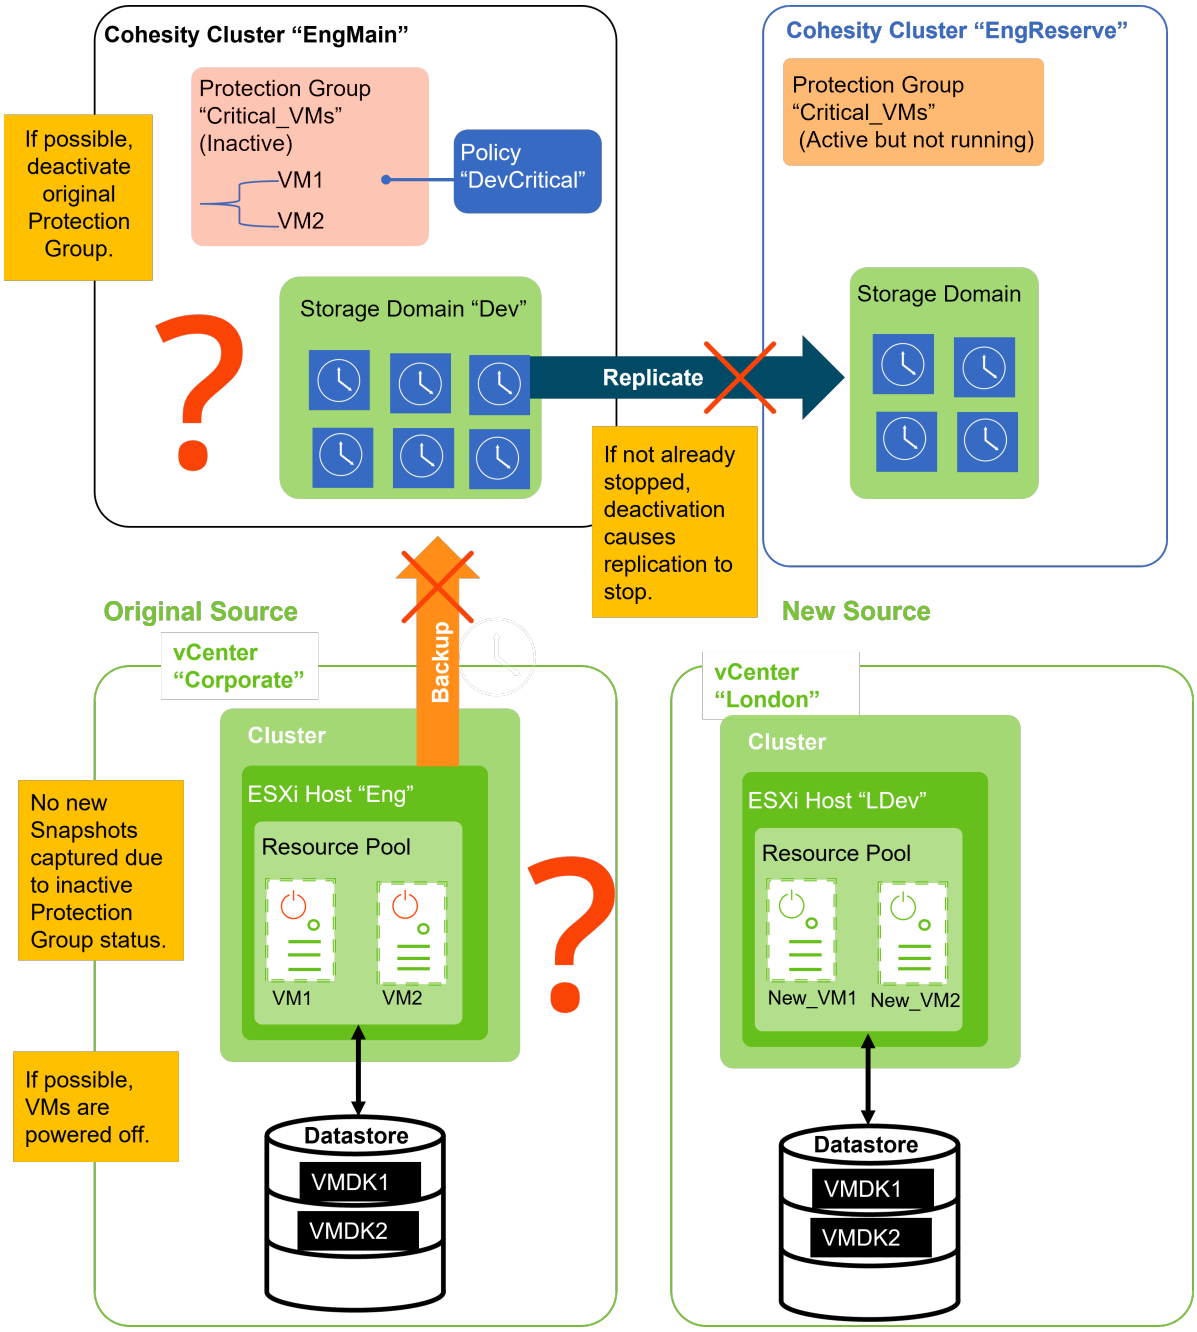

To deactivate the original Protection Group:

1. Log in to the Cohesity cluster where the original Protection Group was running.
2. Navigate to **Data Protection > Protection**.
3. Find the original Protection Group in the list, click on the actions menu ⋮ and select **Deactivate**.
4. Click **Deactivate and Power off VMs**, or click **Deactivate Only** if you do not want

to power off the VMs.

5. Optionally delete the VMs and VM files using the VMware vSphere client.

Restore VM Protection Group to Running

You can restore the Protection Group activated by the Failover, by assigning a new policy to the Protection Group and selecting Objects (VMs) to back up. These actions restore the Protection Group to a running state on the new cluster. To learn more, see Restore Activated Protection Group.

To restore the Protection Group:

1. Navigate to **Data Protection > Protection**.
2. On the **Protection** page, find the inactive Protection Group (created by replication).
3. From the Actions menu (⋮) of that Protection Group, select **Resume**.

This is the last step in the Failover process.

Chapter 28: Manage Recovery

This chapter includes the following topics:

Use Wildcard * in Search for Recovery 4441

Cancel a Recovery 4442

Use Wildcard * in Search for Recovery

You can optionally use the Wildcard `*` in any location in the search string. The following table shows some example searches.

Example Search	Finds	Description
`*`	All the files, Protection Group Names, or Server names	
`test*`	`vol1/usr/bin/test` `vol1/usr/share/setup/test.ps`	Note that this search does not find: `vol1/usr/memtest.ext`

`*test*`	`vol1/usr/bin/test`
	`vol1/usr/share/setup/test.ps`
	`vol1/usr/memtest.ext`	

`*xml`	`usr/local/firefox/searchplugin/yahoo.xml`	When searching for file extensions, do not enter the period in the search string. For example if searching for all files that end in the `.txt` extension, enter `*txt` and not `*.txt`.

Example Search	Finds	Description
`search* *xml`	`usr/searchplugins/dsl.xml`	Enter `search* *xml` with a space and not `search*xml`. However, when you add spaces between search terms, only files that match all the search terms are found, for example searching with the string: `search* *xml` does not find: `usr/searchplugins` Entering just `search` does not find `usr/searchplugins/dsl.xml`.
`*plug* *xml`	`usr/searchplugins/dsl.xml`	Only files that match all the search terms are found.

> **Note:** Ensure to follow the standard regular expression rules. Specify a dot '.' before asterisk '*' to match all characters, for example, L:.*.ndf.

Cancel a Recovery

The recovery of any objects that have not yet been recovered is canceled. Note the following:

- Canceling a recovery might require some manual cleanup. For example, if you cancel a folder recovery, any files copied prior to the cancellation will not be deleted.
- When canceling a VMware VM recovery, the cancellation will occur after the VM is created, the datastore is mounted and the Storage vMotion migration process begins. You will need to manually clean up the VM that was created and unmount the datastore.

Cancel a Running Recovery

1. In the Cohesity Dashboard, select **Data Protection > Recoveries**.
2. Click the recover task you want to cancel. The task details page is displayed.
3. Click **Cancel** in the upper right corner of the details page.
4. Respond to the confirmation prompt.

Chapter 29: Cohesity Agents

Data protection is paramount in today's complex IT environments, whether for physical servers, virtualized infrastructure, or databases. Cohesity provides a comprehensive solution to safeguard data across different systems using specialized software components known as agents. These agents are crucial in streamlining the data protection process across various platforms, ensuring smooth backup and recovery operations tailored to each environment.

What are Cohesity Agents?

Cohesity agents are lightweight software applications installed on target machines running supported operating systems—whether to protect physical servers, virtual environments, or databases. In particular, when Cohesity agents are used to protect filesystem data, they are referred to as **Cohesity Physical Agents**.

Once installed, these agents enable the Cohesity platform to manage backup and recovery through block-based or file-based **Protection Groups**. These Protection Groups determine how data is protected, while the agents facilitate interaction between the Cohesity cluster and the protected workloads.

The Physical agents also contain modules to protect other data, such as MS SQL databases. Additionally, the agents protect other workloads, such as SAP HANA, in conjunction with other Cohesity modules described in the documentation for the respective integrations.

What Do Cohesity Agents Do?

The core purpose of Cohesity agents is to enable seamless data protection by performing key tasks specific to the environment in which they are deployed:

- **Backup and Recovery**: Agents handle the backup of system, application, and database data, securely storing it within the Cohesity platform for fast recovery when needed.

- **Block-Based or File-Based Protection**: The Cohesity cluster can support block-based and file-based backups, depending on the type of agent and the system it protects.
- **Application Consistency**: Cohesity agents capture application-consistent backups for critical workloads like databases, ensuring data is protected without the risk of corruption.
- **Data Management**: Agents manage the scheduling, execution, and monitoring of backups, integrating with the Cohesity platform for centralized control and reporting.
- **Efficiency**: Using technologies such as deduplication and compression, agents optimize data transmission, reducing storage consumption and bandwidth usage.

Cohesity Agents Deployment in Different Scenarios

Physical Server Environments

In physical environments, **Cohesity Agents** are installed directly onto the operating system. These agents ensure that all critical data from the server, including file systems, applications, and system state, is fully protected. Administrators can use block-based or file-based **Protection Groups** to define backup schedules and policies depending on the system's configuration. This is particularly useful for traditional systems or workloads that are not virtualized, as the agents protect the entire physical infrastructure. It is important to note that nothing precludes the agents from being used inside virtual environments, which are described in the next section.

Virtual Environments

In virtualized environments, such as **VMware** or **Hyper-V**, Cohesity supports both agentless and agent-based protection. When using agents, they are typically installed on virtual machines (VMs) requiring more granular control over backup operations, such as specific applications or critical data. Cohesity can also perform agentless protection by interacting with the hypervisor directly, to allow for better management of workloads at scale.

In cases where the agent-based approach makes sense, the Cohesity agent ensures that the VM and its underlying applications are backed up in an application-consistent manner, providing the option for individual VM restores or more complex recovery scenarios.

Databases

Cohesity offers specialized agents for database environments, particularly for SQL Server, Oracle, and other enterprise databases. These agents are designed to work closely with database applications to ensure consistent and reliable backups. For example:

- **SQL Server**: The agent interfaces with the SQL Server application to take application-consistent snapshots, enabling log truncation, point-in-time recovery, and rapid restoration of databases.
- **Oracle Databases**: Similarly, the Cohesity agent integrates with Oracle RMAN (Recovery Manager) to perform both full and incremental backups, ensuring fast recovery and minimal impact on performance.

For databases, ensuring data consistency is critical. Cohesity agents allow administrators to leverage granular control over backup policies, ensuring that the databases can be restored to a specific state without loss of data integrity.

In summary, Cohesity agents are versatile tools designed to protect data in **Physical**, **Virtual**, and **Database** environments. They are essential for enabling tailored data protection strategies, ensuring that regardless of the infrastructure, backups are consistent, efficient, and easily recoverable. Whether protecting a traditional physical server, a modern virtual environment, or an enterprise-grade database, Cohesity agents provide the backbone for robust, reliable data protection.

Install and Manage Cohesity Agents

Installing the Cohesity Agent is required to protect physical servers and some hypervisor environments. The following topics discuss Agent installation and management.

Cohesity Java Agent Upgrade

> **Note:**
> - The Cohesity Java agent upgrade steps are common for all the platforms.
> - Automatic upgrades are available from the Cohesity cluster UI and for manual updation follow the steps listed below.

To upgrade the Cohesity Java Agent follow the steps listed below.

You can download the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Download the installer:
    a. Using the Cohesity cluster UI

       Navigate to **Data Protection** > **Sources** and click **Download Cohesity Agent**.

    b. Using the Cohesity Package Download site

i. Download the agent installer to the protected server.

ii. Refresh the agent from the Cohesity Dashboard to ensure that the current agent is connected to the Cohesity cluster.

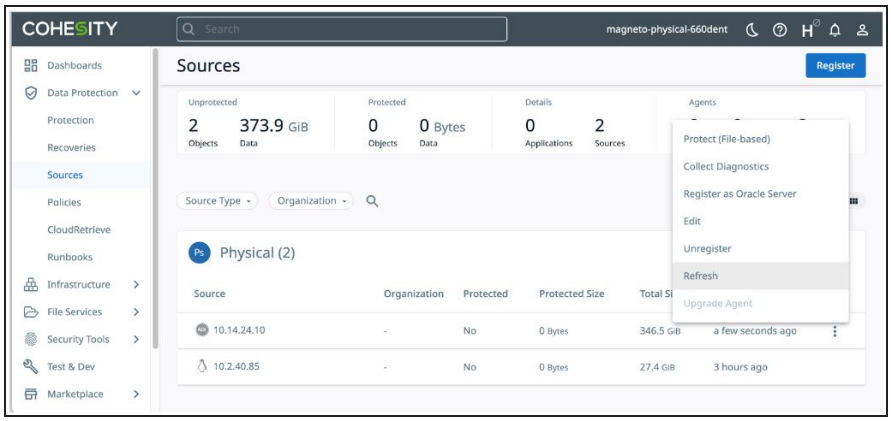

iii. Note down the current agent version and the last refresh time from the Cohesity Dashboard.

iv. Login to the server and upgrade the agent using the platform-specific command as per the below table.

Platform	Upgrade Command
Solaris	`pkgadd -d <installer>`
AIX*	`installp -ac -d <installer> all`
Linux (PowerPC)	`rpm -U <installer>`
Linux (z/Architecture)	`rpm -U <installer>`
Linux (ARM)	`rpm -U <installer>`
HP-UX	`swinstall -s <installer absolute path> -x reinstall=true \*`

*If the AIX agent upgrade fails due to version mismatch, then use the `force install` command:

```
installp -F -ac -d <installer> all
```

v. Refresh the agent from the Cohesity Dashboard and verify that refresh time and agent version are updated.

For more information on how to install and manage agents, see Manage Agents.

Install and Manage the Agent on Windows Servers

> **Note:** If the Windows firewall is active upon installing the Cohesity Agent, you need to add a rule in the firewall to open the port 50051 for communication with the Cohesity nodes. For more information, see Manage Firewall Ports.

Cohesity Agent Component Requirements

Refer to the following table to determine which components to install with the agent service for the server type you want to register/protect.

Server	Install Volume CBT	Install File System CBT
Physical Server without MS SQL	Yes (allows incremental backup) **Note:** If you run file-based backups, then this driver is not required.	No
Physical Server with MS SQL	Yes (allows incremental backup)	Yes (allows file-based SQL backup)
MS Exchange Server	No	Yes
VMware VM (using Installed agent) with MS SQL	No	No
Hyper-V SCVMM server	No	No
Hyper-V 2016 host	No	No
Hyper-V 2012 R2 host	No	Yes
Nutanix AHV	No	No

Free Space Requirement for Block-based Backup

For block-based backup of Windows physical server, the space must be at least equal to the calculation (Initial disk size*Churn rate*Backup window time). If you're running multiple

backups on a cluster, you'll need enough storage capacity to accommodate the snapshots for each of the machines using this calculation.

Download and Install the Agent

Because of weaknesses in the SHA-1 algorithm and to align to industry standards, starting from the Cohesity version 6.7 onwards, sha1 signing of Cohesity artifacts is dropped. Windows Agent installer which was previously signed with sha1 and sha256 hashing algorithms is now signed with only sha256 algorithm. Due to this, legacy OS users will notice the following warning when installing the Cohesity Windows agent.

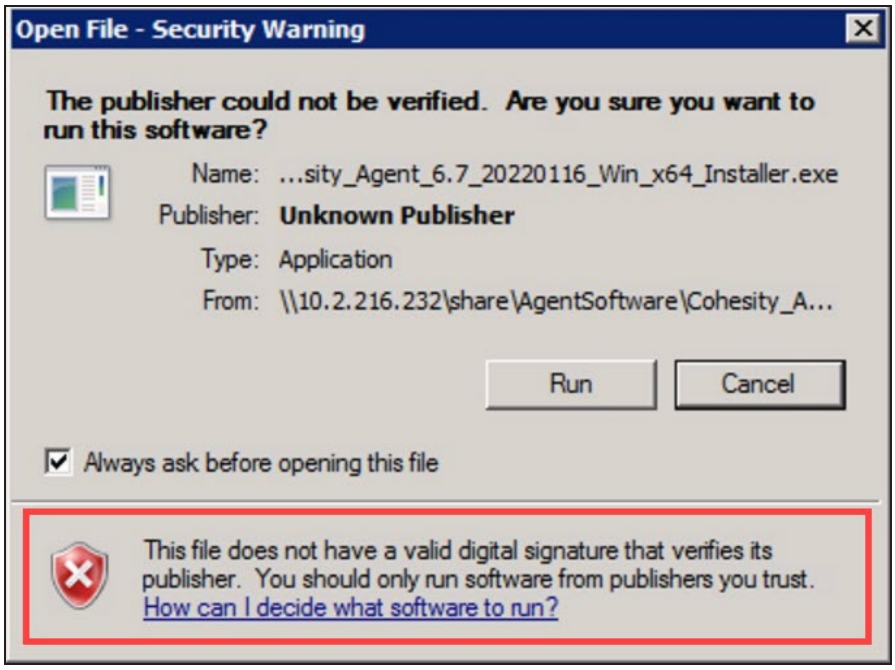

To avoid this warning, install sha2 code signing support on the system where you are installing the agent.

> **Note:** This is only a warning, and you can click **Run** to continue with the installation.

To download and install the Cohesity Windows Agent:

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.

   Note that the Cohesity Windows Agent does not support file path names that are not compliant with UNICODE or UTF-8 encoding. The table below lists the encodings that are not supported.

Language	Locale	Windows Code Page
Simplified Chinese	zh_CN	936
Traditional Chinese.	zh_TW	950
Japanese-SJIS	ja_JP	932
Japanese-EUC	ja_JP	20932
Korean	ko_KR	949

The Cohesity Windows agent installation may fail even if the encodings are not supported.

3. As an administrator with local system privileges on that server, run the executable with **Run as administrator** and complete the installation wizard.Install the **Volume CBT (Changed Block Tracker)** component to perform incremental backups. Installing this component requires a reboot. Without a reboot, only full backups can be performed. Not installing this component means only full backups can be performed.The Cohesity Volume CBT driver creates the **Cohesity_Bitmap.cty** file at the root of each volume. There are multiple ways to backup Windows physical server but the **Cohesity_Bitmap.cty** file will be available on the system only if the Volume CBT driver is installed. This file is used to maintain the CBT data for incremental block-based backups across Windows server reboots. Do not move or delete the "**.cty**" file.

   The **File System CBT (Changed Block Tracker)** component is not required. Install this component if you require file-based SQL backup.

4. The agent starts automatically.

Note:

- The installation wizard may contain references to **user certificates**. This is not currently supported and reserved for future releases.
- This installation does not require a restart. You can wait until a normal maintenance window (such as a Patch Tuesday) to restart. However, until you do, some Protection Groups that depend on Volume CBT will run in full backup mode.
- When installing the Cohesity agents, usage of **Port 21213** is reserved for Cohesity's use with IANA. This is an Early-Access feature. Contact your Cohesity account team to enable the feature.

Silently Install the Cohesity Agent on Windows Physical Servers

You can install the Windows Cohesity agent on multiple Windows servers by running the following command in an Administrator CMD shell or Administrator PowerShell. (The servers must have access to the agent file.)

```
 <InstallerExeFile> [/type=<onlyagent|volcbt|fscbt|allcbt>]
[/log="<absolute-path-to-log-file>"] [/dir="<absolute-path-to-install-
location>"] [/username=<value> [/password=<value>]] [/verysilent]
[/suppressmsgboxes] [/norestart] [/firewallrule=<value>] [/configfile
<value>] /deleteconfigfile <value>] [/enforceusecustomcert=<value>]
[/cohesitycustomcertpath=<Value>]
```

The default name for `<InstallerExeFile>` is `Cohesity_Agent_<Version>_Win_x64_Installer.exe`.

A sample command is shown below:

```
Cohesity_Agent_5.0.0.0_Win_x64_Installer.exe /log="C:\logs\install logs.txt"
/verysilent /suppressmsgboxes /norestart /dir="F:\CohesityAgent"
```

The command parameters are described below.

Parameter	Description
`/type=<onlyagent|fscbt|allcbt|volcbt|onlyagent|volcbtwithsbt|fscbtwithsbt|agentwithsbt>`	Specifies the type of installation. • `volcbt` - Installs the Windows Agent Service and Volume CBT component. • `fscbt` - Installs the Windows Agent Service and File System CBT component. This parameter enables incremental backups of Windows Server 2012 R2 or individual databases hosted by MS SQL Server. • `allcbt` - Installs the Windows Agent Service, Volume CBT component and File System CBT component. This parameter enables incremental backups of Windows Server 2012 R2 or individual databases hosted by MS SQL Server. • `onlyagent` - Installs only the Windows Agent Service and not the Volume CBT or File System CBT components. • `onlyagent`- Installs only the Windows Agent Service and not the Volume CBT or File System CBT components. • `volcbtwithsbt`- Agent, Volume CBT and Oracle SBT Library • `fscbtwithsbt`- Agent, File System CBT and Oracle SBT Library • `allcbtwithsbt`- Agent, File System CBT, Volume CBT and Oracle SBT Library • `agentwithsbt`- Agent and Oracle SBT Library **Note:** The `/type` parameter is supported only for a fresh installation. It is not supported for upgrades. You cannot add or remove components during an upgrade. **Note:** If the `volcbt` parameter is not specified, volume based backups will always be full. Similarly, if the `fscbt` parameter is not specified, SQL and exchange file based backups will be full.
`/log="<absolute-path-to-log-file>"`	Records the events during installation and helps to debug the errors, if any, during the installation process.

Parameter	Description
`/username=<value> [/password=<value>]`	The Cohesity Windows Agent Service is installed, by default, in the Local System account. Use these parameters to change the account under which the service will run. Note that the specified user account must have local administrator privileges on the server and should have the `Log on as a service` right.
`/verysilent`	Suppresses the installation progress window. **Note:** If you specify this option during the agent installation, you should review the default values for all the other parameters such as `/type`, `/dir`, `/username`, etc., and explicitly add them with relevant values if the defaults are not acceptable.
`/suppressmsgboxes`	Suppresses any message boxes if `/verysilent` is specified.
`/norestart`	Prevents the setup from restarting the system following a successful installation, or after a `Preparing to Install` failure that requests a restart.
`/dir="<absolute-path-to-install-location>"`	Specifies the directory path in which the Cohesity Windows Agent Service should be installed. **Note:** The installer will create the necessary folders if they do not exist. If this parameter is not specified, the system will install the Windows Agent Service in the default directory path: `"%SystemDrive%\Program Files\Cohesity"`.

Parameter	Description
`/firewallrule=<yes | no>`	Specifies whether or not to open Agent port in the firewall. The default value is 'yes'. **Note:** The installer creates an inbound rule in the local firewall of the Windows machine for port 50051 and is bound with the Cohesity Agent Service. Any domain policy settings that may negate this effect will have to be handled at the domain policy level. The installer creates rules only in the Windows Firewall application. If any other firewall application is in force on the machine then the user will have to create appropriate rules in the active firewall application.

`/configfile <value>`	Specifies the path to INI file containing credentials for Agent service user account.
	
	Note: If the `configfile` parameter is not specified, the username and password command line parameters are ignored.
	
	To pass the credentials in the config file, follow the steps below:
	
	Create a `config.ini` file with username and password and specify the same in `/configfile` option.
	
	Example
	
	`[credentials]`
	`username=<user_name>`
	`password=<password>`

`/deleteconfigfile <yes | no>`	Specifies whether to delete the configfile upon successful installation.
`[/enforceusecustomcert=<yes>]`	Enforce the user to use a custom certificate while installing the agent. If this is set to YES, then the subsequent fields should have valid values, else the agent will not start. If the custom certificate is not found then the agent would not start.

Parameter	Description
[/cohesitycustomcertpath=<path-to-agent-certificate-in-Cohesity-format>	To be used only when **/customcertstype = cohesitycustomcerts**. Specifies the directory path in which the Cohesity custom certificate for the agent is present.

> **Note:** On Windows physical servers, the Cohesity agent logs are stored in C:/programData/cohesity/.

Upgrade the Agent on VMware

Windows Agent

Cohesity recommends using the agent installer when installing the persistent agent. If the persistent agent is manually installed, use the following workaround to upgrade the agent:

1. Log on to the server from which the persistent agent must be upgraded.
2. Open an elevated command window, and then stop the earlier persistent agent service:

   ```
   net stop cohesityagent
   ```

3. Download the Cohesity agent installer from the Cohesity UI.
4. Double-click the installer .exe file to initiate the installation. Follow the installer UI prompts.
5. Do not choose to install any CBT drivers. Only upgrade the agent binary.

   This procedure upgrades the Cohesity agent and creates an entry in **Control Panel** > **Uninstall Program**.

VMware VMs that are not registered as a source on the Cohesity cluster need to upgrade the installed Cohesity agent manually.

1. Download the Cohesity agent installer from the Cohesity UI.
2. Double-click the installer .exe file to initiate the installation. Follow the installer UI prompts.
3. Do not choose to install any CBT drivers. Only upgrade the agent binary.

Download and Install the Agent on the SCVMM Server

Before registering SCVMM, you must install the Cohesity agent on the SCVMM server or on an existing proxy endpoint that is connected to an SCVMM server.

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges, run the executable and complete the installation wizard.

   Install the agent without additional components.
4. The agent starts automatically. Now you can register the Hyper-V SCVMM server as a source.

Download and Install the Agent on Hyper-V Hosts

Ensure you install the Cohesity agent on any Hyper-V hosts that you want Cohesity to protect.

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate hosts.
3. As an administrator with local system privileges, run the executable and complete the installation wizard.

   For Hyper-V 2016 servers, install the agent without additional components.

   For incremental backups of individual VMs hosted by Hyper-V on Windows Server 2012 R2, install the **File System CBT (Changed Block Tracker)** component.
4. The agent starts automatically.

Guest Windows VM minimum recommended specifications: 2 GB RAM and the equivalent of a 1 GHz processor.

Best practice: For Hyper-V 2016, configure all VMs' Automatic Stop Action to shut down or turn off, instead of save. This results in all powered-on VMs having minimal size .vmrs files. Because Cohesity does not back up .vmrs files that are greater than 10 MB (.vmrs file size, if set to save, scales to the amount of memory assigned to the VM). VMs in the saved state, by definition, generally have .vmrs files greater than 10 MB. This Cohesity Windows Agent registry value entry determines the maximum size of the .vmrs file supported per backup: agent_hyperv_max_vm_config_file_size_mb. Set it to zero to back up .vmrs files of all sizes.

Upgrade the Agent on SCVMM and Hyper-V Hosts

You can use the Cohesity UI to upgrade agents on a non-SCVMM cluster.

1. Select **Data Protection > Sources**.
2. Click the Hyper-V source name.

3. If an upgrade is available, click **Upgrade Agents**.

   This upgrades eligible agents, such as a singleton SCVMM node or a proxy SCVMM.

To upgrade an agent individually, mouse over the corresponding actions menu and click **Upgrade Agent** ( ).

Agents on hosts that belong to an SCVMM cluster must be updated manually. Download a new Cohesity agent and run the installer.

Download and Install the Agent on Nutanix AHV

File or folder-level recovery of VMs in a Nutanix Acropolis Hypervisor requires a Linux or Windows Cohesity agent installed on the target VM. Follow these steps to install the Cohesity agent.

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**. Ensure the agent has been downloaded to the appropriate server.
3. As an administrator with local system privileges on that server, run the executable and complete the installation wizard. Do not install any CBT components.

   If you use the Windows command prompt to install the agent, you can use a command similar to this sample:

   ```
   Cohesity_Agent_6.3.0.0_Win_x64_Installer.exe /type=onlyagent
   ```

   For more information about command parameters, see Silently Install the Cohesity Agent on Windows Physical Servers.
4. The agent starts automatically.
5. Configure the firewall and ports as necessary. For more information, see Manage Firewall Ports.

Upgrade the Windows Agent on a Physical Server

> **Note:** Scheduled upgrades of agents are supported on the Cohesity cluster.

You can use the Cohesity UI to upgrade agents on a physical server.

1. Navigate to **Data Protection** > **Sources**.
2. If an upgrade is available, click **Upgrade Agents**.

To upgrade an agent individually, mouse over the corresponding actions menu (⋮) and click **Upgrade Agent**.

Auto-upgrade of the Cohesity Windows Agent

Cohesity upgrade process supports auto-upgrade of the agents running on sources.

> **Note:** The option to auto-upgrade Cohesity Agent is not available on the Cohesity cluster with multi-tenancy configured. The auto-upgrade option is available only during cluster upgrade.

To enable auto-upgrade:

1. In the Cohesity UI, navigate to **Settings** > **Update**, and select the **Upgrade** tab.
2. Select the **Upgrade Cohesity Agent on all servers in Sources** checkbox to upgrade the agents on the source automatically after the cluster upgrade.

Uninstall the Agent

The following uninstallation methods are available:

- Use the Add/Remove Programs in the Control Panel and Start Menu shortcuts in Windows versions where it is supported.
- Use the uninstaller executable. The details of the uninstaller executable are available in the registry value "HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\{99246986-0300-4B21-9A1C-DE278D8EC715}_is1\UninstallString". The default value is "C:\Program Files\Cohesity\Installer\unins000.exe", assuming C: is the default operating system installation volume and "C:\Program Files\Cohesity" is the default Cohesity Agent install location.

Preserve registry settings when uninstalling the agent

Uninstalling the Cohesity agent normally clears registration settings from the registry. You can preserve registration settings in the registry by using the `/preservesettings` flag when uninstalling from the command line. Preserving registry settings is not supported when uninstalling from the UI.

Sample command:

```
unins000.exe /preservesettings
```

Install and Manage the Agent on Linux Servers

The Cohesity Linux agent is now available with different installer packages, providing support on multiple Linux distributions. The following table lists the installer packages and Linux distributions on which the installer package is supported.

Installer package	Linux distribution
(Default) RPM	RHEL and its derivative
SUSE RPM	SUSE
Script Installer	All supported Linux operating systems

Linux agent has dependencies on the following packages, which must be installed on the Linux server:

Command/Package	Linux distribution				
	RHEL	SUSE	CentOS	Ubuntu	Debian
rsync	rsync	rsync	rsync	rsync	rsync
mount	mount	mount	mount	mount	mount
lsof	lsof	lsof	lsof	lsof	lsof
umount	umount	umount	umount	umount	umount
cp	cp	cp	cp	cp	cp
chown	chown	chown	chown	chown	chown
chmod	chmod	chmod	chmod	chmod	chmod
mkdir	mkdir	mkdir	mkdir	mkdir	mkdir
rm	rm	rm	rm	rm	rm
tee	tee	tee	tee	tee	tee
hostname	hostname	hostname	hostname	hostname	hostname
stat	stat	stat	stat	stat	stat
blkid	blkid	blkid	blkid	blkid	blkid
ls	ls	ls	ls	ls	ls

Command/Package	Linux distribution				
	RHEL	SUSE	CentOS	Ubuntu	Debian
losetup	losetup	losetup	losetup	losetup	losetup
dmsetup	dmsetup	dmsetup	dmsetup	dmsetup	dmsetup
timeout	timeout	timeout	timeout	timeout	timeout
lvs	lvs	lvs	lvs	lvs	lvs
vgs	vgs	vgs	vgs	vgs	vgs
lvcreate	lvcreate	lvcreate	lvcreate	lvcreate	lvcreate
lvremove	lvremove	lvremove	lvremove	lvremove	lvremove
lvchange	lvchange	lvchange	lvchange	lvchange	lvchange
libpcap-progs	null	libpcap-progs	null	null	null
nfs-utils **Note:** Required for Instant Volume Mount, file-folder recovery from block-based backup and VMware backup.	nfs-utils	nfs-utils	nfs-utils	nfs-utils	nfs-utils
wget	wget	wget	wget	wget	wget

Installing the Linux agent as an LDAP user is supported.

Free Space Requirement for Volume-based Backup

For block-based backup of Linux physical server, the space must be at least equal to the calculation (Initial disk size*Churn rate*Backup window time). If you're running multiple

backups on a cluster, you'll need enough storage capacity to accommodate the snapshots for each of the machines using this calculation.

You can run the `vgdisplay` command on the physical server to check the free space in a Volume Group in LVM.

If the volume-based backup fails due to insufficient free space in a Volume Group, then the following error message is displayed on the Cohesity UI.

```
Volume group "rhel" has insufficient free space (0 extents): 1 required.
```

Download and Install the Linux Agent

Installing the agent on a server allows actions such as registering the agent as a source with the Cohesity cluster or performing file-level recovery.

Install the Agent

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**.

   The Download Agents window appears.
3. Follow the installation instructions for the corresponding installer package.
4. Configure the firewall and ports as necessary. For more information, see Manage Firewall Ports.

Install RPM, Debian or SUSE RPM Installer Package

The following table explains the action to be performed for installing the RPM, Debian or SUSE installer package:

> **Note:** In the following commands, replace the Cohesity agent version with the desired agent version.

Action	How to
Download the Agent installer	From the Download Agents window, based on your Linux distribution, select **RPM**, **Debian**, or **SUSE RPM** and download it to the server you want to protect.
Navigate to the downloaded directory	As the root user with local system privileges on that server, change the directory to the location of the installer package.

Action	How to
Install the Custom certificates	**LINUX** If you are using a Cohesity certificate then run the following commands to install the user certificate followed by agent installation: `export ENFORCE_USE_CUSTOM_CERTS=true && export AGENT_CERT_FILE=<path-to-agent-certificate-in-Cohesity-format>`
Install the Agent	Run the following command depending on the installer package: • RPM: `rpm -i el-cohesity-agent-7.4.1-1.x86_64.rpm` or `yum localinstall ./el-cohesityagent-7.4.1-1.x86_64.rpm` • Debian: `dpkg -i cohesity-agent_7.4.1-1_amd64.deb` • SUSE RPM: `rpm -i cohesity-agent-7.4.1-1.x86_64.rpm`
Install the agent as root user	Run the executable file with the following command syntax: `/<path_to_installer_file> -- --install -c 0 -S root -G root` Followings are the command options in details: 1. **-S** this is the user to run the agent, specify root. 2. **-G** this is the group permission the Agent will use for files and directories installed by the agent, specify root. 3. **-c** this boolean switch controls if the OS user and group should be created. 0 means do not create the OS user and group, 1 means the agent installation will create the specified OS user and group. (When choose to run with the root user, specify '-c 0' as 'root' already exists)

Action	How to

Install the agent as non-root user	To start the service as a non-root user, create a new user or use an existing user with sudo permission and run the following command:
	• RPM: `export COHESITYUSER=<username> ; rpm -i el-cohesity-agent-7.4.1-1.x86_64`
	• Debian: `COHESITYUSER=<username> dpkg -i cohesity-agent_7.4.1-1_amd64`
	• SUSE RPM: `export COHESITYUSER=<username> rpm -i cohesity-agent-7.4.1-1.x86_64`
	Where, *<username>* is the user name with sudo permission.
	Note:
	By default, the installation uses the root user permission for all the files, and the service is started as root. Therefore, it is required to add non-root users to the sudoers list by making the following changes in the `/etc/sudoers` file:
	`<user_name> ALL=(ALL) NOPASSWD:ALL`
	`Defaults:<user_name> !requiretty`
	An example of the sudoers file is as given below:
	dgoble ALL=(ALL) NOPASSWD:ALL
	Defaults:dgoble !requiretty
	cohbackup ALL=(ALL) NOPASSWD:ALL
	Defaults:cohbackup !requiretty

Action	How to

	Note:
	You can add sudo privileges for non-root users using the syntax given below:
	`<user_name> ALL=(ALL) NOPASSWD:/path/to/command1, /path/to/command2, /path/to/command3`
	Remove `NOPASSWD` from the command above if you do not want to authenticate every sudo commands with password.
	The example below shows the sudo commands for user cohesityagent:
	`cohesityagent ALL=NOPASSWD: /bin/mount,/bin/umount,/bin/chmod,/bin/chown,/bin/hostname, /bin/mkdir,/bin/rm,/usr/bin/tee,/usr/bin/stat,/bin/cp,/bin/sed,/usr/bin/touch,/sbin/blkid,/usr/bin/lsof,/bin/ls,/sbin/losetup,/sbin/dmsetup,/usr/bin/rsync,/usr/bin/timeout,/usr/bin/getfacl,/usr/bin/setfacl,/tmp/installer_ch_cohesity_linux_ agent,/home/cohesityagent/installer,/home/cohesityagent/cohesityagent/data/agent/tmp/installer_ch_cohesity_linux_ agent,/bin/systemctl,/home/cohesityagent/cohesityagent/software/crux/bin/cohesity_linux_agent.sh`
	The actual paths for each command may vary in each system. Cohesity keeps adding or removing commands across releases.

Location	• Installation directory: `/opt/cohesity` • Log file: `/var/log/cohesity`

Install script installer package

The following table explains the action to be performed for installing the script installer package:

> **Note:** In the following commands, replace the Cohesity agent version with the desired agent version.

Action	How to
Download the Agent installer	From the Download Agents window, based on your Linux distribution, select **Script Installer** and download it to the server you want to protect.

Action	How to
Navigate to the downloaded directory	As the root user with local system privileges on that server, change the directory to the location of the installer package. For SLES 11 SP4, installing the agent as the root user is required.
Make the installer executable	Make the installer executable, for example: `chmod +x cohesity_agent_7.4.1_linux_x64_installer`
Install the Agent	Run the executable: `sudo /cohesity_agent_6.6.0d_u2_linux_x64_installer - --install`
Location	• Installation directory:`/home/<username>/cohesityagent` or `/root/cohesityagent` • Log file: `/home/cohesityagent/cohesityagent/logs` **Note:** You can specify the option `-log-dir <installtion_path>` to specify the directory path where the agent logs will be saved. If this option is not provided, the logs will be saved in the default path

Install Cohesity Agent for Linux PowerPC

> **Note:** To use the Linux PowerPC agent (Java based), Java Runtime Environment (JRE) 8 needs to be installed on the source system.

The following section explains the action to be performed for installing the Linux PowerPC Agent as a ROOT user:

1. Select **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent**.
3. From the **Download Agents** window, select **PowerPC** and download it to the server you want to protect.
4. Change the directory to the location of the installer package and run the following command to install the rpm based installer package for PowerPC distribution:

   ```
   cohesity-java-agent-7.0.1-1.ppc64le.rpm
   ```

   The agent will be installed in the following directory `/opt/cohesity` and the logs files can be found in `/var/log/cohesity`.

5. To uninstall the agent, run the following command:

```
rpm -e cohesity-java-agent
```

For information about all Cohesity Linux agent installer options, see Linux Script Installer Options.

The agent starts after installation completes.

Agent startup behavior following a machine reboot:

- CentOS and RedHat (distributions with the "systemd" init system) - The agent starts automatically.
- Ubuntu (distributions with the "upstart" init system) - The agent starts automatically.

If a Linux server's `/etc/sudoers` file is managed by a deployment engine such as Chef, Puppet or others, this could impact Cohesity cluster's interaction with servers that have the Linux agent installed. Take the corresponding actions below:

Agent Installation - User type	Action required

As default cohesityagent user	The Cohesity Linux agent is installed using the cohesityagent user by default. For default installations, the cohesityagent is created by the installer. During installation, the installer updates the `/etc/sudoers` file to allow cohesityagent sudo and no-tty sudo access. Ensure the following settings in the /etc/sudoers file for the cohesityagent user are preserved:
	`cohesityagent ALL=(ALL) NOPASSWD:ALL`
	`Defaults:cohesityagent !requiretty`
	An example of the sudoers file is as given below:
	#includedir /etc/sudoers.d
	dgoble ALL=(ALL) NOPASSWD:ALL
	cohbackup ALL=(ALL) NOPASSWD:ALL
	Defaults:cohbackup !requiretty

As non default user, example foo	Ensure the above settings in the `/etc/sudoers` file for the foo user are preserved by replacing the occurrences of cohesityagent with foo.
As root user	No changes required

Install Cohesity Agent for Linux on z/Architecture

> **Note:** To use the Linux on z/Architecture agent (Java-based), Java Runtime Environment (JRE) 8+ (JRE 8, 11, 17) needs to be installed on the source system.

The following section explains the action to be performed for installing the Linux on z/Architecture Agent as a ROOT user:

1. Select **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent**.
3. From the **Download Agents** window, select **s390x** and download it to the server you want to protect.
4. Change the directory to the location of the installer package and run the following command to install the rpm based installer package for z/Architecture distribution:

   ```
   rpm -i cohesity-java-agent-0.0.0-1.s390x.rpm
   ```

   The agent will be installed in the following directory `/opt/cohesity` and the logs files can be found in `/var/log/cohesity`.
5. To uninstall the agent, run the following command:

   ```
   rpm -e cohesity-java-agent
   ```

Install Cohesity Agent for Linux on on ARM (aarch64)

> **Note:** To use the Linux on on ARM (aarch64) agent (Java-based), Java Runtime Environment (JRE) 8+ (JRE 8, 11, 17) needs to be installed on the source system.

The following section explains the action to be performed for installing the Linux on on ARM (aarch64) Agent as a ROOT user:

1. Select **Data Protection** > **Sources**.
2. Click **Download Cohesity Agent**.
3. From the **Download Agents** window, select **aarch64** and download it to the server you want to protect.
4. Change the directory to the location of the installer package and run the following command to install the rpm based installer package for z/Architecture distribution:

   ```
   rpm -i cohesity-java-agent-7.3.1-1.aarch64.rpm
   ```

   The agent will be installed in the following directory `/opt/cohesity` and the logs files can be found in `/var/log/cohesity`.
5. To uninstall the agent, run the following command:

   ```
   rpm -e cohesity-java-agent
   ```

cohesity_linux_agent.sh

The table lists how `cohesity_linux_agent.sh` interacts with various `init` systems to issue `start/stop/status` commands.

If Cohesity does not detect any of the `init` systems, the installer uses the **undetected** `init` system commands.

init system	Commands

systemd	`sudo systemctl start cohesity-agent`
	`sudo systemctl stop cohesity-agent`
	`sudo systemctl status cohesity-agent`

upstart	`sudo initctl start cohesity-agent`
	`sudo initctl stop cohesity-agent`
	`sudo initctl status cohesity-agent`

SysV	`sudo /etc/init.d/cohesity-agent start`
	`sudo /etc/init.d/cohesity-agent stop`
	`sudo /etc/init.d/cohesity-agent status`

undetected	`sudo /etc/init.d/cohesity-agent start`
	`sudo /etc/init.d/cohesity-agent stop`
	`sudo /etc/init.d/cohesity-agent status`

Download and Install the Cohesity Agent on RHV/OLVM Hosts

> **Note:** Cohesity supports RHV only until version 6.8.1. While you can still protect your RHV sources, the Cohesity Support team will not provide assistance.

Ensure you install the Cohesity Agent on any RHV/OLVM hosts that you want Cohesity to protect. For installing the agent on RHV/OLVM hosts, follow the steps described in the Download and Install the Cohesity Agent on Linux Physical Servers" section based on the Linux distribution

Run the following command on all the RHV/OLVM hosts after they are connected and configured by the oVirt Manager:

```
# iptables -I INPUT `iptables -L INPUT --line-numbers | grep REJECT | awk '{
print $1 }'` -p tcp --dport 50051 -j ACCEPT

# service iptables save
# service iptables restart
```

If the firewall is enabled on RHV 4.1 hosts, open the following port 50051 to accept traffic.

To open the ports, run the following command:

firewall-cmd --permanent --add-port 50051/tcp

Linux Script Installer Options

The Cohesity Linux agent installer has the following options available.

Option	Description
`--agent-options, -a [options]`	Options to pass to the Cohesity agent. Used with single mode operation only.
`--create-user, -c [0|1]`	Whether to create `--service-user` or not. [Default: 1] Used with `--install` option only.
`--debug, -d`	Enable shell script debugging with set `-x` [Default: off]
`--full-uninstall, -U`	Uninstall the Cohesity Linux agent including service-user account and its home directory contents. Destructive operation. Use with care.
`--help, -h`	Print this help.
`--install, -i`	Install/uninstall the Cohesity Linux agent.
`--install-dir, -I [dir]`	Cohesity agent installation dir. [Default: `/home/<username>/cohesityagent` or `/root/cohesityagent`]
`--install-log-file, -L [file]`	Filename where installer logs should be saved. [Default: */tmp/cohesity_agent_installer_<time_in_seconds>_<pid>.log*]
`--log-dir`	Cohesity agent log installation directory. [Default: `/home/cohesityagent/cohesityagent/logs`]
`--service-group, -G [group]`	Groupname to use for service. [Default: `cohesityagent`]

Option	Description
`--service-user, -S [user]`	Username to use for service. [Default: `cohesityagent`]
`--single-mode, -s`	Install the Cohesity Linux agent for single mode operation. This installs the agent for one time usage/run only. It does not install the Agent to run as service or daemon to run continuously. This copies all required binaries/files to a location and runs the agent. Example use case: You want to recover files to a Linux VM. Cohesity pushes the Cohesity agent and starts it. After performing the usual file/folder recovery by talking to Cohesity Agent and file copying is done, the Agent is removed. The agent is not installed to run as service. It runs for the duration of the recover task only.
`--skip-lvm-check`	**1** - Skip LVM related checks during installation and start the agent in non-lvm mode (only file-based backup will be supported). **0** - Flag is not set. > **Note:** The default value is **0**.
`--skip-mountpoint-check`	Skip checks performed on installation directory's filesystem mountpoint.
`--target`	Extract directly to a target directory. Directory path can be either absolute or relative. [Default: `/tmp`]
`--update-uninstall, -u`	Uninstall the Cohesity Linux agent. Preserve user, group and agent config. Useful when upgrading software.
`--yes-all, -y`	Accept 'yes' answer to all questions in installer. [Default: off] This is useful for performing silent install/uninstall.

Install the Cohesity Agent to run with the Root User

Follow the steps below to install the Cohesity agent to run with the root user:

> **Note:** In the following command, replace the Cohesity agent version with the desired agent version.

For example, run the following command: `sudo ./cohesity_agent_7.4.1_linux_x64_installer -- --install -S root -G root -c 0 -I /opt -y`

The above example does the following:

- `-S root -G root` sets user "root:root" user:group for the agent.
- `-c 0` instructs the installer that user `root` already exists on the system, so do not create it.
- `-I /opt` installs the agent in the /opt directory. If you do not specify the directory, then the default home directory of the user will be used, /root/ in this case.

Sudo Privileges for Linux Users

The Linux system administrator can allowlist set of commands that do not need the root rights to be executed from the Cohesity agent deployed on the Linux servers.

The following tables list the commands that can be allowlisted for Oracle database, volume backups, and file restores for the Cohesity agent.

Oracle Database

Argument	Default Behavior	Allowlist Commands

`--skip-oracle-check`	The default value is 0. We will allowlist commands required by Oracle database.	`mkdir`
		`rm`
		`tee`
		`stat`
		`cp`
		`srvctl (from Oracle home)`

	If set to 1, we will not add any commands to the allowlist.	

Volume Backup

Argument	Default Behavior	Allowlist Commands

`--skip-lvm-check`	The default value is 0. We will allowlist commands required by volume backup.	`lvs`
		`vgs`
		`lvcreate`
		`vgcreate`

	If set to 1, we will not add any commands to the allowlist.	

File Restore

Argument	Default Behavior	Allowlist Commands

`--skip-restore-check`	The default value is 0. We will allowlist commands required by file restore.	`blkid`
		`lsof`
		`ls`
		`losetup`
		`dmsetup`
		`rsync`
		`timeout`

	If set to 1, we will not add any commands to the allowlist.	

> **Note:** During a Linux agent upgrade:
> If **NOPASSWD:ALL** entry is present in the sudoers file, then all commands can be executed without the root rights.
> If **NOPASSWD:ALL** entry is removed from the sudoers file, then you can add the commands listed in the table above for allowlisting.

Upgrade the Linux Agent

You can upgrade the Linux agent either from the Cohesity UI or using the CLI.

> **Note:** In the following commands, replace the Cohesity agent version with the desired agent version.

Upgrade from the Cohesity UI	Upgrade using CLI

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**.
3. Stop the Agent service on the server.
4. Click **Upgrade Agents**.
5. Restart the Agent service.
	1. Navigate to the Agent installation directory.
	2. Run the following command to install the script installer package:
	`sudo cohesity_agent_7.4.1_linux_x64_installer -- --install`
	3. Run the following command to install the RPM package:
	`rpm -U el-cohesity-agent-7.4.1-1.x86_64.rpm`
	4. Run the following command to install the Debian package:
	`dpkg -i cohesity-agent_7.4.1-1_amd64.deb`
	5. Run the following command to install the SUSE RPM package:
	`rpm -U cohesity-agent-7.4.1-1.x86_64.rpm`
	6. Run the following command to install the PowerPC RPM package:
	`rpm -U cohesity-java-agent-7.0.1-1.ppc64le.rpm`
	7. To install the Linux on z/Architecture package on Linux variants running on the x64 platform:
	`rpm -U cohesity-java-agent-0.0.0-1.s390x.rpm`
	8. To install the Linux on ARM (aarch64) package on Linux variants running on the x64 platform:
	`rpm -U cohesity-java-agent-7.3.1-1.aarch64.rpm`

Auto-upgrade of the Cohesity Linux Agent

Cohesity supports auto-upgrade of the Cohesity Linux Agent running on physical servers.

> **Note:** The option to auto-upgrade Cohesity Agent is not available on the Cohesity cluster with multi-tenancy configured.

To enable auto-upgrade:

1. Navigate to **Settings** > **Update**, and select **Upgrade** tab.
2. Select the **Upgrade Cohesity Agent on all servers in Sources** checkbox.

Uninstall the Script-Based Agent

To uninstall script-based agent, follow one of the methods listed below:

- If the Cohesity agent is installed in the default installation directory, run the "uninstall" script that is included with the installed software. For example:

```
cd <bin_directory>
[default=/home/cohesityagent/cohesityagent/software/crux/bin]
sudo ./uninstall
```

- Use the agent installer itself. For example:

```
sudo ./cohesity_agent_7.4.1_linux_x64_installer -- --full-uninstall
```

> **Note:** Replace the Cohesity agent version with the desired agent version.

Uninstall RPM/PKG Agents

To uninstall the RPM/PKG agents, follow the steps listed below:

- To uninstall the SUSE RPM/RPM package, run the following command:

```
rpm -e cohesity-agent
```

- To uninstall the Debian package, run the following command:

```
dpkg -r cohesity-agent
```

- To uninstall the PowerPC package, run the following command:

```
rpm -e cohesity-java-agent
```

- To uninstall the PowerPC Linux on z/Architecture package, run the following command:

```
rpm -e cohesity-java-agent
```

- To uninstall the PowerPC Linux on ARM (aarch64) package, run the following command:

```
rpm -e cohesity-java-agent
```

Sample Ubuntu Linux Installation Output

Sample output for the following command:

```
sudo ./cohesity_linux_agent_installer -- -i -y
```

```
Verifying archive integrity... All good.
```

```
Uncompressing Cohesity Linux Agent  100%
```

```
==> User asked to accept 'yes' answer to all questions in installer.
```

```
------------------------------------------------------------------------------
--
==> UserId = uid=0(root) gid=0(root) groups=0(root)
==> Umask  = 0077
==> Setting umask to 022
==> Umask  = 0022
------------------------------------------------------------------------------
--
==> SUDO_COMMAND_PREFIX=
/sbin/init: unrecognized option '--version'
==> INIT system detected: systemd
==> Checking if user has provided overrides from command line...
==> Service user : cohesityagent
==> Finding service user group based user input ...
==> Service user group : cohesityagent
==> Finding installation directory ...
==> Installation directory for agent: /home/cohesityagent/cohesityagent
==> Generating set_env.sh ...
==> Sourcing /tmp/selfgz30348/set_env.sh
==> Cohesity Linux Agent Software Version: 5.0.1
==> Installed Cohesity Linux Agent will be run with user=cohesityagent.
------------------------------------------------------------------------------
--
==> Installing Cohesity Linux Agent version 5.0.1
------------------------------------------------------------------------------
--
COHESITY_USER=cohesityagent
COHESITY_GROUP=cohesityagent
COHESITY_USER_CREATED=1
COHESITY_HOME_DIR=/home/cohesityagent
COHESITY_INSTALL_DIR=/home/cohesityagent/cohesityagent
```

```
------------------------------------------------------------------------------
--
==> User has given '--yes-all' option. Using that for next question...
Do you want to continue installing? (y/n) yes
==> Continuing installation...
------------------------------------------------------------------------------
--
==> This is x86_64 system
==> LVM: Using lvs from path: /sbin/lvs
==> LVM: 2.02.133(2)(2015-10-30)
==> Rsync: Using rsync from path: /usr/bin/rsync
==> Rsync: rsync  version 3.1.1  protocol version 31
==> Installing on Linux Distribution: "Ubuntu 16.04.2 LTS"
==> No previous Cohesity Agent is detected on this system.
==> Agent process linux_agent_exec is not running.
==> Checking mountpoint for install directory: /home/-
cohesityagent/cohesityagent
==> Parent of installation directory (existing): /home
==> Mountpoint of parent installation directory: /
==> Checking if mountpoint / has required option "rw" set.
(rw,relatime,errors=remount-ro,data=ordered)
==> Checking if mountpoint / has unsupported option "nosuid" set.
==> Checking if mountpoint / has unsupported option "noexec" set.
==> User=cohesityagent does not exist in /etc/passwd
==> Checking if User=cohesityagent exists using getent command.
==> User=cohesityagent not found with getend command.
==> Group=cohesityagent does not exist in /etc/group
==> Group=cohesityagent not found with getent command.
id: ‘cohesityagent’: no such user
==> Group=cohesityagent is not found with id command
==> Adding cohesityagent to sudoers list
```

```
==> Directory structure created under /home/cohesityagent/cohesityagent

==> Copied software contents to /home/cohesityagent/cohesityagent

==> Old installation directory is empty. Skipping copying data from older
installation.

/home/cohesityagent/cohesityagent/software/crux/bin/linux_agent_exec = cap_
dac_override,cap_sys_resource+ep

Synchronizing state of cohesity-agent.service with SysV init with /lib/sys-
temd/systemd-sysv-install...

Executing /lib/systemd/systemd-sysv-install enable cohesity-agent

insserv: warning: current start runlevel(s) (empty) of script `cohesity-
agent' overrides LSB defaults (2 3 4 5).

insserv: warning: current stop runlevel(s) (empty) of script `cohesity-
agent' overrides LSB defaults (2 3 4 5).

Created symlink from /etc/systemd/system/multi-user.target.wants/cohesity-
agent.service to /usr/lib/systemd/system/cohesity-agent.service.

==> Service cohesity-agent added to systemd

==> Starting cohesity-agent...

● cohesity-agent.service - Cohesity Linux Agent

Loaded: loaded (/usr/lib/systemd/system/cohesity-agent.service; enabled;
vendor preset: enabled)

Active: active (running) since Mon 2018-01-29 23:22:01 PST; 5ms ago

Process: 30796 ExecStart=/etc/init.d/cohesity-agent start (code=exited,
status=0/SUCCESS)

Process: 30793 ExecStartPre=/bin/chown -R cohesityagent:cohesityagent
/var/run/cohesity-agent (code=exited, status=0/SUCCESS)

Process: 30788 ExecStartPre=/bin/mkdir -p /var/run/cohesity-agent (code-
e=exited, status=0/SUCCESS)

Main PID: 30927 (linux_agent_exe)

Tasks: 10

Memory: 280.0K

CPU: 142ms

CGroup: /system.slice/cohesity-agent.service
```

```
‣ 30927 /home/cohesityagent/cohesityagent/software/crux/bin/linux_agent_exec
--log_dir=/home/cohesityagent/cohesityagent/data/logs --max_log_size=30 --
stop_logging_if_full_disk=true --logbuflevel=-1 --linux_agent_nfs_moun-
tpoint_base_dir=/home/cohesityagent/cohesityagent/data/agent//tmp/nfs_mounts
--linux_agent_volume_mountpoint_base_dir-
r=/home/cohesityagent/cohesityagent/data/agent//tmp/vol_mounts --linux_
agent_config_file_path-
h=/home/cohesityagent/cohesityagent/data/config/cohesity_agent.cfg --util_
misc_mount_command=/bin/mount --util_misc_umount_command=/bin/umount --util_
misc_findmnt_command=/bin/findmnt --util_misc_sudo_command=/usr/bin/sudo --
yoda_util_vmware_mount_com-
mand=/home/cohesityagent/cohesityagent/software/crux/bin/vm_mounter.sh --
yoda_mount_command=/bin/mount --yoda_umount_command=/bin/umount --yoda_
findmnt_command=/bin/findmnt --yoda_sudo_command=/usr/bin/sudo --timeout_com-
mand=/usr/bin/timeout --yoda_blkid_command=/sbin/blkid --yoda_lsof_com-
mand=/usr/bin/lsof --yoda_ls_command=/bin/ls --linux_sudo_
command=/usr/bin/sudo --linux_chmod_command=/bin/chmod --linux_timeout_com-
mand=/usr/bin/timeout --linux_lvs_command=/sbin/lvs --linux_vgs_com-
mand=/sbin/vgs --linux_lvcreate_command=/sbin/lvcreate --linux_lvremove_
command=/sbin/lvremove --linux_agent_rsync_cmd=/usr/bin/rsync --linux_agent_
wget_command=/usr/bin/wget --yoda_losetup_command=/sbin/losetup --yoda_
dmsetup_command=/sbin/dmsetup --linux_dump_memory_usage=false --yoda_bindfs_
mount=false --libssl_non_fips_path= --libssl_path= -v=0

Jan 29 23:21:59 ubuntu16-server cohesity-agent[30796]: Starting linux_agent_
exec...

Jan 29 23:21:59 ubuntu16-server cohesity-agent[30796]: GNU coreutils version
= 8.25

Jan 29 23:21:59 ubuntu16-server cohesity-agent[30796]: Setting coredump
limit to unlimited for qa build

Jan 29 23:21:59 ubuntu16-server cohesity-agent[30796]: Init system is sys-
temd, will run linux_agent_exec in background

Jan 29 23:21:59 ubuntu16-server su[30918]: Successful su for cohesityagent
by root

Jan 29 23:21:59 ubuntu16-server su[30918]: + ??? root:cohesityagent

Jan 29 23:22:01 ubuntu16-server cohesity-agent[30796]: Parent process pid-
d=30927

Jan 29 23:22:01 ubuntu16-server cohesity-agent[30796]: 30927

Jan 29 23:22:01 ubuntu16-server systemd[1]: Started Cohesity Linux Agent.

Jan 29 23:22:01 ubuntu16-server systemd[1]: Started Cohesity Linux Agent.
```

```
--------------------------------------------------------------------------------
--
==> Cohesity Linux Agent is installed successfully.
--------------------------------------------------------------------------------
--
--------------------------------------------------------------------------------
--
Agent Help (init system: systemd)
--------------------------------------------------------------------------------
--
Agent start:
# sudo systemctl start cohesity-agent
Agent stop:
# sudo systemctl stop cohesity-agent
Agent status:
# sudo systemctl status cohesity-agent
--------------------------------------------------------------------------------
--
cohesity@ubuntu16-server:~/agent$ cohesity@ubuntu16-server:~/agent$ sudo
./cohesity_linux_agent_installer_qa -- -i -y
==> Umask  = 0022
--------------------------------------------------------------------------------
--
==> SUDO_COMMAND_PREFIX=
/sbin/init: unrecognized option '--version'
==> INIT system detected: systemd
==> Checking if user has provided overrides from command line...
==> Service user : cohesityagent
==> Finding service user group based user input ...
==> Service user group : cohesityagent
==> Finding installation directory ...
==> Installation directory for agent: /home/cohesityagent/cohesityagent
```

```
==> Generating set_env.sh ...
==> Sourcing /tmp/selfgz30348/set_env.sh
==> Cohesity Linux Agent Software Version: 5.0.1
==> Installed Cohesity Linux Agent will be run with user=cohesityagent.
------------------------------------------------------------------------------
--
==> Installing Cohesity Linux Agent version 5.0.1
------------------------------------------------------------------------------
--
COHESITY_USER=cohesityagent
COHESITY_GROUP=cohesityagent
COHESITY_USER_CREATED=1
COHESITY_HOME_DIR=/home/cohesityagent
COHESITY_INSTALL_DIR=/home/cohesityagent/cohesityagent
------------------------------------------------------------------------------
--
==> User has given '--yes-all' option. Using that for next question...
Do you want to continue installing? (y/n) yes
==> Continuing installation...
------------------------------------------------------------------------------
--
==> This is x86_64 system
==> LVM: Using lvs from path: /sbin/lvs
==> LVM: 2.02.133(2)(2015-10-30)
==> Rsync: Using rsync from path: /usr/bin/rsync
==> Rsync: rsync  version 3.1.1  protocol version 31
==> Installing on Linux Distribution: "Ubuntu 16.04.2 LTS"
==> No previous Cohesity Agent is detected on this system.
==> Agent process linux_agent_exec is not running.
==> Checking mountpoint for install directory: /home/-
cohesityagent/cohesityagent
==> Parent of installation directory (existing): /home
```

```
==> Mountpoint of parent installation directory: /

==> Checking if mountpoint / has required option "rw" set.

(rw,relatime,errors=remount-ro,data=ordered)

==> Checking if mountpoint / has unsupported option "nosuid" set.

==> Checking if mountpoint / has unsupported option "noexec" set.

==> User=cohesityagent does not exist in /etc/passwd

==> Checking if User=cohesityagent exists using getent command.

==> User=cohesityagent not found with getend command.

==> Group=cohesityagent does not exist in /etc/group

==> Group=cohesityagent not found with getent command.

id: ‘cohesityagent’: no such user

==> Group=cohesityagent is not found with id command

==> Adding cohesityagent to sudoers list

==> Directory structure created under /home/cohesityagent/cohesityagent

==> Copied software contents to /home/cohesityagent/co-bash: cohes-
ity@ubuntu16-server:~/agent$: No such file or directory

hesityagent

==> Old installation directory is empty. Skipping copying data from older
installation.

/home/cohesityagent/cohesityagent/software/crux/bin/linux_agent_exec = cap_
dac_override,cap_sys_resource+ep

Synchronizing state of cohesity-agent.service with SysV init with /lib/sys-
temd/systemd-sysv-install...

Executing /lib/systemd/systemd-sysv-install enable cohesity-agent

insserv: warning: current start runlevel(s) (empty) of script `cohesity-
agent' overrides LSB defaults (2 3 4 5).

insserv: warning: current stop runlevel(s) (empty) of script `cohesity-
agent' overrides LSB defaults (2 3 4 5).

Created symlink from /etc/systemd/system/multi-user.target.wants/cohesity-
agent.service to /usr/lib/systemd/system/cohesity-agent.service.

==> Service cohesity-agent added to systemd

==> Starting cohesity-agent...
```

```
● cohesity-agent.service - Cohesity Linux Agent
```

```
Loaded: loaded (/usr/lib/systemd/system/cohesity-agent.service; enabled;
vendor preset: enabled)
```

```
Active: active (running) since Mon 2018-01-29 23:22:01 PST; 5ms ago
```

Troubleshoot

For troubleshooting information, see Troubleshooting Cohesity Agent issues on Linux physical servers.

Log in to the Cohesity support site to see more knowledge base articles.

Install and Manage the Agent on Hyper-V Linux VMs

To allow file-level recovery to a guest Linux VM running on Hyper-V, the guest VM must have the following installed:

- Cohesity Linux Agent
- Linux Integration Services
- Ubuntu
- CentOS or RHEL

Limitations

- Efficient file-level recovery does not support recovering from multiple Linux volumes.

Free Space Requirement for block-based Backup

For each virtual machine backup, you'll need free space on the volume hosting the virtual hard disk files to allow Hyper-V enough room for differencing disks (AVHDs) during backup. The space must be at least equal to the calculation Initial disk size*Churn rate*Backup window time. If you're running multiple backups on a cluster, you'll need enough storage capacity to accommodate the AVHDs for each of the virtual machines using this calculation.

Cohesity Linux Agent

Follow these steps to download and install the agent.

1. Navigate to **Data Protection > Sources**.
2. Click **Download Cohesity Agent**.

   The Download Agents window appears.

3. Manually install the agent on the Linux VM you want to recover to.

   Follow the installation instructions in this topic: Install and Manage the Agent on Linux Servers.

Linux Integration Services

The Linux Integration Services package must be installed on the Linux VM. For more information, see Linux Integration Services. Without Linux Integration Services installed, the Linux VM hostname and IP addresses are not visible to SCVMM and the VM cannot be selected as a recovery target.

Ubuntu

Ensure the VM is running a supported release. For more information, see Supported Ubuntu Virtual Machines on Hyper-V.

Install these additional packages:

- hv-kvp-daemon-init
- linux-tools
- linux-cloud tools

1. The following example command installs the packages:

   ```
   sudo apt-get update && sudo apt-get install hv-kvp-daemon-init linux-tools-$(uname -r) linux-cloud-tools-$(uname -r) -y
   ```

2. After you install the packages, test the daemons by running the following commands:
   - hv_fcopy_daemon
   - hv_kvp_daemon
   - hv_set_ifconfig
   - hv_vss_daemon
   - hv_get_dns_info
3. Shutdown the VM (not just a reboot).

   ```
   sudo shutdown -h now
   ```

4. Power on the VM using PowerShell or SCVMM.

   The IP address should be visible in the Networking tab.

CentOS or RHEL

Ensure the VM is running a supported release. For more information, see Supported CentOS and Red Hat Enterprise Linux virtual Machines on Hyper-V.

> **Note:** Ensure that **Guest Services** is enabled on Hardware Configuration > Integration Services on the selected VM.

Install the hyperv-daemons package.

1. The following example command installs the package:

   ```
   sudo yum install hyperv-daemons
   ```

2. Enable hypervvsd and hypervkvpd by running the following commands:

   ```
   sudo systemctl enable hypervvssd
   sudo systemctl enable hypervkvpd
   ```

3. Reboot the VM.

   ```
   sudo shutdown -r now
   ```

   The IP address should be visible in the Networking tab.

Install and Manage the Cohesity Java Agent on AIX Servers

You can install the Cohesity Java Agent on AIX server to protect:

- AIX server as a physical server
- Oracle database running on AIX server

Install the Cohesity Java Agent on AIX Physical Server

Prerequisites

Ensure that the AIX physical server meets the following requirements:

- To use the AIX agent (Java based), Java Runtime Environment (JRE) 8 needs to be installed on the source system.
- AIX agent (Java based) being installed should be a higher version than the existing version of GoAgent.
- Install 64-bit version 8 of IBM-supported AIX compatible Java.
- 1.5 GB of free RAM (at run time) to account for JRE overhead.
- Sudo package must be installed.
- Latest version of Bash must be installed. For more information, see AIX Toolbox for Linux Applications.
- Java version for the Agent must be JRE 1.8.0 AIX ppc-64 bit.

Considerations

- Only file-based Protection Groups are supported.
- IVM (Instant Volume Mount) is not supported.
- Cohesity does not support source-side deduplication of AIX server.
- AIX clustering is not supported.

Install the Cohesity Java Agent

For information on supported AIX versions on Cohesity cluster, see Supported Software.

To install the Cohesity Java agent:

You can download the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Download the installer:
   a. Using the Cohesity cluster UI

      Navigate to **Data Protection** > **Sources** and click **Download Cohesity Agent**.

   b. Using the Cohesity Package Download site

      Download the Cohesity Java Agent package (a **.bff** file) from Cohesity Package Download. Ensure the Java Agent is on the AIX server you want to protect.
2. To Install the Custom certificates, follow the below steps:

   ```
   echo ENFORCE_USE_CUSTOM_CERTS=true >> /usr/local/cohesity/set_env.sh
   ```

   If you are using a Cohesity certificate then run the following commands to install the user certificate:

   ```
   echo AGENT_CERT_FILE=<path-to-agent-certificate-in-Cohesity-format> >>
   /usr/local/cohesity/set_env.sh
   ```

   > **Note:**
   > - ENFORCE_USE_CUSTOM_CERTS, if this parameter is set to true, then the agent will try to load the specified certificate and fail to start in case of an invalid file or files.
   > - AGENT_CERT_FILE specifies the path to the agent certificate file in the Cohesity format
3. Install the package:

   Log in to the AIX server as a root user and run the following command to install the package:

   ```
   installp -ad <path_to_downloaded_agent_package> all
   ```

For example:

```
installp -ad cohesity_agent_6.6_aix_java.bff all
```

The Cohesity Java Agent starts automatically. Now you can register the server with the Cohesity cluster. Reboot is not required. The Cohesity Java Agent starts after any future reboots.

Install the Cohesity Java Agent on the Oracle AIX Server

Prerequisites

Before you install the AIX Java agent on the server, make sure the following software is installed on the server:

- The latest Bash package. For more information, see AIX Toolbox for Linux Applications.
- Java version, JRE 1.8.0 AIX ppc-64 bit for the agent.

> **Note:** Before installing the AIX Java agent, Cohesity recommends logging in to SQL Plus and RMAN using appropriate credentials (OS Auth or User Based). If the login is successful, proceed with the installation. If not, obtain the required user credentials and verify that you can log in.

Considerations

The mount point used for Oracle DB cloning is not persistent across reboots. To resolve this issue, manually update the **/etc/filesystem** file using any editor and update the mount path information.

Install the Cohesity Java Agent

To install the Cohesity Java agent on Oracle AIX sever:

You can upgrade the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Follow steps 1 and 2 from Install the Cohesity Java Agent section to download the agent and install Custom certificates.
2. Install the package:

   Log in to the AIX server as an Oracle ROOT user and run the following command to install the package.

   ```
   sudo install_with_oracle_adapter_support=true installp -ad <path_to_
   downloaded_agent_package> all
   ```

   For example:

```
bash-4.2$ id
uid=204(oracle) gid=54321(oinstall) groups=54322(dba),54323(oper)

bash-4.2$ sudo install_with_oracle_adapter_support=true installp -ad
cohesity_agent_6.6_aix_java.bff all
```

The Cohesity Java Agent starts automatically. Now you can register the server with the Cohesity cluster. Reboot is not required. The Cohesity Java Agent starts after any future reboots.

For a non ROOT user, run the following command:

```
sudo install_with_oracle_adapter_support=true install_oracle_adapter_
user_name=<oracle-username> install_oracle_adapter_dba_group_name=<DBA-
GROUP-NAME> installp -ad <path to .bff> all
```

> **Note:**
> - You should set the AIX NFS block size to 512K after installing the AIX java agent on the server. For more information, contact Cohesity Support.
> - `/var/log/cohesity/agent` is the default location where logs are stored.
> - For a non ROOT user, user should have **SETENV** enabled for the installp command in the `/etc/sudoers`.

Upgrade the Cohesity Java Agent

Cohesity supports auto-upgrading the Cohesity Java Agent.

> **Note:** If the existing Cohesity agent version is 6.4 and earlier versions, to upgrade the agent you need to uninstall the agent and install the new Cohesity Java Agent. For Cohesity agent 6.4.1 and higher versions, you can upgrade the agent from the Cohesity UI or CLI.

To upgrade the Cohesity Java agent

You can upgrade the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Using the Cohesity cluster UI

   Navigate to **Data Protection** > **Sources**. Click the action icon (⋮) next to the registered source and then click **Upgrade Agents**, if an upgrade is available.

2. Upgrading the Cohesity agent installer using CLI

   a. Download the latest Cohesity Java Agent package (a .bff file). Ensure the Java Agent is on the AIX server you want to upgrade.

   b. Install the package:

      Log in to the AIX server as a root user and run the following command to upgrade the agent:

      ```
      installp -ac -d <path_to_downloaded_agent_package> all
      ```

      For Example:

      ```
      installp -ac -d cohesity_agent_6.5.1_aix_java.bff all
      ```

Uninstall the Cohesity Java Agent

Uninstallation stops the running Agent and removes the libraries and other artifacts that were created during installation.

To uninstall the Cohesity Java agent, run the following command:

```
installp -u cohesity_agent.rte
```

Cohesity Java Agent Command-line Options

To start, stop and to get the status and version of the Cohesity Java agent, run the following command:

```
/usr/local/cohesity/agent/aix_agent.sh start|stop|status|version
```

Install and Manage the Agent on Solaris Servers

Install the Cohesity Agent on Solaris Physical Servers

To install the Cohesity agent on the Solaris server:

You can download the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Download the installer:
   a. Using the Cohesity cluster UI

      Navigate to **Data Protection** > **Sources** and click **Download Cohesity Agent**.

   b. Using the Cohesity Package Download site

      Download the Solaris 10 Agent here.

      > **Note:** Solaris 10 and 11 agents can also be installed in non-global zones. We can register a non-global zone as a physical server and perform file-level backups or recoveries.

   c. Run the following commands to create a custom installer option file to install with custom certificates:

      ```
      echo ENFORCE_USE_CUSTOM_CERTS=true >> /usr/local/cohesity/set_env.sh
      ```

      > **Note:**
      > - ENFORCE_USE_CUSTOM_CERTS, if this parameter is set to true, then the agent will try to load the specified certificate and fail to start in case of an invalid file or files.
      > - AGENT_CERT_FILE specifies the path to the agent certificate file in the Cohesity format.

   d. Install the package:

      Run the following command to install the package:

      ```
      sudo /usr/sbin/pkgadd -d <cohesity_agent_file_path>
      ```

      If sudo is not installed, install the package as the root user.

      The agent starts automatically after installation and does not require server reboot. Now you can register the server with the Cohesity cluster. The agent will start automatically after every reboot of the server.

   e. Run the following command to know the agent version on the source:

      `/usr/local/cohesity/agent/solaris_agent.sh` version

Additional Agent

This section lists the additional CLI commands that can be used with the Solaris agent.

View Agent status:

```
/usr/bin/svcs -x svc:/application/cohesity-agent:default
```

Start the Agent:

```
/usr/sbin/svcadm enable svc:/application/cohesity-agent:default
```

Stop the Agent:

```
/usr/sbin/svcadm disable svc:/application/cohesity-agent:default
```

Upgrade the Agent

To upgrade the agent using the CLI, follow the steps below:

You can upgrade the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Upgrade the installer:
    a. Using the Cohesity cluster UI

       Navigate to **Data Protection** > **Sources**. Click the action icon (⋮) next to the registered source and then click **Upgrade Agents**, if an upgrade is available.

    b. Using the Cohesity Package Download site
        i. Download the latest Solaris 10 Agent here.
        ii. Run the following command to upgrade the package:

```
sudo /usr/sbin/pkgadd -d <cohesity_agent_file_path>
```

The Agent restarts automatically. Reboot is not required.

Uninstall the Agent

Uninstallation stops the running Agent and removes the libraries and other artifacts that were created during installation.

To uninstall agent, run the following command:

```
sudo /usr/sbin/pkgrm cohesity-agent
```

Install and Manage the Agent on HP-UX Servers

For information on supported HP-UX versions on the Cohesity cluster, see Physical Servers.

Prerequisites

Ensure that the HP-UX server meets the following requirements:

- **OS Version**: HP-UX 11i v3 (B.11.31)
- **Architecture**: Intel Itanium (IA64)
- **Cohesity Version**: Cohesity cluster and agent version must be version 6.8.1 or later.
- **Java Version**: Java 8 installed on HP-UX and it should be installed in either PATH or `/usr` or `/opt`.
- **Minimum Space Requirements**:
  - 200MB in /var and /usr directories
  - 100MB in /tmp directory
  - 10KB in /etc directory
- **Memory Requirements**: Minimum memory of 2GB is recommended.

Consideration

- The Cohesity software version and the Cohesity Agent version should be 6.8.1 or later for HP-UX servers.
- Cohesity supports only file-based backup and recovery on HP-UX servers.
- Cohesity supports only the VxFS filesystem on HP-UX servers.

Download and Install the HP-UX Agent

To download and install the HP-UX agent:

You can download the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Download the installer:
   a. Using the Cohesity cluster UI

      Navigate to **Data Protection** > **Sources** and click **Download Cohesity Agent**.

   b. Using the Cohesity Package Download site

      Download the Cohesity HP-UX Agent package.
2. To install the Custom certificates, follow the below steps:

   ```
   echo ENFORCE_USE_CUSTOM_CERTS=true >> /usr/local/cohesity/set_env.sh
   ```

   If you are using a Cohesity certificate then run the following commands to install the user certificate:

   ```
   echo AGENT_CERT_FILE=<path-to-agent-certificate-in-Cohesity-format> >>
   /usr/local/cohesity/set_env.sh
   ```

> **Note:**
> - ENFORCE_USE_CUSTOM_CERTS, if this parameter is set to true, then the agent will try to load the specified certificate and fail to start in case of an invalid file or files.
> - AGENT_CERT_FILE specifies the path to the agent certificate file in the Cohesity format.

3. Install the package:

   Run the following command to install the package as a Root user:

   ```
   swinstall -s <installer absolute path>  -x reinstall=true \*
   ```

   The Agent starts automatically. Now you can register the Server with the Cohesity cluster. Reboot is not required. The Agent starts after any future reboots.

Backup and Restore on HP-UX Server

Cohesity supports only file-based backup and restores on HP-UX servers.

To backup and restore on HP-UX server:

1. For creating backup jobs on the HP-UX server, see Add or Edit a Protection Group for Physical Servers.
2. To restore files and folders on the HP-UX server, see Recover Files or Folders to the Original Location and Recover Files or Folders to a New Location.

Consideration

- Cohesity supports only file-based backup and restores on HP-UX servers.

Upgrade the Agent

To upgrade the agent, follow the steps below:

You can upgrade the Cohesity agent installer from the Cohesity UI or manually download the agent from the Cohesity Package Download and install the agent using CLI.

1. Upgrade the installer:
   a. Using the Cohesity cluster UI

      Navigate to **Data Protection** > **Sources**. Click the action icon (⋮) next to the registered source and then click **Upgrade Agents**, if an upgrade is available.

   b. Using the Cohesity Package Download site

i. Download the latest Cohesity HP-UX Agent package.

ii. Run the following command to install the package as a Root user:

```
swinstall -s <installer absolute path>  -x reinstall=true \*
```

Uninstall the Agent

Uninstallation stops the running Agent and removes the libraries and other artifacts that were created during installation.

To uninstall the agent, run the following command:

```
swremove cohesity_agent
```

Additional Agent Commands

This section lists the additional HP-UX agent commands:

View Agent status

```
/usr/local/cohesity/agent/hpux_agent.sh status
```

Start the Agent

```
/usr/local/cohesity/agent/hpux_agent.sh start
```

Stop the Agent

```
/usr/local/cohesity/agent/hpux_agent.sh stop
```

Restart the Agent

```
/usr/local/cohesity/agent/hpux_agent.sh restart
```

Deploy Cohesity Agent from the Cohesity UI

Cohesity provides the ability to deploy Cohesity Agents on Windows or Linux servers from the Cohesity UI. The server can be a physical server or a virtual server. This section provides information on deploying Cohesity Agent on physical and virtual servers, and its use cases.

Supported Software

The following table lists the supported server on which you can deploy Cohesity Agent from the Cohesity UI.

Supported Server	Operating systems
Physical server	• Windows • Linux
VMware vCenter Server and ESXi standalone host	• Windows • Linux

For information on the supported version of the servers and operating systems, see Supported Software.

Deploy Cohesity Agent on Physical Servers

For protecting a physical server using Cohesity, you must first deploy the Cohesity agent on the physical server. You can deploy the Cohesity agent on a physical server:

- By downloading the Cohesity agent installer from the Cohesity UI and then running the installer on the server you want to protect.
- Directly from the Cohesity UI.

This section provides information on deploying the Cohesity agent on physical servers from the Cohesity UI.

You can use one of the following options to deploy the Cohesity agent on the physical servers from the Cohesity UI:

- **Enter Hostname/IP**: You can choose this option to deploy the Cohesity agent on multiple physical servers having the same credentials for server authentication.
- **Upload CSV**: You can choose this option to deploy agents on multiple physical servers having different credentials for server authentication.

Plan and Prepare

Before you deploy the Cohesity agent on a physical server, ensure the following prerequisites:

- The required firewall ports are open to enable communication between the Cohesity cluster and the physical server.
- Minimum permission required to deploy the Cohesity agent.
- WinRM is enabled on the Windows physical servers.
- SSH server is installed on the Linux physical servers.

Ports Used for Communication

Ensure the following ports are open based on the server you want to deploy the Cohesity agent:

Port	Usage	Windows	Linux
50051	To enable communication between the Cohesity cluster and theCohesity agent running on the physical server.	Yes	Yes
22	To enable communication between Cohesity nodes and Linux sources.	No	Yes
5985/5986	WinRM HTTP connectivity HTTPS	Yes	No

Minimum Permissions

Ensure you have the following permissions for deploying the Cohesity agent.

For Windows server

For deploying the Windows agent, you can either use the LOCAL SYSTEM account or an account that meets the following requirements to install the Cohesity agent:

- The account must be a member of the local Windows Administrators group. For example, qa01\tme-backup is an Active Directory user account in the data center. If the backup admin plans to use this account, qa01\tme-backup must be part of the local Windows Administrators group on the Windows server.
- The account must have Log on as a service in the User Rights Assignment on the Windows server to deploy the Cohesity agent.

For Linux Server

If you are deploying the Linux agent using an existing user, the user must have sudo permission. For sudo permission, you must add the user to the sudoers list by making the following changes in the **/etc/sudoers** file:

```
<user_name> ALL=(ALL) NOPASSWD:ALL
```

```
Defaults:<user_name> !requiretty
```

Deploy Cohesity agent by Entering Hostname or IP Address

If the physical servers on which you deploy the Cohesity agent have the same credentials for server authentication, then you can perform the following steps:

> **Note:** If the physical servers have different credentials, then you can deploy the Cohesity agent using the Upload CSV option. For more information, see Upload CSV.

To deploy the Cohesity agent on multiple physical servers having the same credentials:

1. Navigate to **Data Protection** > **Sources**.
2. Click the plus icon on the top-right of the page and then select **Deploy Cohesity Agent**.

   The Deploy Cohesity agent page appears.
3. Select **Enter Hostname/IP** as the option to deploy the Cohesity agent.
4. Select the operating system of the server from the drop-down list.
5. In the **Hostname or IP Address** field, enter one or more comma-separated IP addresses or hostnames of the servers with the same type of operating system. Ensure that the servers specified have the same username and password for authentication.
6. Click **Select**.
7. For server authentication, enter the username and password. For Linux server, you can optionally select the **SSH Key** tab and enter the user name and SSH key.
8. Disable the **Register as a Physical Server** option if you do not want to register the servers with the Cohesity cluster. If you disable this option, then to protect the physical servers, you must register the servers as described in the Register a physical server topic. By default, this option is enabled.
9. Optional. Click the down-arrow icon next to the **More Options** setting, and then click the **edit** icon to change the default settings if required:

Option	Description	
	Windows OS	**Linux OS**

Deployed As	By default, the Cohesity agent will be deployed using the **Local System account**.
	
	To change the default account, click the **edit** icon, select **User Account**, and then provide the credentials of a user account having local administrator privileges on the server with "Log on as a service" permission.
		By default, the Cohesity agent will be deployed using the `cohesityagent` user and group. During the agent deployment, the installer updates the **/etc/sudoers** file of the Linux server to allow `cohesityagent` sudo and no-tty sudo access.
		
		To change the default service user, click the **edit** icon, and then provide the user name and group name.
		
		For a new user type:
		
		If you have specified a new user, , then select the **Create the user if it does not exists** checkbox. By selecting this option, the installer updates the **/etc/sudoers** file of the Linux server to allow the specified user sudo and no-tty sudo access.
		
		For an existing user type:
		
		If the user you specified already exists, then you must add the user to the sudoers list by making the following changes in the **/etc/sudoers** file:
		
		`<user_name> ALL=(ALL) NOPASSWD:ALL`
		`Defaults:<user_name> !requiretty`
		
		An example of the sudoers file is as given below:
		
		`#includedir`

Option	Description	
	Windows OS	Linux OS
		`/etc/sudoers.d dgoble ALL=(ALL) NOPASSWD:ALL cohbackup ALL=(ALL) NOPASSWD:ALL Defaults:cohbackup !requiretty`
Deployment Path	By default, the Cohesity agent is deployed in the **C:\Program Files\Cohesity** directory. To change the default path, click the **edit** icon and specify the path.	By default, the Cohesity agent is deployed in the **/opt/cohesity/agent/** directory. To change the default path, click the **edit** icon and specify the path.
Components	• **Volume CBT (Change Blocks Tracker)**: This component is required to perform incremental backups. This component will be installed by default if you have selected to register the server as a physical server. You must reboot the server after installing this component. If you are not rebooting the server after installation or not installing this component, then only full backups will be performed. • **File System CBT (Change Blocks Tracker)**: This component is required for the incremental backup of individual MS SQL databases hosted on a server which is registered as a physical server. This component is not installed by default. To install this file system CBT, click the edit icon and enable the **File System CBT (Change Blocks Tracker)** option.	Not applicable

10. Click **Deploy**.

Cohesity will initiate the Cohesity agent deployment on the physical servers. You can view the status by clicking the notification icon on the top-right corner of the Source page.

Deploy Cohesity Agent by Uploading CSV

If the physical servers on which you deploy the Cohesity agent have different credentials for server authentication, then you can perform the following steps:

> **Note:** If the physical servers have the same credentials, then you can deploy the Cohesity agent using the **Enter Hostname/IP** option. For more information, see Deploy Cohesity agent by Entering Hostname or IP Address.

To deploy the Cohesity agent on multiple physical servers having different credentials:

1. Navigate to **Data Protection** > **Sources**.
2. Click the plus icon on the top-right of the page and then select **Deploy Cohesity Agent**.

   The Deploy Cohesity agent page appears.
3. From the Deploy Cohesity agent page, select **Upload CSV** as the option to deploy Cohesity agent.
4. Click **Download sample** file and download the CSV sample file to know the information and format in which the input must be provided in the CSV file. The following is the format in which the input must be provided in the CSV file:

   ```
   #IP/Hostname,OS Type,Server Username,Server Password,Install
   Location,Volume CBT,File CBT,Service Account Username,Service Account
   Password,SSH Key,
   ```

   Where,

   - `IP/Hostname`: The IP address or a valid hostname on which the Cohesity agent needs to be deployed.
   - OS type: The operating system of the server.
   - `Server Username` and `Server Password`: Credentials of the server on which the Cohesity agent needs to be deployed.
   - `Install Location`: (Optional) The path where the Cohesity agent must be deployed. If left blank, then the agent will be installed in the default location.
   - `Volume CBT` : (Optional) Applicable only for Windows OS types. If set to TRUE, the volume CBT driver will be installed. This field is ignored for non-Windows OS types.
   - `File CBT`: (Optional) Applicable only for Windows OS types. If set to TRUE, the file CBT driver will be installed. This field is ignored for non-Windows OS types.
   - `Service Account Username` and `Service Account Password`: (Optional) Credentials to register Cohesity agent as a service account.
   - `SSH Key`: (Optional) Use the SSH key to authenticate the server to deploy the Cohesity agent.

   For example:

- `#10.0.0.1,linux,admin,test_password,/var/agent_location/,,,,,,`
- `#10.0.0.2,windows,admin,test_password,C:\\Program Files\,,,service_admin,service_password,,`

> **Note:**
> - Each entry should end with a comma.
> - If no value is to be given for an optional field, it should be left blank between separating commas.

5. Click **Browse to upload a CSV file** and upload the CSV file in the format provided in the sample file.
6. Click **Select**.

Cohesity will initiate the Cohesity agent deployment on the physical servers. You can view the status by clicking the notification icon on the top-right corner of the Source page.

Deploy Cohesity Agent on Virtual Servers

By deploying the Cohesity agent on a virtual machine, you can:

- Perform efficient file-level recovery.
- Protect VMs as physical servers.
- Protect the applications running on the VMs.

This topic provides information on how to deploy the Cohesity agent on virtual machines from the Cohesity UI.

Plan and Prepare

Before you deploy the Cohesity agent on a virtual server, ensure the following prerequisites:

- The required firewall ports are open to enable communication between the Cohesity cluster and the physical server.
- Minimum permission required to deploy the Cohesity agent.
- Disable the Windows UAC.
- Understand the considerations.
- VMware tool is installed and running on the VMs of the virtual server.

Ports Used for Communication

For deploying the Cohesity agent on a virtual server, port 50051 must be open.

Minimum Permissions

- The user account that is used for auto-deploying agents in a Windows VMware environment must have administrator privileges on the VM where the agent is deployed. The account must have the "sysadmin" Server Role if the virtual server has SQL Server installed.
- If Windows firewall is used, add an inbound rule to accept the Cohesity agent and TCP connections on local port 50051.
- The user account that is used for file recovery to a Windows VM must have administrator privileges on the target VM. The account must have the "sysadmin" Server Role if the virtual server has SQL Server installed.

Deploy Cohesity Agent on VMs for file-level recovery

To perform efficient file-level recovery of the VMs of a virtual server, you must deploy the Cohesity agent on VMs.

To deploy the Cohesity agent on a virtual server from the Cohesity UI for efficient file-level recovery

1. Navigate to **Data Protection** > **Sources**.
2. Navigate the hierarchy to the VMware vCenter server to which you want to deploy the Cohesity agent.
3. Click the server's actions menu, and select **Deploy Cohesity Agent**.

   The Deploy Cohesity agent page appears.
4. Select the VMs onto which you want to deploy the Cohesity agent, and click **Select**. You can filter the VMs based on the operating system.

   > **Note:** You cannot select VMs of different operating systems. For example, if you have selected Windows VMs then you cannot select Linux VMs and vice-versa.
5. For server authentication, enter the guest username and password.
6. By default, the **Register as a Physical Server** option is disabled and you must not enable this option.

   If you enable this option, then the virtual server will be registered as a physical server and you cannot perform file-level recovery of the VMs.
7. Optional. Click the down-arrow icon next to the **More Options** setting and click the edit icon to change the default settings if required:

Option	Description	
	Windows OS	**Linux OS**

Deployed As	By default, the Cohesity agent will be deployed using the **Local System account**.
	
	To change the default account, click the **edit** icon, select **User Account**, and then provide the credentials of a user account having local administrator privileges on the server with "Log on as a service" permission.
		By default, the Cohesity agent will be deployed using the `cohesityagent` user and group. During the agent deployment, the installer updates the **/etc/sudoers** file of the Linux server to allow `cohesityagent` sudo and no-tty sudo access.
		
		To change the default service user, click the **edit** icon, and then provide the user name and group name.
		
		For a new user type:
		
		If you have specified a new user, , then select the **Create the user if it does not exists** checkbox. By selecting this option, the installer updates the **/etc/sudoers** file of the Linux server to allow the specified user sudo and no-tty sudo access.
		
		For an existing user type:
		
		If the user you specified already exists, then you must add the user to the sudoers list by making the following changes in the **/etc/sudoers** file:
		
		`<user_name> ALL=(ALL) NOPASSWD:ALL`
		`Defaults:<user_name> !requiretty`
		
		An example of the sudoers file is as given below:
		
		`#includedir`

Option	Description	
	Windows OS	**Linux OS**
		`/etc/sudoers.d dgoble ALL=(ALL) NOPASSWD:ALL cohbackup ALL=(ALL) NOPASSWD:ALL Defaults:cohbackup !requiretty`
Deployment Path	By default, the Cohesity agent is deployed in the **C:\Program Files\Cohesity** directory. To change the default path, click the **edit** icon and specify the path.	By default, the Cohesity agent is deployed in the **/opt/cohesity/agent/** directory. To change the default path, click the **edit** icon and specify the path.
Components	• **Volume CBT (Change Blocks Tracker)**: This component is required to perform incremental backups. This component will be installed by default if you have selected to register the server as a physical server. You must reboot the server after installing this component. If you are not rebooting the server after installation or not installing this component, then only full backups will be performed. • **File System CBT (Change Blocks Tracker)**: This component is required for the incremental backup of individual MS SQL databases hosted on a server which is registered as a physical server. This component is not installed by default. To install this file system CBT, click the edit icon and enable the **File System CBT (Change Blocks Tracker)** option.	Not applicable

8. Click **Deploy**.

Cohesity will initiate the Cohesity agent deployment. You can view the status by clicking the notification icon on the top-right corner of the Source page.

Deploy Cohesity Agent to Protect VMs as physical servers

From the Cohesity UI, you can deploy the Cohesity agent on VMs to protect them as physical servers.

To deploy Cohesity agent on VMs to protect the VMs as physical servers:

1. Navigate to **Data Protection** > **Sources**.
2. Click the plus icon on the top-right of the page and then select **Deploy Cohesity Agent**.

   The Deploy Cohesity agent page appears.
3. Select the **Browse vCenter** option to deploy the Cohesity agent on servers.
4. From the **Data Source** drop-down menu, select the virtual server that is already registered with the Cohesity cluster.
5. Select the VMs onto which you want to deploy the Cohesity agent, and click **Select**. You can filter the VMs based on the operating system.

   > **Note:** You cannot select VMs of different operating systems. For example, if you have selected Windows VMs then you cannot select Linux VMs and vice-versa.
6. Provide the guest username and password for authentication. For Linux servers, you can optionally select the **SSH Key** tab and enter the user name and SSH key.
7. By default, the **Register as a Physical Server** option is enabled to register the VMs as physical servers. If you disable this option, then the selected VMs will not be registered as physical servers and you cannot protect the VMs as physical servers.
8. Optional. Click the down-arrow icon next to the **More Options** setting and click the **edit** icon to change the default settings if required:

Option	Description	
	Windows OS	**Linux OS**

Deployed As	By default, the Cohesity agent will be deployed using the **Local System account**.
	
	To change the default account, click the **edit** icon, select **User Account**, and then provide the credentials of a user account having local administrator privileges on the server with "Log on as a service" permission.
		By default, the Cohesity agent will be deployed using the `cohesityagent` user and group. During the agent deployment, the installer updates the **/etc/sudoers** file of the Linux server to allow `cohesityagent` sudo and no-tty sudo access.
		
		To change the default service user, click the **edit** icon, and then provide the user name and group name.
		
		For a new user type:
		
		If you have specified a new user, , then select the **Create the user if it does not exists** checkbox. By selecting this option, the installer updates the **/etc/sudoers** file of the Linux server to allow the specified user sudo and no-tty sudo access.
		
		For an existing user type:
		
		If the user you specified already exists, then you must add the user to the sudoers list by making the following changes in the **/etc/sudoers** file:
		
		`<user_name> ALL=(ALL) NOPASSWD:ALL`
		`Defaults:<user_name> !requiretty`
		
		An example of the sudoers file is as given below:
		
		`#includedir`

Option	Description	
	Windows OS	Linux OS
		`/etc/sudoers.d dgoble ALL=(ALL) NOPASSWD:ALL cohbackup ALL=(ALL) NOPASSWD:ALL Defaults:cohbackup !requiretty`
Deployment Path	By default, the Cohesity agent is deployed in the **C:\Program Files\Cohesity** directory. To change the default path, click the **edit** icon and specify the path.	By default, the Cohesity agent is deployed in the **/opt/cohesity/agent/** directory. To change the default path, click the **edit** icon and specify the path.
Components	• **Volume CBT (Change Blocks Tracker)**: This component is required to perform incremental backups. This component will be installed by default if you have selected to register the server as a physical server. You must reboot the server after installing this component. If you are not rebooting the server after installation or not installing this component, then only full backups will be performed. • **File System CBT (Change Blocks Tracker)**: This component is required for the incremental backup of individual MS SQL databases hosted on a server which is registered as a physical server. This component is not installed by default. To install this file system CBT, click the edit icon and enable the **File System CBT (Change Blocks Tracker)** option.	Not applicable

9. Click **Deploy**.

Cohesity will initiate the Cohesity agent deployment. You can view the status by clicking the notification icon on the top-right corner of the Source page.

Deploy Cohesity Agent to register the Applications Running on the VMs

By deploying the Cohesity agent on a VM, you can register and protect applications such as Oracle database, MS SQL database, or a physical server running on a VM. While registering

the application with the Cohesity cluster, the Cohesity agent is also deployed on the VMs.

Register Oracle Database with Agents Deployed

Currently, from the Cohesity UI you can only deploy Linux Agent on the Oracle database.

To register Oracle database:

1. Navigate to **Data Protection** > **Sources**.
2. Navigate the hierarchy to the VMWare vCenter server of the VM that is running the Oracle database.
3. Mouse over the server's actions menu, and select **More Options** > **Register Oracle**.

   The Register Oracle page appears.
4. Optionally, from the **Data Source** drop-down menu, you can select a different vCenter server that is running Oracle database.
5. Select the VMs onto which you want to deploy the Cohesity agent, and click **Select**.
6. Enter the user name and password for guest authentication.
7. Select the authentication type to connect to the Oracle database:
   1. **Use Server Authentication**: If you want to run Oracle SQL commands as an OS authenticated user, then the installation should be performed as an Oracle OS user. By default, this option is enabled.
   2. **Use Database authentication**: Enter the database user name and password.

   For more information, see Credentials and Privileges for Oracle Database.

   > **Note:** For registering the Oracle database, as it is mandatory to register the virtual server as a physical server, the **Register as a Physical Server** option is enabled by default. You cannot change this setting.
8. Optional. Click the down-arrow icon next to the **More Options** setting and click the **edit** icon to change the default settings if required:
   1. **Deployment Path**: By default, the Cohesity agent is deployed in the **/opt/cohesity/agent/** directory. To change the default path, click the **edit** icon and specify the path.
   2. **Deploy As**: By default, the Cohesity agent will be deployed using the `cohesityagent` user and group. During the agent deployment, the installer updates the **/etc/sudoers** file of the Linux server to allow cohesityagent sudo and no-tty sudo access.

      To change the default user type, click the **edit** icon, and then provide the user name and group name for the agent.

If you have specified a new user, , then select the **Create the user if it does not exists** checkbox. By selecting this option, the installer updates the **/etc/sudoers file** of the Linux server to allow the specified user sudo and no-tty sudo access.

If the user you specified already exists, then you must add the user to the sudoers list by making the following changes in the **/etc/sudoers** file:

```
<user_name> ALL=(ALL) NOPASSWD:ALL
Defaults:<user_name> !requiretty
```

An example of the sudoers file is as given below:

```
#includedir /etc/sudoers.d
dgoble ALL=(ALL) NOPASSWD:ALL
cohbackup ALL=(ALL) NOPASSWD:ALL
Defaults:cohbackup !requiretty
```

9. Click **Register**.

   The oracle database running on the VM will be registered along with the agent deployed.

Cohesity will initiate the Cohesity agent deployment. You can view the status by clicking the notification icon on the top-right corner of the Source page.

Register MS SQL Database with Agents Deployed

To register MS SQL Database:

1. Navigate to **Data Protection** > **Sources**.
2. Navigate the hierarchy to the server that is running the MS SQL database.
3. Mouse over the server's actions menu, and select **More Options** > **Register MS SQL**.
4. Optionally, from the **Data Source** drop-down menu, you can select a different server that is running the MS SQL database.
5. Select the VMs onto which you want to deploy the Cohesity agent, and click **Select**. Since MS SQL is supported only on Windows operating systems, only the VMs running Windows OS are displayed.
6. Provide the MS SQL database username and password for authentication.
7. Enable the **Register as a Physical Server** option if you want the virtual server to be registered as a physical server. By default, this option is disabled.
8. Optional. Click the down-arrow icon next to the **More Options** setting and click the

edit icon to change the default settings if required:

Option	Description
Deployed As	Deploy Cohesity agent using any of the following account: • **Local System**: By default the agent will be deployed using the Local System account. • **User Account**: Use the user credentials of a user account having local administrator privileges on the server with "Log on as a service" permission. For more information, see Requirements for Microsoft SQL Server Protection.
Deployment Path	By default, the Cohesity agent is deployed in the **C:\Program Files\Cohesity** directory. To change the default path, click the **edit** icon and specify the path.
Components	• **Volume CBT (Change Blocks Tracker)**: This component is required to perform incremental backups. You must reboot the server after installing this component. If you are not rebooting the server after installation or not installing this component, then only full backups will be performed. This component is not installed by default. To install volume CBT, click the edit icon and enable the File System CBT (Change Blocks Tracker) option. • **File System CBT (Change Blocks Tracker)**: This component is required for the incremental backup of individual MS SQL databases hosted on a server which is registered as a physical server. This component will be installed by default if you have selected to register the virtual server as a physical server.

9. Click **Register**.

Cohesity will initiate the Cohesity agent deployment. You can view the status by clicking the notification icon on the top-right corner of the Source page.

Register Physical Servers with Agents Deployed

To register physical server:

1. Navigate to **Data Protection** > **Sources**.
2. Navigate the hierarchy to the server where you want to register the physical server.
3. Mouse over the server's actions menu, and select **More Options** > **Register Physical**.

   The Register Physical Server page appears.
4. Optionally, from the **Data Source** drop-down menu, you can select a different server where you want to register the physical server.

5. Select the VMs onto which you want to deploy the Cohesity agent, and click **Select**. You can filter the VM based on the operating system.

   > **Note:** You cannot select VMs of different operating systems. For example, if you have selected Windows VMs then you cannot select Linux VMs and vice-versa.

6. Provide the username and password for authentication.
7. Optional. Click the down-arrow icon next to the **More Options** setting and click the edit icon to change the default settings if required:

Option	Description	
	Windows OS	**Linux OS**

Deployed As	By default, the Cohesity agent will be deployed using the **Local System account**.
	
	To change the default account, click the **edit** icon, select **User Account**, and then provide the credentials of a user account having local administrator privileges on the server with "Log on as a service" permission.
		By default, the Cohesity agent will be deployed using the `cohesityagent` user and group. During the agent deployment, the installer updates the **/etc/sudoers** file of the Linux server to allow `cohesityagent` sudo and no-tty sudo access.
		
		To change the default service user, click the **edit** icon, and then provide the user name and group name.
		
		For a new user type:
		
		If you have specified a new user, then select the **Create the user if it does not exists** checkbox. By selecting this option, the installer updates the **/etc/sudoers** file of the Linux server to allow the specified user sudo and no-tty sudo access.
		
		For an existing user type:
		
		If the user you specified already exists, then you must add the user to the sudoers list by making the following changes in the **/etc/sudoers** file:
		
		`<user_name> ALL=(ALL) NOPASSWD:ALL`
		`Defaults:<user_name> !requiretty`
		
		An example of the sudoers file is as given below:
		
		`#includedir`

Option	Description	
	Windows OS	Linux OS
		`/etc/sudoers.d dgoble ALL=(ALL) NOPASSWD:ALL cohbackup ALL=(ALL) NOPASSWD:ALL Defaults:cohbackup !requiretty`
Deployment Path	By default, the Cohesity agent is deployed in the **C:\Program Files\Cohesity** directory. To change the default path, click the **edit** icon and specify the path.	By default, the Cohesity agent is deployed in the **/opt/cohesity/agent/** directory. To change the default path, click the **edit** icon and specify the path.
Components	• **Volume CBT (Change Blocks Tracker)**: This component is required to perform incremental backups. This component will be installed by default if you have selected to register the server as a physical server. You must reboot the server after installing this component. If you are not rebooting the server after installation or not installing this component, then only full backups will be performed. • **File System CBT (Change Blocks Tracker)**: This component is required for the incremental backup of individual MS SQL databases hosted on a server which is registered as a physical server. This component is not installed by default. To install this file system CBT, click the edit icon and enable the **File System CBT (Change Blocks Tracker)** option.	Not applicable

8. Click **Register**.

Cohesity will initiate the Cohesity agent deployment. You can view the status by clicking the notification icon on the top-right corner of the Source page.

Troubleshoot Cohesity Agent

This topic provides the troubleshooting information for Cohesity Agents.

Collect Diagnostic Information for Cohesity Agents

For Microsoft Windows servers, you can find the logs for Cohesity Agents at `C:\ProgramData\Cohesity\`.

If the **ProgramData** folder is hidden or not visible, then from the top-menu of the **C** directory, click **View**, and then check the **Hidden items** check box.

The hidden folders and files will be displayed.

Diagnostic data is subject to the following default settings. For more information, refer to the table below.

Default Settings - Linux (x64)	Value
Maximum size of a log file before a new file is created	30MB
Log file cleanup interval (frequency at which older logs are cleaned up)	1 Hour
Maximum retention time for a log file	30 days
Maximum space allowed to store log files	50MB
Maximum number of core archives to retain	5
Minimum batch size for log files before cleanup kicks in (at least this many log files should be present for cleanup based on time or size to occur)	8

Default Settings - Windows	Value
Max log retention time for a log file	7 days
Log file cleanup interval (frequency at which older logs are cleaned up)	1 Hour
The maximum amount of space allowed for storing logs (in MB)	10MB
The maximum number of core archives to retain	5
The maximum number of Cristie log archives to retain	2
The maximum number of Windows log archives to retain	500
The maximum number of days for storing diagnostics file	1 day
The maximum amount of space allowed for storing diagnostics file (in MB)	100MB

Default Settings - JAVA agents (AIX, Solaris, HP-UX, Linux PowerPC LE)	Value
Maximum size of a log file before a new file is created	30MB
Max number of agent log files to retain	20
Max log retention time for a log file	7 days
Log file cleanup interval (frequency at which older logs are cleaned up)	1 Hour

> **Note:** Contact Cohesity Support to modify the default diagnostic settings.

Find the Cohesity Agent Version

To find the version of the Cohesity Agent installed on Windows, Linux, and AIX servers, perform the following steps:

1. Navigate to **Data Protection > Sources**
2. Click **Register > Physical Server**.
3. Select a server. If the installed Agent's version is lower than the Cohesity cluster version, it will display **Upgrade available**.

Prerequisites

Ports 80 and 443 are open from the client to the cluster so that the Cohesity agent upgrade from UI can occur.

Troubleshooting Agent Issues

The following table provides details about Agent issues that you might encounter:

Issue	Cause	Workaround	

Registering a Cohesity Agent on a new Cohesity cluster fails with the following error: `Host is already registered with another cluster: <old cluster-id>, <new cluster-id>`	Agent is already registered to a different Cohesity cluster ID.	1. Navigate to **Data Protection > Sources > Register > Physical Server**. 2. cluster_ fqdn/protection/sources/new/physical `?forceRegister=true`" to the end of the URL and press the **Enter** key. You will be redirected to a new registration page. 3. Enter the Agent information and ensure that the Force Agent Registration check box is selected.	
Installing or upgrading the Cohesity Agent either manually or automatically fails with the following error: `WinHttpSendRequest() failed : 0x2ee7`	This error indicates a DNS issue in Windows. While installing the Cohesity Agent, the Agent confirms that it can resolve the host name of the Cohesity cluster. If this check fails, this error is returned.	Ensure DNS forward and reverse lookups are working on the server or VM.	

Login to the Cohesity Support site to see more Knowledge Base Articles.

Upgrading the Cohesity Agent on a Windows Physical Server may Become Stuck

Under specific conditions, upgrading the Cohesity agent on a Windows physical server may become stuck. To identify if a server is affected by this issue, read the following:

Check if the agent logs (under %programdata%\Cohesity) contain "Cannot start the agent because the port 50051 is in use," indicating that the agent is in a crash loop.

The upgrade issue can occur if all of the following are true:

- The Cohesity agent is running under a specific user account (not the default Local System account).
- The Cohesity agent service is the only service in the running state at the time of upgrade using the specified account.
    - The specified user account has not been logged into through other means (such as RDP) at the time of the upgrade.
- The "Do not forcefully unload the user registry at user logoff" policy is "Not Configured" (the default) or set to "No."

To get the agent out of this state:

1. Open Windows Task Manager and click the Details tab.
2. Locate the "installer.exe" and "installer.tmp" processes (with the Cohesity icon).
3. Right-click "installer.tmp" and select "End Task". This stops both processes ("installer.exe" and "installer.tmp").

To avoid future occurrences:

Use the Group Policy Editor (gpedit.msc) to enable the "Do not forcefully unload the user registry at user logoff" policy. The policy is located under Computer Configuration > Administrative Templates > System > User Profiles.

Chapter 30: NoSQL and Hadoop Service

NoSQL and Hadoop Service is an enterprise data management solution to protect Hadoop clusters and NoSQL distributed databases on a single web-scale platform. The NoSQL and Hadoop Service expands Cohesity's DataProtect coverage to support protection for Couchbase, MongoDB, Cassandra, Hbase, Hive, HDFS, all managed through a single policy framework and global user interface. See the following topics for details:

Supported Software 4517
Features 4518
Prerequisites 4519
Process Overview 4519
Install NoSQL and Hadoop Service 4520
Run the NoSQL and Hadoop Service 4522
Verify the Status of NoSQL and Hadoop Service 4523
Using NoSQL and Hadoop Data sources 4523
Post Installation Procedures 4523
Managing NoSQL and Hadoop Services 4539
Additional Configuration for NoSQL and Hadoop Services 4546
Monitor the NoSQL and Hadoop Service Health 4552
FAQs on NoSQL and Hadoop Services 4560

Supported Software

Cohesity cluster Version: 6.5.1 or higher

For details on the supported versions for Cassandra, Couchbase, MongoDB and Hadoop:

Application	Details
Cassandra	See Supported Versions for Cassandra DataProtect Solution.
Couchbase	See Supported Versions for Couchbase DataProtect Solution.
MongoDB	See Supported Versions for MongoDB DataProtect Solution.
Hadoop	See Supported Versions for Hadoop DataProtect Solution.

Features

- Supports protection of NoSQL databases such as Couchbase, MongoDB, Cassandra, and Hadoop (HDFS, Hive, and Hbase).
- Scalable backup and recovery architecture supporting petabytes of data.
- Meet business RPOs/RTOs with parallel architecture to support rapid backup and recovery.
- Supports a simple agentless deployment.
- Offers a unified and automated policy-driven model involving no manual steps.
- Supports a flexible backup and recovery method from a database to a table or a collection level.
- Offers incremental forever backup.
- Recovery is topology-independent.
- Improved data resiliency with erasure coding.
- Reduced data footprint with global variable-length deduplication, compression.
- Native cloud integrations for simplified archival to the public cloud for long-term retention.
- One-click archival to public cloud for long term retention using native cloud integrations.
- An API-first architecture approach with RESTful API support.
- Integration with Active Directory, LDAP, and Kerberos.
- Encryption of data in-flight and at rest.
- Customizable role-based access control.
- Supports protection of NoSQL databases such as Cassandra and Couchbase configured using IPv6 addresses.

Prerequisites

- The NoSQL and Hadoop service is designed to handle petabytes of data efficiently by scaling out to perform parallel backups and restores. As a result, the services require a higher amount of CPU and memory depending on the RPO & RTO needs. Compute Nodes in the Cohesity cluster provide CPU and memory resources to run resource-intensive apps and services. In order to install the NoSQL and Hadoop service, the Cohesity cluster must have compute nodes added. See Manage Compute Nodes for details.
- Data and Compute nodes on the Cohesity cluster must have connectivity to all the primary cluster nodes of your data source (Hadoop cluster nodes or Cassandra cluster nodes or MongoDB cluster nodes) for successful discovery, backup and recovery.
- For more information on managing apps and the prerequisites for apps management, see Manage Apps .

Process Overview

A quick overview of the process is as follows:

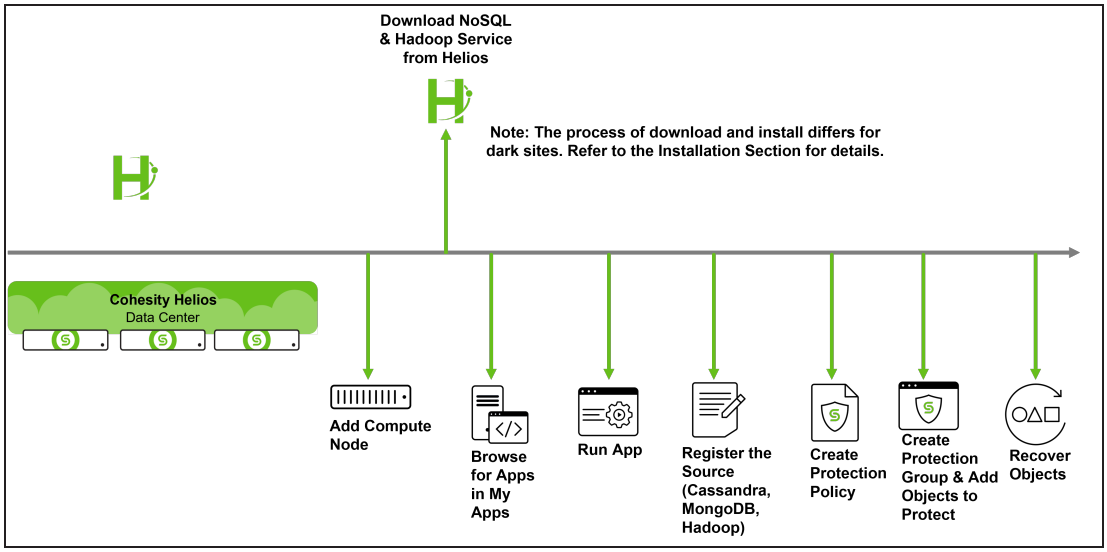

1. To use NoSQL and Hadoop service, you must first add a compute node.
2. Download the NoSQL and Hadoop Service from Helios, and Run App.
3. Configure the **QoS Policy** and **Instant Size** based on the business requirement and run the application.
4. Ensure that the advanced configuration settings are enabled for NoSQL and Hadoop solutions. Contact Cohesity Support to enable this capability.

> **Note:** Note: You can perform this step even before you initiate the app installation process. However, to be able to perform any operations, you must have the advanced configuration settings enabled.

5. Perform various operations such as register, back up and recover Cassandra, Couchbase, Hadoop and MongoDB solutions.

Install NoSQL and Hadoop Service

> **Important:** Ensure you have an understanding of the Cohesity Apps and met the prerequisites required to install an App. For more information, refer to Manage Apps Manage Apps.

> **Note:**
> - The NoSQL and Hadoop Service must be installed and run as a local admin. Using an SSO or AD user to install the app might cause issues while trying to launch the container instances.
> - Upgrading the NoSQL app from version 5.4.11 to version 7.0.1 is not supported. Install the NoSQL app version 7.0.1 when you upgrade Cohesity from version 6.8.1 to version 7.0.1.
> - All the workflows and policies configured in version 5.4.11 will be preserved when you install the NoSQL app version 7.0.1.
> - To preserve the SSL/Kerberos settings done in version 5.4.11, perform the following step on any node, before stopping the NoSQL app in version in 5.4.11:
>
> ```
> imanis_cli.sh -c config save
> ```

To download the NoSQL and Hadoop Service from Helios and install it on your cluster:

1. Log in to Helios with valid credentials.
2. Navigate to **Marketplace** > **My Apps** > **Browse Apps**.

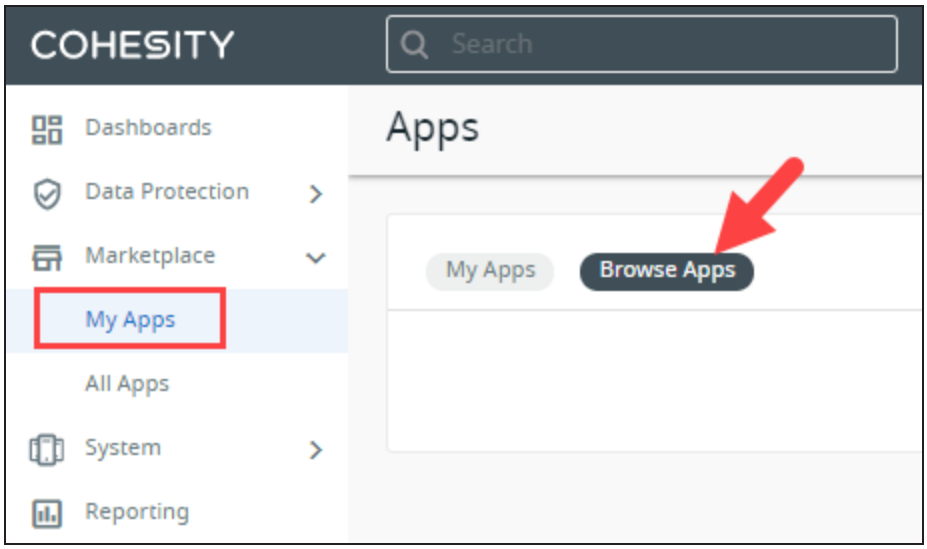

3. Search for the **NoSQL and Hadoop Service** app, and click **Install App**.
4. Select the Cohesity cluster from the list of registered clusters on which you wish to run the app, and click **Install**.

   The service is installed successfully.

To download the NoSQL and Hadoop Service from Helios and install it on a dark site:

1. Log in to Helios with valid credentials.
2. Navigate to **Marketplace** > **My Apps** > **Browse Apps**.
3. Search for the **NoSQL and Hadoop Service** app, and click the title of **NoSQL and Hadoop Service**.
4. Select the **Download** tab (at the end).
5. Select the cluster version and click **Get App versions**.
6. Click on the latest version of the app displayed, and click **Download** to initiate downloading of the app package.
7. Transfer the app package to the dark site where the NoSQL and Hadoop service needs to run from.
8. You can then log in to the Cohesity cluster, navigate to **Marketplace** > **My Apps**, click **Upload App** and select the downloaded app package.
9. Click **Upload and Install**.

   The service is installed successfully.

> **Note:** To use MongoDB CDP, you need to ensure that the set of prerequisites under Configuring Continuous Data Protection (CDP) for MongoDB is satisfied. This should be done before running the NoSQL and Hadoop Services app.

Run the NoSQL and Hadoop Service

1. Log in to the Cohesity cluster, and navigate to **Marketplace** > **My Apps** > **My Apps** tab.
2. If the app upload is successful, the NoSQL & Hadoop Service App is listed and the **Run App** button appears.
3. Click **Run App** to start running the app instance on your cluster.

   > **Note:** You can run only one instance of the app per cluster.

4. Based on the requirement, select **QoS Policy**, **Views**, and **Instant Size**. Enter a brief description (optional) and click **Run App**.

   An instance of the app will be initiated.

Parameters	Description
QoS Policy	This decides how much % of resources of compute nodes are allocated to the NoSQL and Hadoop service. **Note:** • You can set this value as Max if this is the only app instance running on your cluster. However, if you wish to use other apps at the same time, you can set a lower QoS policy. • Max QoS allows 80% of total allocatable resources for NoSQL App. • High QoS allows 50% of total allocatable resources for NoSQL App. • It is recommended that you allow maximum QoS for NoSQL app for better performance.

Parameters	Description
Instance Size	This indicates the size of the container. It is recommended to use a large instance size. Each instance of nodemanager container will be of size • 16GB/4core (small) • 32GB/8core (medium) • 40GB/10core (large) App nodes (compute only nodes) should have enough resources to fit at least 3 small instances of nodemanagers. **Note:** Minimum of 3 instances of nodemanagers are required for NoSQL and Hadoop services to run.

Verify the Status of NoSQL and Hadoop Service

1. Navigate to **Marketplace** > **My Apps** and then click the **All Instances** tab.
2. Make sure the instance status for **NoSQL and Hadoop Service** is **Running**.

> **Note:** If the app status is displayed as Unhealthy, then confirm if the compute node is up and running. Otherwise, Contact Cohesity Support.

Using NoSQL and Hadoop Data sources

To work with the NoSQL and Hadoop data sources, see:

- Working with Cassandra DataProtect Solution
- Working with Couchbase DataProtect Solution
- Working with MongoDB DataProtect Solution
- Working with Hadoop DataProtect Solution

Post Installation Procedures

After the successful installation of NoSQL and Hadoop Service, you must run some of the post installation scripts to ensure that you can seamlessly work with the NoSQL and Hadoop solutions. Some of the post installation procedures to be followed are:

- To Verify Installation
  - Verify Services
  - Verify UI
- To Manage the NoSQL and Hadoop Services
  - Retrieve NoSQL and Hadoop Services Version
  - Work with Kerberos Security enabled primaries for data protection
- Manage SSL enabled NoSQL primary clusters
  - Certificate file formats supported
  - Work with NoSQLandHadoop Services using SSL tool

> **Note:**
> - In the examples below, you can run the `kubectl.sh` and the `imanis_cli` commands from the Cohesity cluster node. Once you execute the `kubectl.sh exec` command, you get logged into the Imanis pod. Subsequently, all the `talena` based commands can be run from that Imanis pod.
> - `imanis_cli.sh` tool must be executed from the Cohesity cluster node where Imanis pods are hosted or running.

To Verify Installation

Before using NoSQL and Hadoop services, you must verify if the installation has been successful. You can either verify the installation on the NoSQL and Hadoop services as well as verify if the GUI is functioning correctly.

Verify Services

Log in to any one Cohesity cluster node and if the installation is successful, the following will be displayed:

```
$ kubectl.sh get pods -A | grep imanis
athena-34                imanis-api-server-0
 1/1     Running             0               13h
athena-34                imanis-historyserver-0
 1/1     Running             0               13h
athena-34                imanis-nodemanager-0
 1/1     Running             0               13h
athena-34                imanis-nodemanager-1
 1/1     Running             0               13h
athena-34                imanis-nodemanager-2
 1/1     Running             0               13h
```

```
athena-34                imanis-nodemanager-3
 1/1     Running             0               13h
athena-34                imanis-nodemanager-4
 1/1     Running             0               13h
athena-34                imanis-nodemanager-5
 1/1     Running             0               13h
athena-34                imanis-nodemanager-6
 1/1     Running             0               13h
athena-34                imanis-resourcemanager-0
 1/1     Running             0               13h
athena-34                imanis-zookeeper-0
 1/1     Running             0               13h
athena-34                imanis-zookeeper-1
 1/1     Running             0               13h
athena-34                imanis-zookeeper-2
 1/1     Running             0               13h
```

In the above example, NoSQL and Hadoop Service is running in 12 kubernetes pods. The namespace name is **athena-104**. You can log in to any of the pod (usually, imanis-nodemanager-0 is used as it has more memory and cpu to run additional load).

```
$ kubectl.sh exec -it imanis-nodemanager-0   -n athena-104 /bin/bash
```

```
[hdfs@imanis-nodemanager-0 /]$ talena-services.sh status all
```

Running Smoke Tests

If all services are running, then you can run the following smoke tests for verification:

```
$ talena-services.sh smoke all
```

If any smoke test fails, then it needs attention before proceeding to use the app.

Verify UI

To verify if the NoSQL and Hadoop data sources are enabled in Cohesity for the first time:

1. Navigate to **Data Protection** > **Sources**.
2. Click the plus icon (+) on the top-right of the page and then verify if the Couchbase, Hadoop, Cassandra and Mongo DB options are displayed in the menu.

To Manage the NoSQL and Hadoop Services

This section provides basic command line options that you can use to manage the NoSQL and Hadoop services.

```
$ imanis_cli.sh
INFO : Checking reachability of pod imanis-nodemanager-0
--------------------------------------------------------------------
```

```
Usage: imanis_cli_main.sh  -c command_name [ command_args ]

-c command_name : valid command_name is one of con-
fig,explorer,hosts,joblog,kerberos,patch,ssl,vip
command_args : "options_string_enclosed_in_double_quotes"
-----------------------------------------------------------------------
```

Retrieve NoSQL and Hadoop Services Version

The following section discusses the procedure to determine the version or release number of NoSQL and Hadoop Services.

1. Navigate to **Marketplace** > **My Apps**.
2. Select the **My Apps** tab to verify the version of NoSQL and Hadoop service.

Work with Kerberos Security enabled primaries for data protection

If your primary cluster to be protected is not Kerberos security enabled, then you can ignore this section. Otherwise, please go through this section before attempting to register a data source in Cohesity. Typical flow of this process is:

a. Obtain Kerberos principal name(s), the corresponding keytab file(s) and Kerberos configuration file (krb5.conf) for your kerberos realm.
b. Ensuring the Kerberos KDC server hostname is resolvable from NoSQL and Hadoop Service pods.
c. Adding the keytab and the kerberos realm to NoSQL and Hadoop Service using a command line tool called imanis_cli.sh.
d. Register the data source in Cohesity.

Obtain Kerberos Principal(s), keytab file(s) and krb5.conf

Verify with your primary cluster administrators about one or more kerberos principals to be used by NoSQL and Hadoop service for connecting to their cluster.

> **Note:**
> NoSQL and Hadoop service recommends using a single principal exclusively for NoSQL and Hadoop service access. This is to ensure better control over what access permissions are granted to this principal and can also be useful for auditing purposes. For example, one can create a principal such as imanis@REALMNAME in the KDC server for exclusive use by NoSQL and Hadoop Service software.
>
> However, in some cases, while protecting Hadoop clusters like CDH or HDP, some deployments may prefer to create 3 separate kerberos principals for allowing access to hdfs, hbase and hive respectively as follows:
>
> - imanis_hdfs@REALMNAME - Kerberos principal used for HDFS workflows
> - imanis_hbase@REALMNAME - Kerberos principal used for HBase workflows
> - imanis_hive@REALMNAME - Kerberos principal used for Hive workflows
>
> In such cases, you may end up with 3 different keytab files corresponding to these 3 principals as opposed to a single keytab file in case of using a single principal.

The key thing to ensure is that the principal(s) given to NoSQL and Hadoop Service have full access permissions granted on the primary cluster, for listing, reading files, tables, keyspaces, databases, collections and write access while restoring a backup from NoSQL and Hadoop Service.

> **Note:** Replace REALMNAME with the actual realm name used in your kerberos cluster. You can use any other existing principal name such as hdfs@REALMNAME , instead of imanis. New principles are created in KDC servers by Kerberos admin.

Ensure Kerberos KDC server hostname is resolvable from NoSQL and Hadoop Service pods

We must ensure that the Kerberos KDC server hostname is resolvable from the NoSQL and Hadoop Service kubernetes pods. This is generally not a problem if the DNS is setup properly and we use FQDN for the Kerberos KDC server. However, if this is not the case, then one can also add this hostname and IP address to /etc/hosts on the pods running NoSQL and Hadoop Service.

Use imanis_cli.sh -c hosts add option to perform this step:

```
$ pwd /home_cohesity_data/athena/mounts/_internal_view_mount/imanis
$ cat $PWD/hosts
10.14.20.237 im39-vm1.cohesity.com
$ imanis_cli.sh -c hosts add $PWD/hosts
```

Add a keytab and Kerberos realm to NoSQL and Hadoop Service

You need to perform this step every time you want to register a new kerberos enabled data source with NoSQL and Hadoop Service. However, if the principal that you plan to use is the same as other data sources that have already been registered, then you can skip this step.

Run a command line tool called imanis_cli.sh on any one Cohesity cluster node. The following options are supported by imanis_cli.sh for the kerberos operations:

```
$ imanis_cli.sh -c kerberos usage:
 imanis_cli.sh -c kerberos addrealm
 imanis_cli.sh -c kerberos deleterealm realm_name
 imanis_cli.sh -c kerberos validatekeytab [local keytabpath]
 imanis_cli.sh -c kerberos import krb5_keytab_dirpath
```

For adding a new kerberos principal, we use the add realm option. There are two ways in which the information about the new Kerberos realm to be added is given as input to this script.

- By pointing the script to an existing krb5.conf file. The conf file will be read and parsed to extract the relevant information. Download the krb5.conf file from your primary node to the node on NoSQL and Hadoop Services
- In case you do not have access to krb5.conf file, you can type the required information manually by confirming the prompts. Download the keytab file(s) and krb5.conf (if you have one) to one of the Cohesity cluster nodes under the directory as follows:

```
cd /home_cohesity_data/athena/mounts/_internal_view_mount/imanis/
mkdir kerberos
cd kerberos
scp ip_or_hostname:/path/to/keytab_and_krb5_conf_files .
```

Here, replace the ip_or_hostname with the IP address or hostname of the server where the keytab and krb5.conf files are present. This is usually your primary cluster itself or any other server where administrators have generated these files and saved.

> **Note:** It is important to download and keep these files under this specific mount point only. This cannot be changed, however you can create any sub directory.
>
> `/home_cohesity_data/athena/mounts/_internal_view_mount/imanis/`

You must have one or more keytabs and/or krb5.conf file in the present working directory as follows:

```
$ pwd
/home_cohesity_data/athena/mounts/_internal_view_mount/imanis/kerberos
$ ls
```

```
imanis.keytab  krb5.conf
```

If **krb5.conf** file is available:

```
$ imanis_cli.sh -c kerberos addrealm
INFO : Checking reachability of pod imanis-nodemanager-0

INFO: Checking nodes reachability

WARN : All values from [libdefaults] section in krb5.conf on NoSQL will not
be changed
       Manual intervention would be required to change values from [lib-
defaults] section(if required)
Do you wish to continue (yes/no) yes ?

INFO: Kerberos configs can be read from krb5.conf OR
      User can interactively provide Kerberos configs
      addrealm need details like KDCServer/Realm/DomainRealms
Do you wish to use krb5.conf to get Kerberos configs (yes/no) yes ?

NOTE: Please download krb5.conf and keytabs from primary to local directory
      Path should always begin with /home_cohesity_data/athena/mounts/_
internal_view_mount

Provide local directory path where krb5.conf and keytabs downloaded from
primary : /home_cohesity_data/athena/mounts/_internal_view_moun-
t/imanis/im20-vm4/

Default Realm : im20-vm4.realm


KDC servers : im20-vm4.cohesity.com

Domain Realm list : im20-vm5.cohesity.com,im20-vm6.cohesity.com,im20-vm4.-
cohesity.com
Do you wish to continue (yes/no) yes ?

INFO: Validating domain_realms
Please enter FQDN of one of your datasource cluster node : im20-vm6.-
cohesity.com
INFO : domain_realm validated against im20-vm6.cohesity.com
INFO: Successfully read keytab: /cohesity/mount/sharedview/imanis/im20-
vm4/im20-vm4.keytab
INFO: Checking KDC server(im20-vm4.cohesity.com) reachability from imanis-
api-server-0
```

```
INFO: Merged keytab will be available at location /op-
t/talena/conf/talena.keytab
INFO: Updating hadoop conf files
INFO: Verifying added realm im20-vm4.realm
Imanis api-server restart is needed after addrealm

Would you like to restart api-server now (yes/no) yes ?
Stopping API server...
API server  stopped successfully
Starting API server...
API server is running on [ imanis-api-server-0]
INFO: addrealm completed for im20-vm4.realm
```

If **krb5.conf** file is not available :

```
$ imanis_cli.sh -c kerberos addrealm
Fetching ImanisData cluster config from athena-6:imanis-nodemanager-0

INFO: Checking nodes reachability

INFO: Kerberos configs can be read from krb5.conf OR
User can interactively provide Kerberos configsaddrealm need details like
KDCServer/Realm/DomainRealms
Do you wish to use krb5.conf to get Kerberos configs (yes/no) yes ? no

Please specify following kerberos configs.

Primary Realm name   : im39-vm1.realm
Primary KDC hostname : im39-vm1.cohesity.com
Primary Domain Realm list from domain_realm section (| separated) :
t14vm3.example.com|t14vm2.example.com|t14vm1.example.com
NOTE: Please download keytabs from primary to local directory
Path should always begin with /home_cohesity_data/athena/mounts/_internal_
view_mount
Provide local directory path where keytabs downloaded from primary : /home_
cohesity_data/athena/mounts/_internal_view_mount/imanis/kerberos

INFO: Successfully read keytab: /co-
hesity/mount/sharedview/imanis/kerberos/keytab

INFO: Checking KDC server(im39-vm1.cohesity.com) reachability from imanis-
api-server-0
INFO: Checking KDC server(im39-vm1.cohesity.com) reachability from imanis-
historyserver-0
INFO: Checking KDC server(im39-vm1.cohesity.com) reachability from imanis-
```

```
nodemanager-0
INFO: Checking KDC server(im39-vm1.cohesity.com) reachability from imanis-
resourcemanager-0
INFO: Checking KDC server(im39-vm1.cohesity.com) reachability from imanis-
zookeeper-0


INFO: keytabs will be merged with /opt/talena/conf/talena.keytab

INFO: Updating hadoop conf files

INFO: Verifying added realm im39-vm1.realm

Imanis api-server restart is needed after addrealm

Would you like to restart api-server now (yes/no) yes ?yes
Stopping API server...
API server  stopped successfully
Starting API server...
API server is running on [imanis-api-server-0]

INFO: addrealm completed for im39-vm1.realm
```

> **Note:** Primary Domain Realm list from domain_realm section (separated by |) :This is the list of hosts that you would like to be translated to the given kerberos realm. Typically, these are the FQDNs of your primary cluster nodes.

This script can also be used to make changes to the keytab for an existing data repository, without having to modify the realm information. If the Kerberos realm information is already present with the NoSQL and Hadoop Services, then only keytab is processed by the script and changes are incorporated.

Register the data source in Cohesity

Optionally, you can verify the kerberos configuration as follows before proceeding to the registration step.

```
$ imanis_cli.sh -c kerberos validatekeytab

Fetching ImanisData cluster config from athena-6:imanis-nodemanager-0
Keytab validation completed successfully
```

Manage SSL enabled NoSQL primary clusters

In general, two entities (client, server programs in our case) that are communicating with one another must ensure the following:

1. Client to ensure that it is talking to an authentic server.
2. Server to ensure that it is talking to an authentic client.
3. The data client-server exchange must not be seen by anyone else in the middle.

See FAQs on NoSQL and Hadoop Services for details on the most frequently asked questions about SSL enabled NoSQL primary clusters.

Certificate file formats supported

NoSQLandHadoop uses PEM format for certificates and keys for importing.

PEM or Privacy Enhanced Mail is a Base64 encoded DER certificate. So you can open it in your favorite editor to check. Generally when a PEM encoded file is opened in a text editor, it contains very distinct headers and footers as below:

A private key in PEM format is as shown below:

```
-----BEGIN RSA PRIVATE KEY-----
Proc-Type: 4,ENCRYPTED
DEK-Info: DES-EDE3-CBC,E1537934256F0AAE
aNLbkgsTZLESaRTPvpgh3lMRm4ltJ9rOZaZCn3msw22iVUec5u+dpr9FM8jY6y7J
m4MB203AfRexnFwmL3FrIki4vylRmVp7HWMxQLLYBxBPPPEi7QVB6EE6GNKEf973
...
...
...
cy3sLxHa7+IGD/c4pnETQ+9Zna/wrHEyt8WDtFzJ4sL9f1u3peh2ip9OG+Gh6F6Y
SDK129oSj5FF/uY///zkzAneHZwDSuAbNJfJfrKMUDZEWA292DqGJI03MFaHlq2f
-----END RSA PRIVATE KEY-----
```

A certificate file in PEM format is as shown below:

```
-----BEGIN CERTIFICATE-----
MIIGATCCA+mgAwIBAgIJAL0seBEmD2e7MA0GCSqGSIb3DQEBCwUAMIGWMQswCQYD
VQQGEwJJTjELMAkGA1UECAwCTUgxDTALBgNVBAcMBFBVTkUxFTATBgNVBAoMDENv
...
...
...
 dklbn4GF8PmGiNJ5u2WEecfQi2QMOfPl1GcMlhx8aVeRVC3jEDC5W0VEe+4mqwNh
0ULBsWbOLwlTFt3LhEKMh30Duiw0utxaPDoT0Pu5H+L5+/qt50QqVIdHwgSw31ZM
DLnNGkDq670rSjatNncpSAc0VElO/TdhBxhpZxZgMSi40x5IN++dAjKnPbha5Tcz
eEuBSUs=
-----END CERTIFICATE-----
```

> **Note:** One pem file can contain multiple certificates and key. You can copy the certificate and key part into two separate files or concatenate key.pem and certificates.pm into s single file as may be needed, using text editors.

Work with NoSQLandHadoop Services using SSL tool

NoSQL and Hadoop services provide a CLI tool called imanis_cli.sh which can be found on the Cohesity nodes where NoSQL and Hadoop Service is installed. This tool is used to set up keystore and truststore on the NoSQL and Hadoop Service app. Following are the options to imanis_cli.sh -c ssl command.

Log in to one of the Cohesity cluster nodes where NoSQL and Hadoop Service is installed and run the following:

```
$ imanis_cli.sh -c ssl

Fetching ImanisData cluster config from athena-1:imanis-nodemanager-0

Usage :
imanis_cli.sh -c ssl import port hostname1 [ hostname2 hostname3 ...]
download certificate(s) from remote server(s) and
import them into Imanis Truststore automatically
arguments:
port   : SSL port where server is listening
hostname : list of hostnames of the servers to download SSL certificates

imanis_cli.sh -c ssl import-manual alias1 /path/to/certificate/file1 alias2
/path/to/certificate/file2 ...
Import SSL certificates manually into Imanis truststore
NOTE: You need to copy the certificates to one of the Imanis nodes manually
Use this only if the automatic import fails
arguments:
alias : short hostname from where the certificate is importedfile  : path to
certificate file
```

```
imanis_cli.sh -c ssl generate-keystore
Generate self-signed SSL certificates and use it for Imanis keystore
```

```
imanis_cli.sh -c ssl import-keystore /path/to/key/file /path/to/cert/file
Import CA signed SSL certificate+key into Imanis keystore,
Use this when self-signed certificate is not good enough
arguments:
/path/to/key/file : path to file containing private key in PEM format
/path/to/cert/file: path to corresponding certificate file in PEM format
NOTE: If you have a single file containing both key and certificate,
```

```
then repeat the same file path for both arguments
```

```
imanis_cli.sh -c ssl list [truststore|keystore]
List the certificates from Imanis Truststore/Keystore

imanis_cli.sh -c ssl delete alias
Delete the certificate for given alias from Imanis Truststore

imanis_cli.sh -c ssl delete [truststore|keystore]Delete the trust-
store/keystore, if it exists
```

Setting up truststore:

The truststore contains certificates from the SSL enabled servers to which we want to connect to. There are a couple of ways in which you can build the truststore:

a. **Importing certificates from servers automatically**:

   In this case, you run the imanis_cli.sh tool by specifying the port and the list of server hostnames or IP addresses. The tool will connect to the specified port on each server and download the SSL certificate presented by the server.

   IPv4 Example: you have three servers in your MongoDB cluster with IP address

   10.14.20.87, 10.14.20.88 10.14.20.89 and they are listening on port 27221

```
$ imanis_cli.sh -c ssl import  27221 10.14.20.87 10.14.20.88
10.14.20.89

Fetching ImanisData cluster config from athena-1:imanis-nodemanager-0
Downloading certificate from 10.14.20.87:27221 to
/tmp/10.14.20.87.crt.970
Downloading certificate from 10.14.20.88:27221 to
/tmp/10.14.20.88.crt.970
Downloading certificate from 10.14.20.89:27221 to
/tmp/10.14.20.89.crt.970
Importing SSL certificate into Imanis...
 truststore does not exist. A new truststore will be created.
you will be prompted for a password. Please ensure that
you remember this password as it will be required later
for importing more certificates to this Truststore.

Enter the Truststore password : ********

Re-Enter the Truststore password : ********
```

```
Importing /tmp/10.14.20.87.crt.970 to Imanis Truststore
Successfully imported /tmp/10.14.20.87.crt.970

Importing /tmp/10.14.20.88.crt.970 to Imanis Truststore
Successfully imported /tmp/10.14.20.88.crt.970

Importing /tmp/10.14.20.89.crt.970 to Imanis Truststore
Successfully imported /tmp/10.14.20.89.crt.970

Copying truststore to /opt/talena/conf/talena.truststore
```

IPv6 Example:

```
[support@cb146-00505684fff6-node-1 ~]$ imanis_cli.sh -c ssl import
18091 2400::3000:0:0:100:1025 2400::3000:0:0:100:1026
2400::3000:0:0:100:1027
INFO : Checking reachability of pod imanis-nodemanager-0
Downloading certificate from [2400::3000:0:0:100:1025]:18091 to
/tmp/2400::3000:0:0:100:1025.crt.39833
Downloading certificate from [2400::3000:0:0:100:1026]:18091 to
/tmp/2400::3000:0:0:100:1026.crt.39833
Downloading certificate from [2400::3000:0:0:100:1027]:18091 to
/tmp/2400::3000:0:0:100:1027.crt.39833
Importing SSL certificate into Imanis...
Enter the Truststore password : *******
Re-Enter the Truststore password : *******
Importing /tmp/2400::3000:0:0:100:1025.crt.39833 to Imanis Truststore
Multiple certificates found in /tmp/2400::3000:0:0:100:1025.crt.39833
Generating /tmp/certs.39833/2400::3000:0:0:100:1025.crt.39833-0 ..
Generating /tmp/certs.39833/2400::3000:0:0:100:1025.crt.39833-1 ..
Importing /tmp/2400::3000:0:0:100:1026.crt.39833 to Imanis Truststore
Multiple certificates found in /tmp/2400::3000:0:0:100:1026.crt.39833
Generating /tmp/certs.39833/2400::3000:0:0:100:1026.crt.39833-0 ..
Generating /tmp/certs.39833/2400::3000:0:0:100:1026.crt.39833-1 ..
Importing /tmp/2400::3000:0:0:100:1027.crt.39833 to Imanis Truststore
Multiple certificates found in /tmp/2400::3000:0:0:100:1027.crt.39833
Generating /tmp/certs.39833/2400::3000:0:0:100:1027.crt.39833-0 ..
Generating /tmp/certs.39833/2400::3000:0:0:100:1027.crt.39833-1 ..
Copying truststore to /opt/talena/conf/talena.truststore
Certificates imported successfully
[support@cb146-00505684fff6-node-1 ~]$
```

> **Note:**
> - The port should be same on all nodes.
> - The port number may be different on your setup. Consult with your administrators for your primary cluster.
> - For Hadoop, it may be necessary to use different ports for different hosts (Namenode port and Datanode port, for example). In that case, you need to run this tool multiple times.

b. **Importing certificates manually**:

Although automatic import is the preferred option as it reduces the chances of errors in manual download/import steps, in some cases, you may want to manually download the certificates from a remote location and import them into NoSQL and Hadoop Service truststore. In such cases, use the below option:

```
$ imanis_cli.sh -c ssl import-manual im2-vm4 $PWD/im20-
vm4.cohesity.com.cer im2-vm5 $PWD/im20-vm5.cohesity.com.cer im2-vm6
$PWD/im20-vm6.cohesity.com.cer
Fetching ImanisData cluster config from athena-1:imanis-nodemanager-0
Importing SSL certificate into Imanis...


Enter the Truststore password : ********

Re-Enter the Truststore password : ********

Importing /home_cohesity_data/athena/mounts/_internal_view_
mount/imanis/certs/im20-vm4.cohesity.com.cer to Imanis Truststore
Successfully imported /home_cohesity_data/athena/mounts/_internal_view_
mount/imanis/certs/im20-vm4.cohesity.com.cer

Importing /home_cohesity_data/athena/mounts/_internal_view_
mount/imanis/certs/im20-vm5.cohesity.com.cer to Imanis Truststore
Successfully imported /home_cohesity_data/athena/mounts/_internal_view_
mount/imanis/certs/im20-vm5.cohesity.com.cer

Importing /home_cohesity_data/athena/mounts/_internal_view_
mount/imanis/certs/im20-vm6.cohesity.com.cer to Imanis Truststore
Successfully imported /home_cohesity_data/athena/mounts/_internal_view_
mount/imanis/certs/im20-vm6.cohesity.com.cer
```

```
Copying truststore to /opt/talena/conf/talena.truststore
```

> **Note:** In most cases, you do not have to do this step, as a self-signed certificate is automatically generated, when you try to import any certificate into NoSQL and Hadoop Service truststore for the very first time.

Setting up keystore

A keystore must have a private key and a corresponding certificate. To set up the keystore:

a. **Generate a self-signed key+certificate pair**: This is the default setup. When you import a truststore for the very first time, the script will automatically generate a self-signed keystore, if it does not exist. A self signed certificate will also be generated and can be found at `$INSTALL_DIR/conf/talena.crt` on any of the NoSQL and Hadoop pod.

```
$ imanis_cli.sh -c ssl generate-keystore
Fetching ImanisData cluster config from athena-1:imanis-nodemanager-0
Generating random keystore password
Generating a keystore and self-signed SSL certificate for Imanis
Successfully configured keystore
```

b. **Import a key+CA signed certificate**

In some environments, the company security policy may mandate you to use the CA signed certificate instead of self-signed. In such a case, you need to use the import-keystore option to import a pair of key and CA signed certificate.

```
$ imanis_cli.sh -c ssl import-keystore $PWD/Imanis.key $PWD/Imanis.crt

Fetching ImanisData cluster config from athena-1:imanis-nodemanager-0
Verifying private key and certificate
Keystore already exists on Imanis nodes
Saving it as  /opt/talena/conf/talena.keystore.062220-050744 and /op-
t/talena/conf/talena.crt.062220-050744
Importing SSL KeyStore into Imanis...
You will be prompted to enter a password for the keystore
Please enter the password that was used for generating the key

Enter the Keystore password : *********

Re-Enter the Keystore password : *********
You will be prompted to enter a passphrase for the private key if there is
no passphrase, then just hit enter
Enter the Keystore passphrase : ******
```

```
Re-Enter the Keystore passphrase : ******

private key and certificate look good. Going ahead with import
Converting  /cohesity/mount/sharedview/imanis/certs-root/Imanis.key into
PKCS12 format /tmp/keystore.45617
Importing /tmp/keystore.45617 to /opt/talena/conf/talena.keystore
Picked up _JAVA_OPTIONS: -Djava.io.tmpdir=/tmp -Dimanis.namespace=athena-1 -
Dlog.rotation.dir=/cohesity/mount/sharedview/imanis/5.0.0-REL1002/rotated-
logs
Importing keystore /tmp/keystore.45617 to /op-
t/talena/conf/talena.keystore...
Entry for alias 1 successfully imported.
Import command completed:  1 entries successfully imported, 0 entries failed
or cancelled

Warning:
The JKS keystore uses a proprietary format. It is recommended to migrate to
PKCS12 which is an industry standard format using "keytool -importkeystore -
srckeystore /opt/talena/conf/talena.keystore -destkeystore /op-
t/talena/conf/talena.keystore -deststoretype pkcs12".
Successfully imported .

Restart Imanis ap-server, after importing the appropriate SSL certificates.

Would you like to restart now ? (yes/no) yes ? yes
Stopping API server...
API server  stopped successfully
Starting API server...
API server is running on [imanis-api-server-0]
```

Note:

- The key and certificate files should be placed under directory starting from path /home_cohesity_data/athena/mounts/_internal_view_mount/imanis/ (see the usage for this tool).
- Keystore passphrase is the private key passphrase that was set when the private key was generated. If it was not set, then just hit enter for the pass phrase.
- Restarting the api-server is mandatory for the changes to become effective.
- Ignore the warning about PKCS12 format.

Managing NoSQL and Hadoop Services

Upgrade NoSQL and Hadoop Services

After upgrading Cohesity to the latest version, to continue using the features provided by the NoSQL and Hadoop Service, you must upgrade to the latest version of the NoSQL & Hadoop Service.

To upgrade NoSQL and Hadoop services from an earlier release to the latest one:

Note: It is recommended that all scheduled "NoSQL and Hadoop Service" protection runs are paused before the upgrade. Also verify that no recovery is in progress.

> **Note:**
> NoSQL and Hadoop app upgrade from older versions to 5.2.1 is not supported. The app should be reinstalled.
>
> To reinstall app version 5.2.1, do the following:
>
> 1. Run the following command to save the configurations from the existing version:
>
>   # imanis_cli.sh -c config save
>
> 2. Uninstall existing app version
> 3. Install NoSQL and Hadoop Services 5.2.1
> 4. Launch NoSQL and Hadoop Services5.2.1

1. Download the App package that you want to upgrade.
2. Navigate to **Marketplace** > **My Apps**.
3. Select **My Apps** tab, and click **Upload App** on the top right corner.
4. Click **Select File** and browse the location for the **NoSQL and Hadoop Services App** and select it.
5. Click **Upload and Install**.

   The progress of the upload is displayed and once the upload is completed, the app is installed automatically.

   > **Note:** Please wait for the installation to complete and the version to be shown as the desired upgraded version.

6. Once the installation is complete, click **All Instances** and click **Pause** to pause the current active instance that is running.

   Wait for the Status to be shown as **paused**.

   > **Note:** To continue using data protection for MongoDB, you need to configure the MongoDB CDP service. For more information, see Managing NoSQL and Hadoop Services.

7. Click **Resume** to launch services using the latest upgraded version of NoSQL and Hadoop Services App.

> **Note:** Please wait for around 10-15 minutes for the status to be shown as running. Click on the drop down to see the upgraded version of NoSQL and Hadoop Services App.

8. Make sure that all NoSQL and Hadoop services are up and running.

Perform the following steps to upgrade the NoSQL app:

1. Save the current configuration:

   ```
   imanis_cli.sh -c config save
   ```

2. Pause the current active instance of the app.
3. Upload and install the new NoSQL app.
4. Resume the app after a successful installation.

   The new app instance displays.

5. When the new app is resumed, by default, it restores ssl and krb5 configurations from the config backup. But in case it fails, restore the config using:

   ```
   imanis_cli.sh -c config restore
   ```

   For more information, see Config Save and Restore.

6. To verify whether the app is in running state, use the following:

   ```
   kubectl.sh get pods -A | grep "imanis"
   ```

   Ensure all pods are in running state.

7. Perform a smoke test before running any jobs.
   1. Log in to any one of the imanis pods.
   2. Run the following and make sure all tests PASS.

      ```
      talena-services.sh smoke all
      ```

Install a Patch for NoSQL and Hadoop Service

This section describes the procedure of installing a patch for NoSQL and Hadoop Services.

1. Download the patch, example PATCH-00.tar.gz.
2. Copy the file at location `/home_cohesity_data/athena/mounts/_internal_view_mount/imanis`.
3. Run the patch application utility from a Cohesity node as a *cohesity* user.

   ```
   # imanis_cli.sh -c patch -f /home_cohesity_data/athena/mounts/_
   internal_view_mount/imanis/PATCH-159.tar.gz
   ```

4. Verify if the patch is applied to the system once the patch is successfully applied:

```
# imanis_cli.sh -c patch -i
```

Scale Nodemanager Pods

Sometimes, you may want to increase (scale up) the number of nodemanager replicas to use the additional resources available after adding a new compute node. In such cases, you can utilize the additional resources by launching more nodemanager replicas using the command given below:

```
$ imanis_cli.sh -c scale --help
INFO : Checking reachability of pod imanis-nodemanager-0
-------------------------------------------------------------------------
Performing scale operation of node manager pods.
-------------------------------------------------------------------------
Usage:
To scale up number of nodemanager pods,
imanis_cli.sh -c scale up --nmcount
To scale down number of nodemanager pods,
imanis_cli.sh -c scale down --nmcount
--nmcount: Number of expected node manager pods.
-------------------------------------------------------------------------



$ imanis_cli.sh -c scale up 4
INFO : Checking reachability of pod imanis-nodemanager-0
-------------------------------------------------------------------------
Performing scale operation of node manager pods.
-------------------------------------------------------------------------
NM pods scaled up to 4 successfully.
INFO: Scale up performed successfully
-------------------------------------------------------------------------
```

You can also reduce the number of running nodemanager pods by reducing the replica count using the following command.

```
$ imanis_cli.sh -c scale down 2
INFO : Checking reachability of pod imanis-nodemanager-0
-------------------------------------------------------------------------
Performing scale operation of node manager pods.
-------------------------------------------------------------------------
NM pods scaled down to 2 successfully.
INFO: Scaled down performed successfully
-------------------------------------------------------------------------
```

Configuring Continuous Data Protection (CDP) for MongoDB

Ensure that the NoSQL and Hadoop Service app is installed and running.

To enable MongoDB CDP service, use the following command:

```
$ imanis_cli.sh -c mongocdp enable --cpu <cpu-cores> --memory <memory-GB> --
replica <num-replicas>
```

- **cpu-cores**: Number of CPU cores allocated to each MongoCDP service instance
- **memory-GB**: Amount of memory allocated to each MongoCDP service instance
- **num-replicas**: Number of MongoCDP service instances. Depending on the load on MongoDB primary and the recovery point objective, one or more MongoDB CDP services can be provisioned. The use of multiple smaller CDP services can help RPO in case of nodes going down.

For example, the below command enables MongoDB CDP with single replicas of 1core/2GB memory.

```
$ imanis_cli.sh -c mongocdp enable --cpu 1 --memory 2 --replica 1
INFO : Checking reachability of pod imanis-nodemanager-0
INFO : Log File : /tmp/imanis-cli-workdir.10686/cli/imanis-mongo-cdp.log
INFO : Enabling MongoCDP pods deployment
INFO : Reading allocatable cluster resources
INFO : Total allocatable Memory(GB) : 239
INFO : Total allocatable CPUs : 47
INFO : Number of MongoCDP replicas : 1
INFO : CPUs per MongoCDP replica : 1
INFO : Memory per MongoCDP replica(GB) : 2
INFO : Total CPUs for MongoCDP replicas : 1
INFO : Total Memory for MongoCDP replicas(GB): 2
INFO : Total Memory for Other Services(GB) : 15
INFO : Total CPUs for Other Services : 10
INFO : Small-Nodemanager instance config : 16GB/4 CPUs
INFO : Medium-Nodemanage instance config : 32GB/8 CPUs
INFO : Large-Nodemanager instance config : 40GB/10 CPUs
Please confirm above configurations for 'NoSQL & Hadoop Service' app
deployment with MongoCDP Pods
Do you want to continue (yes/no) yes ? yes
```

```
--------------------------------------------------------------------------------
----
mongocdp : service(s) will be stopped (if running)
--------------------------------------------------------------------------------
----
service : mongocdp is being stopped
Stopping MongoCDP services on imanis-mongocdp-0
MongoCDP services stopped successfully
Setting up MongoCDP services...
Configuring MongoCDP services...
MongoCDP services configured successfully
Deleting stale pod directory:
/cohesity/mount/sharedview/imanis/cdp/mongo/v1/pods/imanis-mongocdp-1
Starting MongoCDP service on imanis-mongocdp-0
MongoCDP service is running on [imanis-mongocdp-0]
MongoCDP service is running
Setting up MongoCDP service successfull
Requested number of pods with given memory and cpu cores is updated
successfully
```

To verify, once the app is Running, run the following command to check the pod status:

```
$ kubectl.sh  get pods  -A
```

Example:

```
$ kubectl.sh get pods -A
NAMESPACE                NAME
 READY   STATUS    RESTARTS   AGE
athena-6                 imanis-api-server-0
 1/1     Running   0          102m
athena-6                 imanis-api-server-1
 1/1     Running   0          102m
athena-6                 imanis-historyserver-0
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-0
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-1
 1/1     Running   0          102m
athena-6                 imanis-nodemanager-2
 1/1     Running   0          102m
```

```
athena-6                  imanis-nodemanager-3
 1/1     Running   0          102m
athena-6                  imanis-nodemanager-4
 1/1     Running   0          102m
athena-6                  imanis-resourcemanager-0
 1/1     Running   0          102m
athena-6                  imanis-resourcemanager-1
 1/1     Running   0          102m
athena-6                  imanis-zookeeper-0
 1/1     Running   0          102m
athena-6                  imanis-zookeeper-1
 1/1     Running   0          102m
athena-6                  imanis-zookeeper-2
 1/1     Running   0          102m
cohesity-agentinstall     agentinstall-controller-0
 2/2     Running   0          12d
cohesity-coredns          coredns-8566b9b44-4hrrk
 1/1     Running   0          12d
cohesity-coredns          coredns-8566b9b44-97nvj
 1/1     Running   0          12d
cohesity-coredns          coredns-8566b9b44-sxrw5
 1/1     Running   0          12d
cohesity-infraoperator    infra-operator-controller-manager-787dbddd76-n7tcc
 2/2     Running   0          12d
cohesity-infraoperator    infra-operator-controller-manager-787dbddd76-qqd96
 2/2     Running   0          12d
cohesity-magnus           magnus-actor-f8b5fbcc-q6fgb
 1/1     Running   0          12d
cohesity-magnus           magnus-server-0
 1/1     Running   0          12d
```

To disable MongoDB CDP:

1. Run the following command:

   $ imanis_cli.sh -c mongocdp disable

2. Check the pod status from Cohesity node with command "kubectl.sh get pods -A" and verify that no mongocdp pods are running.

> **Note:** Once MongoCDP service is disabled, MongoDB data protection would no longer work.

Uninstall NoSQL and Hadoop Services

To terminate and uninstall the NoSQL and Hadoop services, see Cohesity Helios documentation > MarketPlace Apps > Uninstall an App for details.

Additional Configuration for NoSQL and Hadoop Services

This topic describes some of the additional configuration that can be done to NoSQL and Hadoop Service.

Configure NoSQL and Hadoop Services to Use Updated VIPs

Use the following commands to list and configure the latest VIPs.

Commands	Description
`imanis_cli.sh -c vip`	Provides VIP option to update required config files with latest VIPs.
`imanis_cli.sh -c vip list`	Lists existing VIPs used by NoSQL and Hadoop Services.
`imanis_cli.sh -c vip update`	Configure NoSQL and Hadoop Services to use the latest list of VIPs. **Note:** You need to restart the service for the VIP changes.

Megafile Support for Cassandra

Megafile allows fast backup and recovery for large files, helping reduce backup times as well as downtime when recovering from a disaster. Megafile distributes files across all nodes in a cluster. Megafile breaks large files into smaller chunks for parallel backup and recovery across nodes. The minimum size of these chunks is optimized to maximize performance.

Cassandra backup and recovery of size tiered tables generally contains data files which are in tens or hundreds of GBs. These data files are backed up or recovered as Megafiles. Cassandra uses different approaches for backup and recovery. The Megafile approach is used only if `scp` is being used for backup and recovery.

By default, Megafile support for Cassandra backup and recovery is disabled.

Enable or Disable Megafile for Cassandra Backup

Run the following command to enable Megafile for Cassandra backup:

```
imanis_cli.sh -c megafile --connector cassandra-backup --workdir /tmp/imanis
--enable true
```

where, `-workdir` specifies the directory on the primary cluster nodes where Cassandra can copy the binaries.

Run the following command to disable Megafile for Cassandra backup:

```
imanis_cli.sh -c megafile  --connector cassandra-backup  --enable false
```

Enable or Disable Megafile for Cassandra Recovery

A file is restored as a Megafile only if it is backed up as a Megafile. Otherwise, Cassandra uses the normal recovery workflow to recover Cassandra objects.

Run the following command to enable Megafile for Cassandra recovery:

```
imanis_cli.sh -c megafile --connector cassandra-restore --workdir /tm-
p/imanis --enable true
```

where, `-workdir` specifies the directory on the primary cluster nodes where Cassandra can copy the binaries.

Run the following command to disable Megafile for Cassandra recovery:

```
imanis_cli.sh -c megafile --connector cassandra-restore --enable false
```

Megafile Status

The pulse log shows the following when Megafile is enabled for backup and recovery:

```
- Megafile support is enabled
```

The pulse log shows the following after a protection run or recovery:

```
CASSANDRA_MEGAFILES_COPIED=2 CASSANDRA_MEGAFILES_FAILED=2 CASSANDRA_
MEGAFILES_DATA_COPIED=4931999383
```

where

- `CASSANDRA_MEGAFILES_COPIE`D - Number of Megafiles successfully copied.
- `CASSANDRA_MEGAFILES_FAILED` - Number of Megafiles for which copy failed.
- `CASSANDRA_MEGAFILES_DATA_COPIED` - Total data copied as Megafile.

Job Log Collection Tool

The Job Log Collection tool is used to collect logs for a specific protection run for a specific duration.

Note:

- You must log in to the Cohesity cluster as the Support User
- The Support User must have write permission to `outdir`.
- Ensure the `outdir` has enough space to collect the protection run logs.

Perform the following steps to collect the protection run logs:

1. Log in to the Cohesity cluster where NoSQL and Hadoop Service app is running.
2. Run the following command to check if the NoSQL and Hadoop Service is running:

```
$ kubectl.sh get pods -o wide -A | grep imanis
```

3. Run the following command to collect the protection run logs for a specific duration:

```
imanis_cli.sh -c joblog --jobname <jobname> --outdir <dirname> [--
fromdate <fromdate>] [--todate <todate>]
```

Example:

```
$imanis_cli.sh -c joblog --jobname Hdfs-backup1 --outdir /XYZ --fromdate
2021_09_01
```

Add or Delete Host Entries to /etc/hosts on all Pods

You can add or delete host entries in the `/etc/hosts` file on all pods.

Important:

- Do not manually edit the `/etc/hosts` file. You must use `imanis_cli.sh` command to add or delete host entries in the `/etc/hosts` file.
- Any manual edits made into the `/etc/hosts` file are lost after pod restart or app upgrade.

Add Host Entries

Perform the following steps to add host entries to the `/etc/hosts` file:

1. Log in to the Cohesity cluster as the Support User. The cluster must have the NoSQL & Hadoop Service app running.
2. Prepare a file with the host entries at the location `/home_cohesity_data/athena/mounts/_internal_view_mount`.

> **Note:**
> - The host files must have new line separated entries for hosts in <IP><Hostname> format. For example, 10.14.25.101 server1
> - The host files path must always begin with `/home_cohesity_data/athena/mounts/_internal_view_mount`.

Example:

```
$ cat /home_cohesity_data/athena/mounts/_internal_view_mount/newhosts
10.14.25.101 server1
10.14.25.102 server2
10.14.25.103 server3
```

3. Run the following command to add the host entries:

```
$ imanis_cli.sh -c hosts add <hosts-file-path>
```

Example:

```
$ imanis_cli.sh -c hosts add /home_cohesity_data/athena/mounts/_
internal_view_mount/newhosts
```

Delete Host Entries

Perform the following steps to delete host entries in the `/etc/hosts` file:

1. Log in to the Cohesity cluster as the Support User. The cluster must have the NoSQL & Hadoop Service app running.
2. Prepare a file with the host entries at the location `/home_cohesity_data/athena/mounts/_internal_view_mount`.

> **Note:**
> - The host files must have new line separated entries for hosts in <IP><Hostname> format. For example, 10.14.25.101 server1
> - The host files path must always begin with `/home_cohesity_data/athena/mounts/_internal_view_mount`.

Example:

```
$ cat /home_cohesity_data/athena/mounts/_internal_view_mount/newhosts
10.14.25.101 server1
```

```
10.14.25.102 server2

10.14.25.103 server3
```

3. Run the following command to delete the host entries:

   ```
   $ imanis_cli.sh -c hosts delete <hosts-file-path>
   ```

   Example:

   ```
   $ imanis_cli.sh -c hosts delete /home_cohesity_data/athena/mounts/_
   internal_view_mount/newhosts
   ```

List Host Entries

You can also list the host entries available in the `/etc/hosts` file.

To list the host entries, run the following command:

```
imanis_cli.sh -c hosts list
```

Explorer Log Collection

Explorer tool is used to collect the NoSQL & Hadoop Service app specific logs from pods, CHDFS and Cohesity nodes.

> **Important:**
> - The Support User must have write permission to `dest_dir`.
> - Ensure the `dest_dir` has enough space to collect the logs.

Perform the following steps to collect the logs:

1. Log in to the Cohesity cluster as the Support User. The cluster must have the NoSQL & Hadoop Service app running.
2. Run the following command:

   ```
   imanis_cli.sh -c explorer [-n days] [-p string] [-f] [-s list] dest_dir
   ```

   where

   - `-n`: collect log files which are <= n days old. Applicable only for logs, system.

     > **Note:** Mutually exclusive with either full or incremental backup.

   - `-f`: full backup of logs and hdfs:/system files (incremental by default).
   - `-s list`: subset of explorers to run (comma separated). Valid values are:

    - devops
    - chdfs
    - logs
    - machinfo
    - system
    - config
- `-p string`: prefix string to the output tar ball name.
- `dest_dir`: destination path where explorer output is written.

Example:

```
$ imanis_cli.sh -c explorer -n 10 /home/cohesity/data/timecapsules/
```

Config Save and Restore

After terminating the app, any modification made to any of the config files is lost. This tool allows you to restore configs from previous instances of an app if it is saved with the save option.

- config save is used to save config and patches from the running imanis instance.
- config restore is used to restore config/patches from the previous instance to the current running instance.

Save Config and Patches

Perform the following steps to save config and patches from the running instance:

1. Log in to the Cohesity cluster where the NoSQL & Hadoop Service app is running.
2. Run the following command:

   ```
   $ imanis_cli.sh  -c config
   ```

3. Run the following command to save config and patches:

   ```
   $ imanis_cli.sh  -c config save
   ```

Restore Config and Patches

Perform the following steps to restore config and patches to the running instance:

1. Log in to the Cohesity cluster where the NoSQL & Hadoop Service app is running.
2. Run the following command:

   ```
   $ imanis_cli.sh  -c config
   ```

3. Run the following command to restore config and patches:

   ```
   $ imanis_cli.sh  -c config restore
   ```

Optionally you can also restore specific saved config or configs. Run the following command to restore specific saved config or configs to the running instance:

```
$  imanis_cli.sh -c config restore conf hadoop conf agents
```

Monitor the NoSQL and Hadoop Service Health

To ensure that the NoSQL and Hadoop Service is healthy, Cohesity constantly monitors the following services running on the Kubernetes pods: `zookeeper`, `YARN`, `api-server`, and `mongocdp` (if enabled for MongoDB protection). These services are required to enable and perform the backup and restore operations.

If any of these services encounter issues, an error message with an error code appears on the Cohesity UI, which can be used for troubleshooting. The status of the NoSQL and Hadoop Service is displayed as either **Running** or **Not Healthy** depending on the status of these services. The status of the NoSQL and Hadoop Service specifies the following:

- **Running** - The NoSQL and Hadoop Service is healthy and running.
- **Not Healthy** - The NoSQL and Hadoop Service is not healthy and have some issues.

To view the health of the NoSQL and Hadoop service from the Coheity UI, navigate to **Marketplace** > **My Apps** and then click the **All Instances** tab.

Running State:

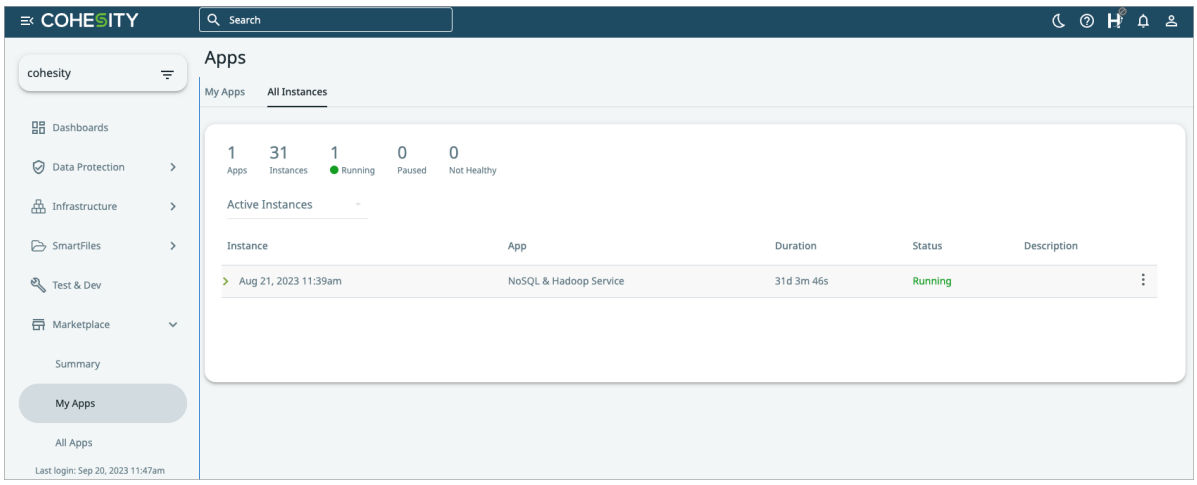

Not Healthy State:

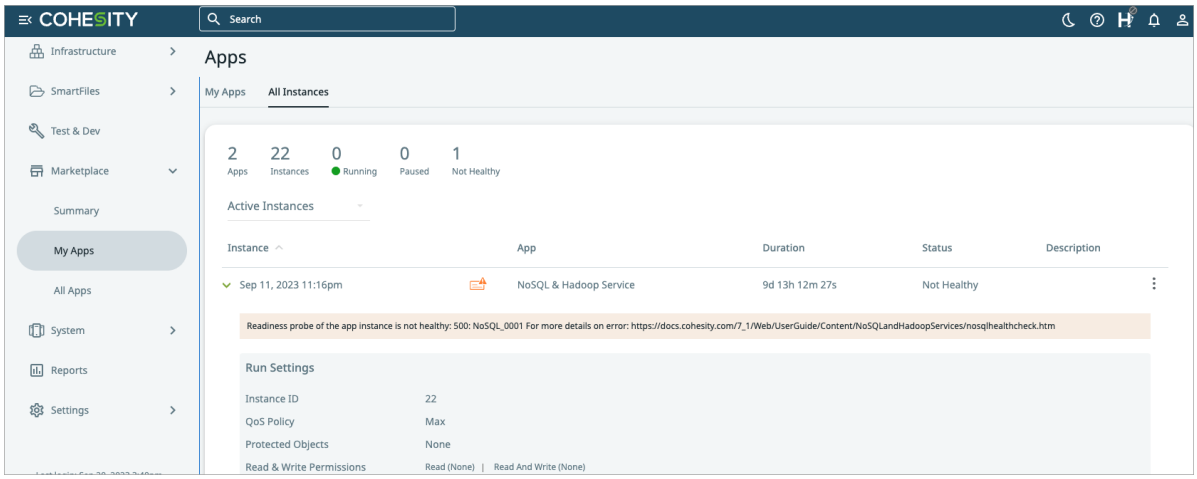

NoSQL and Hadoop Service Error Codes

This section covers the NoSQL and Hadoop Service error codes and the troubleshooting steps.

> **Important:** Ensure that you have Host Shell access on the Cohesity cluster to perform the troubleshooting steps.

NoSQL_0000

Issue:

The NoSQL and Hadoop service health check feature cannot report the health of the service. The NoSQL and Hadoop service might be in a **Running** or **Not Healthy** state, and the health check feature cannot detect the service's current status.

Troubleshooting:

Contact Cohesity Support if the issue persists.

NoSQL_0001

Issue:

The NoSQL and Hadoop service, `zookeeper`, is down on one or more of the pods: `imanis-zookeeper-0`, `imanis-zookeeper-1`, and `imanis-zookeeper-2`.

Troubleshooting:

1. Login to the pod using the following command:

   ```
   kubectl.sh exec -it -n <athena-namespace> imanis-zookeeper-0 --
   /bin/bash
   ```

2. Run the following command to find the pod on which the `zookeeper` service is down:

```
talena-services.sh status zookeeper
```

3. Log in to the pod on which the `zookeeper` service is down and verify `zookeeper` log files in `$LOG_DIR` to determine the reason for the issue.

   For example: The log files to look for, if `zookeeper` is down on `imanis-zookeeper-0` pod:

   ```
   $LOG_DIR/zookeeper-hdfs-server-imanis-zookeeper-0.log
   $LOG_DIR/zookeeper-hdfs-server-imanis-zookeeper-0.out
   ```

4. Also, you can collect the logs from the Cohesity cluster for further debugging by running the following command using the `cohesity` user on the cluster:

   ```
   imanis_cli.sh -c explorer -n 1 -s logs -l zookeeper /home/cohesity
   ```

5. Try restarting the `zookeeper` service to fix the issue. Run the following command on any of the pods to restart the `zookeeper` service:

   ```
   talena-services.sh restart zookeeper
   ```

6. Check the status of the `zookeeper` service to verify if the issue is resolved using the following command:

   ```
   talena-services.sh smoke zookeeper
   ```

Contact Cohesity Support if the issue persists.

NoSQL_0002

Issue:

The NoSQL and Hadoop group of services `YARN`, is down on one or more of the mentioned pods: `imanis-resourcemanager-0`, `imanis-resourcemanager-1`, `imanis-nodemanager-*`, and `imanis-historyserver-0`.

Troubleshooting:

1. Login to the pod using the following command:

   ```
   kubectl.sh exec -it -n <athena-namespace> imanis-nodemanager-0 --
   /bin/bash
   ```

   > **Note:** You can use `imanis-resourcemanager-0` instead of `imanis-nodemanager-0`.

2. Run the following command to find the pod on which the `YARN` service is down:

   ```
   talena-services.sh status yarn
   ```

3. Login to the pod on which the `YARN` service is down and verify the log files specific to yarn (`*nodemanager*`, `*resourcemanager*`, `*historyserver*`) in `$LOG_DIR` to

determine the reason for the issue.

For example: The log files to look for, if `resourcemanager` is down on `imanis-resourcemanager-0`:

```
$LOG_DIR/hadoop-hdfs-resourcemanager-imanis-resourcemanager-0.log
$LOG_DIR/hadoop-hdfs-resourcemanager-imanis-resourcemanager-0.out
```

Also, you can collect the logs from the Cohesity cluster for further debugging by running the following command using the `cohesity` user on the cluster:

```
imanis_cli.sh -c explorer -n 1 -s logs -l yarn /home/cohesity
```

4. Try restarting the `YARN` service to fix the issue. Run the following command on any of the pods to restart the `YARN` service:

```
talena-services.sh restart yarn
```

5. Check the status of the `YARN` service to verify if the issue is resolved using the following command:

```
talena-services.sh smoke yarn
```

Contact Cohesity Support if the issue persists.

NoSQL_0003

Issue:

The NoSQL and Hadoop group of services `api-server`, is down on one or more of the mentioned pods: `imanis-api-server-0` and `imanis-api-server-1`.

Troubleshooting:

1. Login to the pod using the following command:

```
kubectl.sh exec -it -n <athena-namespace> imanis-api-server-0 --
/bin/bash
```

> **Note:** You can use `imanis-api-server-1` instead of `imanis-api-server-0`.

2. Run the following command to find the pod on which the `api-server` service is down:

```
talena-services.sh status api-server
```

3. Login to the pod on which the `api-server` service is down and verify the `api-server` log files in `$LOG_DIR` to determine the reason for the issue.

   For example: The log files to look for, if `api-server` is down on `imanis-api-server-0`:

```
$LOG_DIR/api-server-imanis-api-server-0.log
$LOG_DIR/api-server-imanis-api-server-0.log
```

Also, you can collect the logs from the Cohesity cluster for further debugging by running the following command using the `cohesity` user on the cluster:

```
imanis_cli.sh -c explorer -n 1 -s logs -l api-server /home/cohesity
```

4. Try restarting the `api-server` service to fix the issue. Run the following command on any of the pods to restart the `api-server` service:

```
talena-services.sh restart api-server
```

5. Check the status of the `api-server` service to verify if the issue is resolved using the following command:

```
talena-services.sh smoke api-server
```

Contact Cohesity Support if the issue persists.

NoSQL_0004

Issue:

The NoSQL and Hadoop group of services `mongocdp`, is down on one or more mongocdp pods. For example `imanis-mongodp-0`, `imanis-mongodp-1` etc.

Troubleshooting:

1. Login to the pod using the following command:

```
kubectl.sh exec -it -n <athena-namespace> imanis-mongocdp-0 --
/bin/bash
```

2. Run the following command to find the pod on which the `mongocdp` service is down:

```
talena-services.sh status mongocdp
```

3. Log in to the pod on which the `mongocdp` service is down and verify the `mongo*` log files in `$LOG_DIR` to determine the reason for the issue.

   For example: The log files to look for, if `mongocdp` is down on `imanis-mongocdp-0`:

```
$LOG_DIR/mongodb-cdp-imanis-mongocdp-0.log
$LOG_DIR/mongodb-cdp-imanis-mongocdp-0.out
```

   Also, you can collect the logs from the Cohesity cluster for further debugging by running the following command using the `cohesity` user on the cluster:

```
imanis_cli.sh -c explorer -n 1 -s logs -l mongocdp /home/cohesity
```

4. Try restarting the `mongocdp` service to fix the issue. Run the following command on any of the pods to restart the `mongocdp` service:

```
talena-services.sh restart mongocdp
```

5. Check the status of the `mongocdp` service to verify if the issue is resolved using the following command:

```
talena-services.sh smoke mongocdp
```

Contact Cohesity Support if the issue persists.

NoSQL_0011

Issue:

The NoSQL and Hadoop service `zookeeper` smoke test has failed.

Troubleshooting:

1. Log in to the pod using the following command:

```
kubectl.sh exec -it -n <athena-namespace> imanis-zookeeper-0 --
/bin/bash
```

2. Verify `zookeeper` log files in `$LOG_DIR` to determine the reason for the issue.

   For example: The log files to look for:

```
$LOG_DIR/zookeeper-hdfs-server-imanis-zookeeper-0.log
$LOG_DIR/zookeeper-hdfs-server-imanis-zookeeper-0.out
```

   Also, you can collect the logs from the Cohesity cluster for further debugging by running the following command using the `cohesity` user on the cluster:

```
imanis_cli.sh -c explorer -n 1 -s logs -l zookeeper /home/cohesity
```

3. Try restarting the `zookeeper` service to fix the issue. Run the following command on any of the pods to restart the `zookeeper` service:

```
talena-services.sh restart zookeeper
```

4. Check the status of the `zookeeper` service to verify if the issue is resolved using the following command:

```
talena-services.sh smoke zookeeper
```

Contact Cohesity Support if the issue persists.

NoSQL_0012

Issue:

The NoSQL and Hadoop service `chdfs` smoke test has failed.

Troubleshooting:

1. Collect the logs from the Cohesity cluster for debugging by running the following command using the `cohesity` user on the cluster:

```
imanis_cli.sh -c explorer -n 1 -s logs -l yarn /home/cohesity
```

2. Try restarting the `YARN` service to fix the issue. Run the following command on any of the pods to restart the `YARN` service:

   ```
   talena-services.sh restart yarn
   ```

3. Check the status of the `YARN` service to verify if the issue is resolved using the following command:

   ```
   talena-services.sh smoke chdfs
   ```

Contact Cohesity Support if the issue persists.

NoSQL_0013

Issue:

The NoSQL and Hadoop service `mongocdp` smoke test has failed. The `mongocdp` service runs on pods `imanis-mongocdp-*`.

Troubleshooting:

1. Log in to the pod using the following command:

   ```
   kubectl.sh exec -it -n <athena-namespace> imanis-mongocdp-0 --
   /bin/bash
   ```

2. Verify the `mongo*` log files in `$LOG_DIR` to determine the reason for the issue.

   For example: The log files to look for, if `mongocdp` is down on `imanis-mongocdp-0`:

   ```
   $LOG_DIR/hadoop-hdfs-resourcemanager-imanis-resourcemanager-0.log
   $LOG_DIR/hadoop-hdfs-resourcemanager-imanis-resourcemanager-0.out
   ```

   Also, you can collect the logs from the Cohesity cluster for further debugging by running the following command using the `cohesity` user on the cluster:

   ```
   imanis_cli.sh -c explorer -n 1 -s logs -l mongocdp /home/cohesity
   ```

3. Try restarting the `mongocdp` service to fix the issue. Run the following command on any of the pods to restart the `mongocdp` service:

   ```
   talena-services.sh restart mongocdp
   ```

4. Check the status of the `mongocdp` service to verify if the issue is resolved using the following command:

   ```
   talena-services.sh smoke mongocdp
   ```

Contact Cohesity Support if the issue persists.

NoSQL_0020

Issue:

The NoSQL and Hadoop service's health check fails if the disk usage goes above 95% of the disk space.

Troubleshooting

Contact Cohesity Support if the issue persists.

NoSQL_0022

Issue:

There is a problem while fetching and applying configuration parameters. Refer to the `talena_getGflags.log-<date-time>` log file in the NoSQL App logs for failure details.

The api-server may be running or may not be healthy. The health check feature has some issues determining the current status.

Troubleshooting

1. Log in to the pod using the following command:

   ```
   kubectl.sh exec -it -n <athena-namespace> imanis-api-server-0--
   /bin/bash
   ```

2. Check the configuration fetch log files in $LOG_DIR to identify the cause of the issue:

   ```
   $LOG_DIR/talena_config_fetch<date-time>.log
   ```

3. Try restarting the api-server to re-fetch and apply the configuration. Run the following command on any of the pods to restart the api-server service:

   ```
   talena-services.sh restart api-server
   ```

NoSQL_0023

Issue:

This error occurs when there is a problem updating the core-site.xml configuration file. The update failure may happen during:

- Adding or updating VIP configurations (automated in the background)
- Modifying Kerberos-related properties (imanis_cli.sh -c kerberos)

The error is generated when the automated script responsible for updating the core-site file fails and requires manual intervention to resolve.

Troubleshooting

1. Connect to the API server pod using the following command:

   ```
   kubectl.sh exec -it -n <athena-namespace> imanis-api-server-0 --
   /bin/bash
   ```

2. Navigate to the rotated log directory and examine the update VIPs logs.

   ```
   cd $ROTATED_LOG_DIR
   ```

   ```
   ls -la update-vips*.log
   ```

3. Analyze the log files to identify the root cause of the failure.

   Contact Cohesity Support for assistance with log analysis.

4. Identify and fix the underlying issue with the core-site.xml file based on the log analysis.

   Contact Cohesity Support for further assistance.

5. Collect the log files followed by clean up.

After fixing the issue:

1. Gather log files that contain reported errors.

   Use the following command to list files marked with the retain flag:

   ```
   export LOG_RETAIN_MARKER="IMANIS_RETAIN_LOG"; find $ROTATED_LOG_DIR -
   name "update-vips*" -type f -mmin +10 -exec sh -c 'grep -qi "'"$LOG_
   RETAIN_MARKER"'" "$1" && echo "$1"' _ {} \;
   ```

   These logs can be collected for further investigation.

2. Delete update-VIP log files.

   ```
   rm -f $ROTATED_LOG_DIR/update-vips*.log
   ```

> **Important:** If these log files are not deleted after resolving the issue, the health check will incorrectly identify the core-site.xml file as unhealthy.

FAQs on NoSQL and Hadoop Services

The following are some of the frequently asked questions about managing SSL enabled NoSQL primary clusters:

1. What is alias which is used while importing the certificates ? Does it have to match the hostname ?

   Alias is just a name or a tag for the entry in the truststore. It is used only by users to identify the certificate they are looking for. It could be any unique string for a given certificate entry. Usually the hostname of the machine to which the certificate belongs is preferred as it makes it easy to identify the same.

2. Can I import the certificates into NoSQL and Hadoop Service in PKCS12 format ?

   No, the certificates should be in PEM format for importing into NoSQL and Hadoop service.

3. My primary server certificate is signed by an intermediate CA and the intermediate CA certificate is signed by another root CA. So do I need to import these two certificates as well into NoSQL and Hadoop service?

Yes, root CA, intermediate CA certificates should also be imported into NoSQL and Hadoop service truststore. If you import using an automatic method, then usually all the certificates in the chain are downloaded from the server and imported into NoSQL and Hadoop service truststore.

4. How do I know what format my certificate is in ?

   If you cat the file and find text wrapped in headers such as `-----BEGIN CERTIFICATE----- or -----BEGIN RSA PRIVATE KEY-----`, then the certificate is in PEM format.

5. How do I verify if my private key and certificate match each other ?

   ```
   openssl x509 -noout -modulus -in /path/to/cert | openssl md5
   ```

   ```
   openssl rsa -noout -modulus -in /path/to/key   | openssl md5
   ```

   If both the sums match, then they match each other.

   > **Note:** When importing a keystore into NoSQL and Hadoop service, this check is done by the script.

6. How do I check if I have already imported a certificate into NoSQL and Hadoop Service truststore ?

   ```
   # imanis_cli.sh -c ssl list truststore
   ```

   and check if the entry exists.

7. I am importing a CA signed certificate and the corresponding private key into NoSQL and Hadoop service as keystore. What should be the hostname (CN) ?

   The hostname for NoSQL and Hadoop Service app usually start with `imanis-*.imanis-cluster.*.svc.cluster.local`. But these hostnames are not resolvable outside the Athena app. So for hostname, you can use the Cohesity hostnames on which NoSQL and Hadoop Service app has been deployed.

8. My server certificates have expired. How do I remove the old certificates and re import new ones ?

   Using imanis_cli.sh -c ssl list command, you can list the entries and delete the ones which you want using the specific alias as follows:

   ```
   # imanis_cli.sh -c ssl delete alias
   ```

9. If I delete keystore or truststore while the protection runs are in progress, what happens ?

   Protection runs may fail. This should be done only when you are sure that there are no protection runs connecting to SSL enabled primaries on your NoSQL and Hadoop cluster.

10. What is the difference between password and passphrase for the private key I am importing ? How do I verify if my password and passphrase are correct before importing ?

    Password is set for the keystore or truststore that you import into NoSQL and Hadoop service. Please set a good password and remember it.

    Passphrase is always associated with a private key. It is usually set when you generate a private key. The passphrase you enter should match with the passphrase that was set for the given key. You can not change this afterwards.

11. When do I need to import certificates into NoSQL and Hadoop Service?

    You need to do it before you register your data source into Cohesity UI.

12. Can I only import certificates from my server into NoSQL and Hadoop service? Is importing or generating keystore mandatory ?

    In some environments, there is no requirement for the client to present its certificate to server and server having to carry out the validation. In such environments, you do not need to import NoSQL and Hadoop service certificate into your primary cluster.

    However, in some cases, like in cassandra, the setting required_client_auth=true in yaml configuration file mandates that validation be done both ways. In this case, you have to copy the certificate from NoSQL and Hadoop node into your primary nodes and import it using the methods prescribed by your respective primary setup.

13. Which one should I use ? A self-signed certificate or a CA signed certificate for NoSQL and Hadoop Service?

    This depends on the security requirements of your organization. Self-signed certificate is good enough, if the requirement is to encrypt the data transferred between the NoSQL and Hadoop running cluster and your primary cluster being protected. But if the requirement is not only encrypted data transmission on the wire but also to validate that the server and client are authentic, then you need CA signed certificates.

Chapter 31: Test & Dev (Cloning)

This chapter includes the following topics:

- Test & Dev Overview 4564
- Clone Views 4565
- Clone VMs 4567
- Clone Databases 4583
- Test & Dev Using CloudSpin 4589
- Clone vCloud Director (vCD) VMs 4589

Test & Dev Overview

Cohesity empowers developers to instantiate the latest backup of their production application stack and run it directly off Cohesity, providing a unified foundation for copy data management. Instant, zero-space clones enable businesses to quickly spin up test/dev environments from a backup or existing file share, enabling rapid test and development from actual data without any capacity overhead. Because clones are designed with performance in mind, access time is identical regardless of the number of snapshots or clones that are made.

Cohesity provides the ability to clone objects from snapshots created by a Protection Group. The cloned objects are copied to a new location, rehydrated and restored to a running state. For example when a VMware object is cloned in an on-premises Cohesity cluster, new VM files (such as VMDKs) are created from snapshots and stored in a view on the Cohesity cluster. This view becomes the datastore for these VM files. The Cohesity cluster creates new VMs in the selected resource pool by mounting the view containing the new VMDKs as the datastore on the ESXi hosts. During cloning, VMware is responsible to select the ESXi hosts that have access to the required Datastore based on load balancing.

Note the following about Cohesity Cloning:

- When cloning in an on-premises Cohesity cluster, in a single Clone task you can only clone Protection Groups and VMs that belong to the same Storage Domain and Source.
- You can select only one Snapshot of a VM to clone. While configuring a Clone task, you cannot select two Snapshots that were created from the same VM.

For instructions on how to use Cohesity to Clone VMs, see Clone VMs.

The following example depicts cloning Snapshots of VMware VMs. You want to clone a set of Engineering VMs for use by the QA department for testing. You can select a mix of individual Snapshots and Job Snapshots to clone. In the figure below, the following Snapshots are cloned:

- The individual Snapshot created of VM1 on 03/21/2025 03:00AM
- The individual Snapshot created of VM3 on 03/21/2025 05:10AM

After the testing is complete, you may want to tear down the cloned VMs. For more information, see About Tearing Down Cloned VMs.

The action steps in the following procedure are represented in the figure below with blue numbered squares:

.

T1) From Snapshots, the Cluster creates new VMDKs in the selected View as shown by VMDK4 and VMDK5. This View becomes the Datastore for these VM files.

T2) The Cohesity cluster creates new VMs by mounting the View containing the new VMDKs from the Datastore onto ESXi Hosts in the selected Resource Pool. In this

example, VM4 and VM5 are created from mounting VMDK4 and VMDK5 onto the ESXi Host QA.

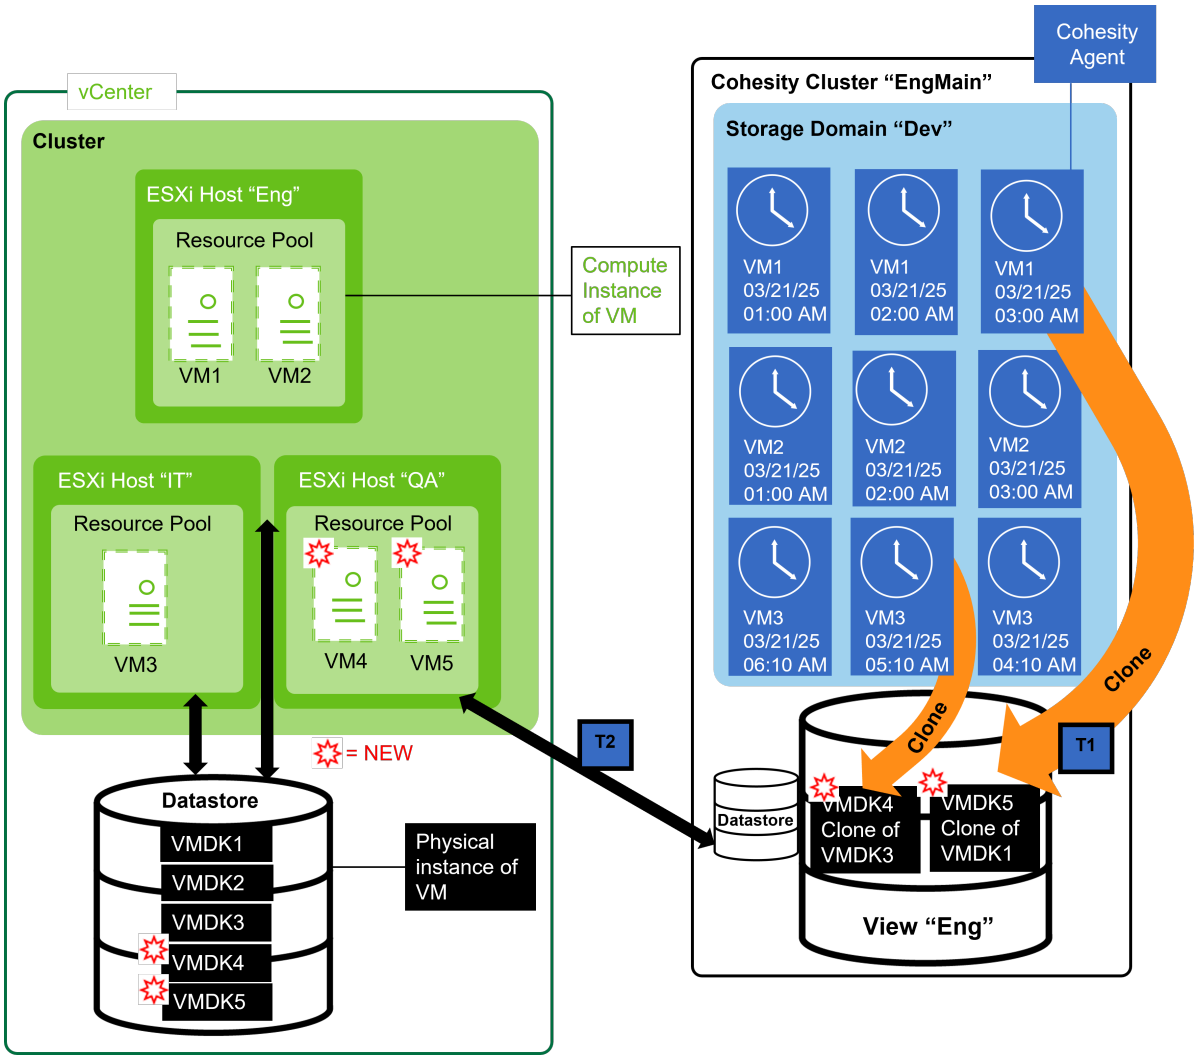

Clone Views

You can make a clone of an existing view. After cloning, the two versions of the data are unique. Changes made to the data in the new view are not replicated in the original view and changes made to the original view are not replicated in the new view. Data within a DataLock view is immutable. You can also clone a view from a view Snapshot. For more information, see Clone a View From a Snapshot.

Before you can clone Protected Objects (Views), Snapshots must exist. To create Snapshots, create and run a Protection Group.

To clone a view:

1. Navigate to **Test & Dev**.
2. Click **Clone** located at the top right of the page.
3. Click **Views**.
4. Select Views to Clone.
5. Set Clone Options.

Select Views to Clone

In this section you select Views to clone.

Before you can clone Protected Objects (Views), Snapshots must exist. To create Snapshots, create and run a Protection Group.

To select Views to clone:

1. Search for Views to clone. Enter characters of the View name and a list of items displays. (You can optionally specify the wildcard character * in the search.) Select an item from the list. You can also optionally narrow the search results by specifying filter criteria. For example, you can filter the search results by a specific Protection Group. Click **Change Filters**, specify Filter by criteria and click **Add**.
2. Click a View.

Set Clone Options

In this section you select the options for the clone task.

To set clone options:

1. Select the clone's **Purpose**.

   Select **Clone as DataLock View** if you want to clone a view as write once read many. Data within a DataLock view is immutable, and most view properties cannot be edited. (Note that view-level DataLock is different than file-level DataLock, as described in File-level DataLock for NFS and SMB Views.)
2. In the **Task Name** field, you can optionally change the default name for the clone task.
3. **View Name:** If you are creating or cloning a new view, specify a name for the view. Forward slashes (/) are not allowed in the view name.

   > **Important:** Renaming a view changes its default share name. Clients won't be able to mount the view unless they use the new share name. Clients using a share alias are not affected by the rename.

For a DataLock view, select a date for lock expiration. During the lock period, data within the view cannot be edited by any user.

During the lock period, users with Data Security privileges can delete the view (unless File DataLock is enabled for the view, in which case the view cannot be deleted), extend the lock period (it cannot be reduced) , add IPs in the allowlist and change the description. After the lock period expires, the view can be deleted by any user that has permission to delete it.

4.
5. The Storage Domain is already set. The new view is always created in the same storage domain as the original view.
6. **QoS Policy:** A QoS Policy determines to which storage media data is written and the priority of I/O when contention occurs. Select an appropriate QoS policy for the view. By specifying a priority and the type of workload you expect, the Cohesity cluster can optimize how it processes data in regard to latency, throughput, and priority.

   For details and recommendations, see QoS Policies for Cohesity Views.
7. To have privileges to mount and access a view, a system must be specified in the allowlist subnet. For more information, see Add Subnets to Allowlists.
8. Select the Snapshot to clone. At the bottom of the page, the View that will be cloned is listed. By default, the latest Snapshot of the View is cloned. You can optionally click Edit ( ) for the Selected Object and choose the exact Snapshot to clone.

After cloning, the two versions of the data are unique. Changes made to the data in the new view are not replicated in the original view and changes made to the original view are not replicated in the new view. Data within a DataLock view is immutable.

> **Tip:** The cluster provides notifications when clone and recover tasks that you initiated complete. To view them, click the notifications icon in the top right corner of the Cohesity UI. You can drill down on tasks for additional details and clear notifications as needed.

Clone VMs

Cohesity provides the ability to clone objects from snapshots created by a Protection Group. The cloned objects are copied to a new location, rehydrated and restored to a running state. For example when a VMware object is cloned in an on-premises Cohesity cluster, new VM files (such as VMDKs) are created from snapshots and stored in a view on the Cohesity cluster. This view becomes the datastore for these VM files. The Cohesity cluster creates new VMs in the selected resource pool by mounting the view containing the new VMDKs as the datastore on the ESXi hosts. During cloning, VMware is responsible to select the ESXi hosts that have access to the required Datastore based on load balancing.

To learn more, see Test & Dev Overview and About Tearing Down Cloned VMs.

> **Tip:** You can also watch the VMware VM Test & Dev (Cloning) video to learn how to clone VMware objects.

You can clone or migrate any DataProtect VM to any Cohesity supported cloud.

Note the following about Cloning VMs:

Prerequisite

For Windows VMs, prepare the VM that is ready for cloning by running the Microsoft Sysprep (System Preparation) tool. This process generalizes the virtual disk and generates a new machine SID.

Considerations:

- Attached CDROM or floppy drives are not recovered or cloned. When recovering or cloning a Snapshot of a VM that had an attached CDROM or floppy drive with data (such as an ISO file), when the VM is recovered or cloned the attached drives are not created and the data is not available.
- Cloning or recovering VMs may fail if the maximum number of mounted NFS Datastores on a VMware ESXi Host is exceeded. Till 6.7, the default value is 8 and from 7.0, the default value is 32. If an ESXi Host already has the maximum number of NFS datastores mounted, the Cohesity View (datastore) does not mount on the ESXi Host during the Cloning or Recovery process. For more information including instructions on increasing the limit, see Increasing the default value that defines the maximum number of NFS mounts on an ESXi/ESX host (2239).
- During VM Recovery or the Cloning of VMs, the Cohesity cluster stores the VMDK and VMX files in different directories on the Cohesity datastore (View), due to a VMware vSphere issue. The VMDK files are placed in the *taskname*/*vmname* directory and the VMX files are placed in the *vmname* directory, where *taskname* is the Task Name you specify for the Recover task and the *vmname* is the name of the VM. For example, if a Clone is created with the Task Name of `QA` that creates clones of the VMs called `Build` and `Test`, the VMDK and VMX files are stored in the following directories in the Cohesity datastore (View): `QA/Build/Build.vmdk`, `QA/Test/Test.vmdk`, `Build/Build.vmx` and `Test/Test.vmx`. However, when recovering a VM, after bringing up the VM on the Cohesity datastore (View), the Cohesity cluster does a Storage vMotion migration of the VM back to its original primary datastore. After the Storage vMotion migration completes, both the VMX and VMDK files end up in the *vmname* directory and the VMDK files are no longer stored in the *taskname* directory. This VM recovery example assumes that the VMDK files were originally located in the default directory (and not in some special directory) when the backup was performed.

- For some conditions, the VMware vSphere reports an error that there is a mismatch between the actual Content ID of Parent VM and the Parent Content ID stored by the Child VM. If you see this error, follow the instructions to fix this issue provided in the following VMware Knowledge Base article: Resolving the CID mismatch error: The parent virtual disk has been modified since the child was created (1007969).
- During the recovery of a single VM that has multiple networks or datastores, VMware does not check if the VM and the ESXi Host's network and datastore are compatible when assigning the VM to the ESXi Host. This can cause the recovery to fail.
- Due to a VMware issue, specific VM recovery scenarios can cause storage migration (and recovery) to fail.
    - Using default settings to recover a VM to the same datastore. If the recovered VM belongs to a local datastore, migration may fail because VMware does not register the local datastore to all ESX Hosts within the Resource Pool.
    - Using custom settings to recover a VM to a selected Resource Pool and datastore. If the recovered VM belongs to a datastore that is not mounted on all ESX Hosts in the Resource Pool, migration may fail because VMware does not check the compatibility of the ESX Host (within the Resource Pool) with the given datastore while registering the VM.
- When the Cohesity cluster clones a VM, the Cluster makes a request to the vCenter Server to create the new VM. The vCenter Server selects an ESXi host from the specified resource pool to host the new VM. However, the VMware vSphere version of the ESXi host, which was selected by the vCenter Sever, may be incompatible with the VMware vSphere version of the original VM. When this happens, the clone of the VM fails with the following error: `A component of the virtual machine is not accessible on the host.`
- If VMs with the same VM name (but different relative VMDK paths) are cloned, the Cohesity cluster retains the original VM name for one VM and renames the rest by adding *_X* to the end of the VM names, where *X* is an increasing number such as 1, 2, or 3. For example if you clone three VMs called QA, the Cluster names the VMs as QA, QA_1 and QA_2. However, cloning VMs with the same relative path is not supported.
- Cloning not supported with the same relative path and VMDK file name. You cannot clone VMs that have the same relative path and the same VMDK name from two different primary datastores. For example, Protection Groups create Snapshots of `test.vmdk` stored in `qa/testing` on `datastore1` and `test.vmdk` stored in `qa/testing` on `datastore2`. If you attempt to Clone these two Snapshots, the Clone fails with an error message. The Clone fails because relative path and VMDK names are used to store both VMDKs in the Cohesity datastore and the Cluster cannot store them in same location. In addition, if two VMs have reference the same Hard disk, the Clone also fails. For example if VM1 included a Hard Disk that was then removed (not deleted) from VM1 and added to VM2. These two VMs cannot be Cloned because the relative path and reference to the Hard Disk is the same.

- Linked clones of VMware Virtual Machines can be backed up, recovered and cloned by the Cohesity cluster but the linking between the VMs is not preserved. The Cohesity cluster supports the Backup, Recovery and Cloning of linked clones of VMware Virtual Machines. However, the linking between the VMs is not preserved and full Snapshots of the VMs are created. In a single Recover task, you can recover a single linked clone VM but not multiple linked clone VMs.
- Cloning VMs within a vApp is not currently supported; they will be cloned under a specified Resource Pool similar to regular VMs.
- vApp VMs can be recovered to their original location (under the same vApp where the original VM that was backed up was present). When vApp VMs are recovered to an alternate location, they can be recovered to the specified Resource Pool only and the user must then move the VMs to the desired vApp.
- Currently, when a VM with multiple NICs is cloned or recovered to an alternate location, Cohesity creates the VM with only one NIC. This will be fixed in a future release.
- For VMware VMs, Virtual RDM disks are recovered as flat disks.
- Cohesity supports recovery only from VM to VM. Recovering VM to physical server and vice-versa is not supported.
- A workaround is needed for Hyper-V VM recovery in SCVMM configurations using test and development (cloning) scenarios, since the initial recovery may report issues with unsupported configurations. To avoid this, the Cohesity file share needs to be added to the recognized file shares for the host/cluster on SCVMM. Once the file share is successfully configured and a VM refresh is completed, the unsupported configuration issue of the VM is resolved. If a clone view is already added to known file shares and is used to perform subsequent clones, this workaround is not necessary.

  > **Note:** The file share must adhere strictly to the following format:
  > \\<clustervip>\<storageDomain>.<cloneViewName>.fs
  > Example file share path: \\dp-test.domain.com\DefaultStorageDomain.new.fs

Before you can clone Protected Objects (VMs), Snapshots must exist. To create Snapshots, create and run a Protection Group.

To learn more, see Replicate and Clone VMs in AWS Cloud Edition and Clone VMs from On-Premises to AWS Cloud Service (On Demand CloudSpin) .

Best Practice for Cloning VMs

Cohesity recommends the following best practice for cloning a VM:

Do not delete the existing VM source while cloning.

To Clone VMs:

1. Verify the following requirements prior to cloning:

   Cloning Snapshots to an Azure Cloud is supported for VMware and Hyper-V VMs with the following operating systems:

   - Linux CentOS 7 on LVM
     - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
     - When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

       ```
       UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4 defaults,nofail 0 0
       ```

       For detailed instructions, see Edit /etc/fstab to use UUID.
   - Windows
     - You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
     - For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.
     - When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. The Cohesity agents installed must be of the same version as the Cohesity Cluster version. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.
     - Due to a Windows issue in builds 14393 and build 15063, a deployed Windows 10 or Windows 2016 VM may not be connected to the network. Resolve this by installing fix KB4073562 for build 14393 or KB4016240 for build 15063 or by updating Windows. For more information, see Cannot connect remotely to a Windows 10 or Windows Server 2016 VM in Azure because of netvsc.sys.
     - Make the disks online by default using the following SAN policy commands:
       - `diskpart`
       - `san policy=OnlineAll`

       For more information, see the SAN Policy documentation.
   - Red Hat Enterprise Linux

- You can clone as many disks as the selected VM supports. For example, if you select a Standard_DS1_v2 VM, four disks can be cloned.
- For the list of supported RHEL versions, see Supported Software.
- When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the Protection Group, as shown by the following example entry in a `/etc/fstab`:

  ```
  UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
  defaults,nofail 0 0
  ```

  For detailed instructions, see Edit /etc/fstab to use UUID.

- The disk size of the VMware VMs must be a multiple of 1 GB.
- Cohesity does not support the CloudSpin of VMs whose OS Disk is formatted with XFS version 5 to Azure. This limitation applies to both Hyper-V and VMware VMs, regardless of their operating system type.
- Cohesity does not support CloudSpin of VMware VMs from VMware Cloud (VMC) to Azure.

Cloning Hyper-V VMs to Azure

- VMs that have one to five disks can be cloned. Each disk must be smaller than 2 TB.
- SUSE Linux VMs are *not* supported.
- All the Hyper-V versions supported by Cohesity for backup is supported for cloning to Azure. For the list of supported Hyper-V versions, see Supported Software.
- Generation 2 VMs are not supported by all VM sizes in Azure. For information on Generation 2 supported VM series, see Microsoft document.
- Cloning of Hyper-V VMs to Azure Stack Hub is not supported.

Cloning VMware VMs to an AWS Cloud has the following requirements:

- Cloning snapshots to an AWS Cloud is supported for VMware VMs with specific Linux operating systems. For information on supported operating systems, see Supported Software
- Cloning Snapshots to an AWS Cloud is supported for VMware VMs with the Windows operating system.
- The disk size of the VMware VMs must be a multiple of 1 MB.
- VMs that have one to ten disks can be cloned.
- When converting and cloning or copying Snapshots of a Linux VM with multiple disks, the disks must be mounted by UUID prior to being backed up by the

Protection Group, as shown by the following example entry in a `/etc/fstab`:

```
UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
defaults,nofail 0 0
```

For detailed instructions, see Edit /etc/fstab to use UUID.

For AWS, be aware of the following requirements:

- In addition, the fstab entry for **CentOS 7.4**, **RHEL 7.4**, **RHEL 6.7**, **SUSE 11** and **Ubuntu 16** operating systems must specify the `nofail` option as shown in the following example:

  ```
  UUID=e12e623a-c6ee-4718-9b1a-802456e72258 /userdata ext4
  defaults,nofail 0 0
  ```

- In addition, the fstab entry for the **Ubuntu 14** operating system must specify the `nobootwait` option as shown in the following example:

  ```
  UUID=2b262696-41a6-47c9-b173-d02717ee0e10 /mnt/expvol ext4
  defaults,nobootwait 0 0
  ```

- When backing up Windows VMs with multiple disks, before a Protection Group captures the Snapshot, the Cohesity agent and VMware tools must be installed on the VM and the VM must be running. For more information, see Install and Manage the Agent on Windows Servers. In addition, port 50051 must be open in your Firewall. For more information, see Manage Firewall Ports.

  > **Note:**
  > **For Convert and Deploy**: It is mandatory to install the Cohesity agent on the VMs. Windows does not show the data disks if the Cohesity agent is not installed, irrespective of the SAN/Disk Policy.
  >
  > **For Disaster Recovery**: It is mandatory to install the Cohesity agent on multi-disk Windows VMs. Windows does not show the data disks if the Cohesity agent is not installed, irrespective of the SAN/Disk Policy. In addition, the Protection Group capturing the Snapshot must have **Cloud Migration** enabled. For detailed instructions, see Add or Edit a Protection Group for Virtual Servers.

- For a Windows VM, if you want to access the newly created VM in the Cloud service using Remote Desktop, enable the Windows Remote Desktop on the VM prior to capturing the Snapshots of a Windows VM.

2. Navigate to **Test & Dev**.
3. Click **Clone** located at the top right of the page.
4. Click **VMs**.

5. Select VMs to Clone.
6. Set Clone Options.

Select VMs to Clone

In this section you select VMs to clone.

Before you can clone Protected Objects, Snapshots must exist. To create Snapshots, create and run a Protection Group.

To select VMs to clone:

1. Search for VMs or Protection Groups that contain the Snapshots to clone. Enter characters of the VM name or Protection Group name and a list of items displays. (You can optionally specify the wildcard character * in the search.) Select an item from the list. You can also optionally narrow the search results by specifying filter criteria. For example, you can filter the search results by a specific Protection Group. Click **Change Filters**, specify Filter by criteria and click **Add**.
2. Select a Protection Group or a VM.

   After selecting a VM or protection run, only protection runs and VMs with Snapshots stored in the same Storage Domain are enabled for selection and an automatic locked filter for the Storage Domain appears. Select a Protection Group or a VM.
3. You can optionally select more protection runs or VMs and click **Add to Cart**.
4. You can also optionally edit the cart and remove selected items by clicking the **Cart Icon** ( ).
5. Click **Continue**.

Set Clone Options

In this section you select the options for the clone task.

1. In the **Task Name** field, you can optionally change the default name for the clone task.
2. Select snapshots to clone. The set of VMs that will be cloned are listed. By default, the latest snapshots of the VMs are cloned. You can optionally click Edit ( ) for each selected object and choose the exact snapshot to clone for each VM. To view a VM's consistency information, click Edit. If one type displays, this means all VMs in the snapshot are that type. If the Mixed type displays, this means the snapshot contains VMs of different types, such as crash-consistent and not powered on. You can display individual VMs to see each VM's type.

3. **Preserve Tags**: Toggle this option to preserve the tags in the cloned VMs from the source.
4. **Preserve Custom Attributes**: Toggle this option to preserve the custom attributes in the cloned VMs from the source.
5. Under **Clone Location**, in the **Source** drop-down, select a registered source to clone the VMs to. The options displayed depend on the selected source type.
   vCenter Source—Select a vCenter server to clone VMs to.
   a. In the **Resource Pool** drop-down, select a resource pool to store the cloned VMs.
   b. Optionally, in the **VM Folder** drop-down, select a VM folder in the datacenter in which to create the cloned VMs.

      > **Note:** Not applicable when cloning to a standalone ESX host.

   c. In the **View** drop-down, create a new view by entering a new name or select an existing view. Only views in the same storage domain as the selected objects to clone are listed in the drop-down. A Cohesity view provides a storage location with NFS, SMB, and S3 mount paths in a storage domain on the Cohesity cluster. You can use the NFS protocol to mount the NFS mount path onto systems. You can mount the SMB mount path onto systems as an SMB share. You can use the view to store data such as files, backup snapshots, and cloned VMs. Views can only be mounted on any machines specified in the allowlist subnets. A view can also act as a datastore during the VM cloning process and during VM recovery. You can select the SWIFT protocol and a create SWIFT-only view for OpenStack environments. You can create a view for the following three categories:File Shares Backup Target Object Services You can select the **Read/Write** protocol and, optionally, select the **Read-only** protocol. You can create a Cohesity view only by using a template. The template can be a predefined template or a custom template. For more information, see:Create Views Using a Predefined TemplateCreate Views Using a Custom TemplateIf you are using Kerberos (AD or MIT), you must create a view using the same storage domain that is mapped to AD/MIT-Kerberos + LDAP.

      > **Note:** If the VMware VM is cloned to a new View, by default, the new View is created with Restore High QoS policy. The priority of this QoS policy is the same as TestAndDev High.

   vCloud Director Source—Select a vCloud Director source to the VMs to.
   a. In the **Source** drop-down, select a resource pool in vCloud Director to store the cloned VMs.

b. In the **Organization VDC** drop-down, select an organization VDC for the cloned VMs.

c. In the **vApp** drop-down, select the vApp for the cloned VMs.

d. In the View drop-down, create a new view by entering a new name or select an existing view.

Hyper-V Source—Select a Hyper-V source to clone VMs to.

a. In the **Compute** drop-down, select a Hyper-V host to store the cloned VMs.

> **Note:** Not applicable when cloning to a standalone Hyper-V host.

b. In the **View** drop-down, create a new view by entering a new name or select an existing view. Only views in the same storage domain as the selected objects to clone are listed in the drop-down. A Cohesity view provides a storage location with NFS, SMB, and S3 mount paths in a storage domain on the Cohesity cluster. You can use the NFS protocol to mount the NFS mount path onto systems. You can mount the SMB mount path onto systems as an SMB share. You can use the view to store data such as files, backup snapshots, and cloned VMs. Views can only be mounted on any machines specified in the allowlist subnets. A view can also act as a datastore during the VM cloning process and during VM recovery. You can select the SWIFT protocol and a create SWIFT-only view for OpenStack environments. You can create a view for the following three categories:File Shares Backup Target Object Services You can select the **Read/Write** protocol and, optionally, select the **Read-only** protocol. You can create a Cohesity view only by using a template. The template can be a predefined template or a custom template. For more information, see:Create Views Using a Predefined TemplateCreate Views Using a Custom TemplateIf you are using Kerberos (AD or MIT), you must create a view using the same storage domain that is mapped to AD/MIT-Kerberos + LDAP.

Azure Subscription Source—Select a Subscription ID in Microsoft Azure to clone the VMs to.

a. **Resource Group**—Select a resource pool in Microsoft Azure to store the cloned VMs.

b. **Compute**—Select the Microsoft Azure VM Size Type such as **Standard_D1** to use when creating the cloned VMs.

c. **Availability Set**—Optionally, select the availability set to which the VM has to be cloned. For more information, see Azure Availability Sets.

d. **Storage Account**—Select the Microsoft Azure VM Size to use when creating the cloned VMs.

e. **Storage Container**—Select the Microsoft Azure storage container to store the cloned VMs.

AWS Source—Select an AWS Source to clone the VMs to.

a. In the **Region** drop-down, select an Amazon region where the new cloned VM will run, for example: `us-west-1`. For a list of supported AWS Regions, see Regions and Availability Zones.

b. In the **Instance Types** drop-down, select instance type (such as `m4.2xlarge`) of the new cloned VMs. The instance type defines the CPU, memory, storage and network capacity of the new VMs. For more information, see Amazon EC2 Instance Types.

6. Specify the networking options. The options displayed depend on the selected source type.

vCenter Source

- **Detach network**—For each cloned VM, detach the virtual network interface cards (NICs) from the VM. For each cloned VM, keep the virtual network interface cards (vNIC) disconnected from the network
- **Attach to a new network**—For each cloned VM, attach the vNIC to a new network. The new network can be a standard port group or a distributed port group.

> **Note:** The network options displayed depend on the clone location options that were selected.

vCloud Director Source

- **Detach network**—For each cloned VM, detach the virtual network interface cards (NICs) from the VM. For each cloned VM, keep the virtual network interface cards (vNIC) disconnected from the network
- **Attach to a new network**—For each cloned VM, attach the vNIC to a new network. The new network can be a standard port group or a distributed port group.

> **Note:** The network options displayed depend on the clone location options that were selected.

Auto Select: When enabled, the recovery task automatically selects the correct VLAN.

Azure Subscription Source

a. In the **Virtual Network** drop-down, select the Microsoft Azure Virtual Network for the cloned VMs.

b. In the **Subnet** drop-down, select the Microsoft Azure Virtual subnet for the cloned VMs.

AWS Source

a. In the **VPC** drop-down, select the AWS Virtual Private Cloud that will contain the cloned VMs.

b. In the **Subnets** drop-down, select the subnet for the cloned VMs.

c. In the **Network Security Group** drop-down, select the security group to apply to the instances of the new cloned VM.

d. When cloning a VMware VM that has 1 TB or greater disks, Cohesity launches a temporary EC2 instance and copies data from the Cohesity cluster to the instance. The instance can use the same VPC and subnet that's used for the clone, or you can provide a different VPC and subnet, but they must be reachable from the cluster.

e. You can optionally add **Tags** to the resources. These tags are key-value pairs that will be assigned to all the temporary and permanent resources created during Cloudspin.

7. Toggle **Rename Cloned VMs** on if you want to add Prefix and/or Suffix strings to the VM names of the new VMs created by this task. For example, you could specify Prefix `Test_` and Suffix `_QA` to the VM name `WinOS` and when the VM is cloned, the new VM that is created is called `Test_WinOS_QA`.
8. Deselect **Leave cloned VMs powered off** if you want the cloned VMs to be powered on after they are created.
9. Select **Continue clone even if errors occur when cloning VMs** if you want the clone task to continue even if errors occur when cloning the VMs.
10. **Auto Select**: When enabled, the cloning task automatically selects the correct VLAN.
11. **Interface Group**: Select a configured interface group from the Interface Group drop-down.
12. **Exclude Disk Handling**: This option is under Recovery Options, and when selected, the following two options are displayed:
    - **Create Excluded Disk as an Empty Disk**: Enable this option if you want Cohesity to create blank disks for every disk you excluded from the backup. The metadata of the excluded disks will be used to create the blank disks. These blank disks are available post-recovery and only contain the metadata of the original disk, but no data. You can configure and run databases and applications on these disks, saving you from the effort of creating new disks. This feature is especially helpful if you have many VMs with excluded disks.
    - **Create Independent Disk as an Empty Disk**: By default, independent disks are not backed up. Enable this option if you want Cohesity to create blank disks for such independent disks. These blank disks are available post-recovery and only contain the metadata of the original disk, but no data.

> **Note:** Excluded physical RDM disks cannot be created as blank disks, even when disk exclusion handling is enabled.

When cloning to a cloud service such as AWS or Azure, the cloning process detaches the network before the clone and reattaches the network after the clone completes. After the VM is running in the cloud, you may need to reapply settings to the network.

> **Tip:** The cluster provides notifications when clone and recover tasks that you initiated complete. To view them, click the notifications icon in the top right corner of the Cohesity UI. You can drill down on tasks for additional details and clear notifications as needed.

About Tearing Down Cloned VMs

Cohesity provides the ability to clone objects from snapshots created by a Protection Group. The cloned objects are copied to a new location, rehydrated and restored to a running state. For example when a VMware object is cloned in an on-premises Cohesity cluster, new VM files (such as VMDKs) are created from snapshots and stored in a view on the Cohesity cluster. This view becomes the datastore for these VM files. The Cohesity cluster creates new VMs in the selected resource pool by mounting the view containing the new VMDKs as the datastore on the ESXi hosts. During cloning, VMware is responsible to select the ESXi hosts that have access to the required Datastore based on load balancing.

Cohesity also provides the ability to tear down (delete) the set of VMs created by a clone task. The teardown procedure of a VMware VM clone is described below.

You create Clones of VMs by creating a clone task. For example, you might clone a set of development VMs for QA to use for testing as shown in the following figure.

VMs Created by a Clone Task

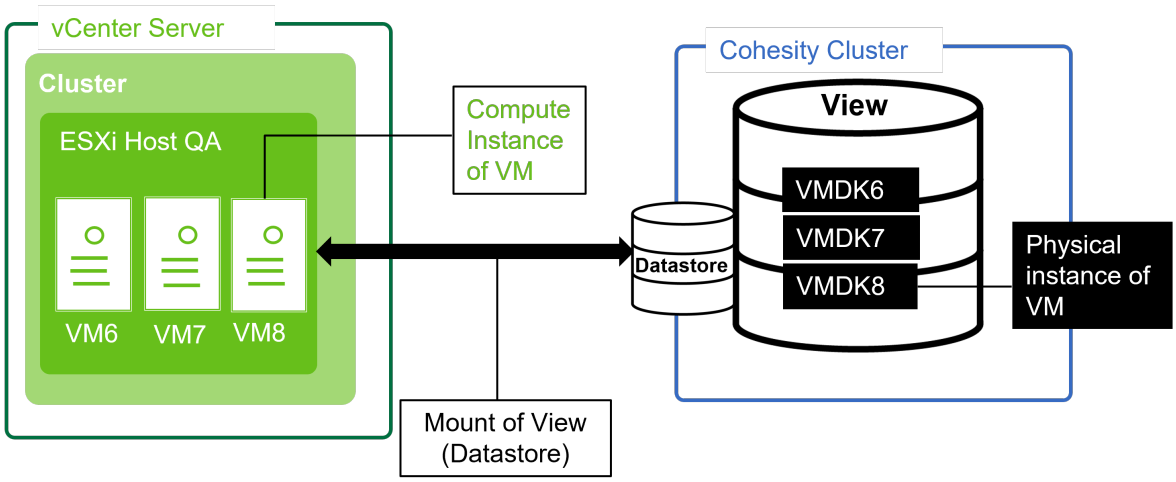

After the testing is completed, you may want to tear down the cloned VMs created by the clone task. You can initiate a clone teardown as described in Tear Down Clones.
The actions that occur during a teardown are described in the following procedure:

1. The Cohesity cluster deletes the compute instances of the cloned VMs in the ESXi host.
2. The VM files (such as VMDK and VMX) are deleted on the view acting as a datastore.
3. If no VMs are using the view (datastore), the view is umounted from all ESXi hosts.

The steps in the preceding procedure are represented in the following figure with blue numbered squares: T1 .

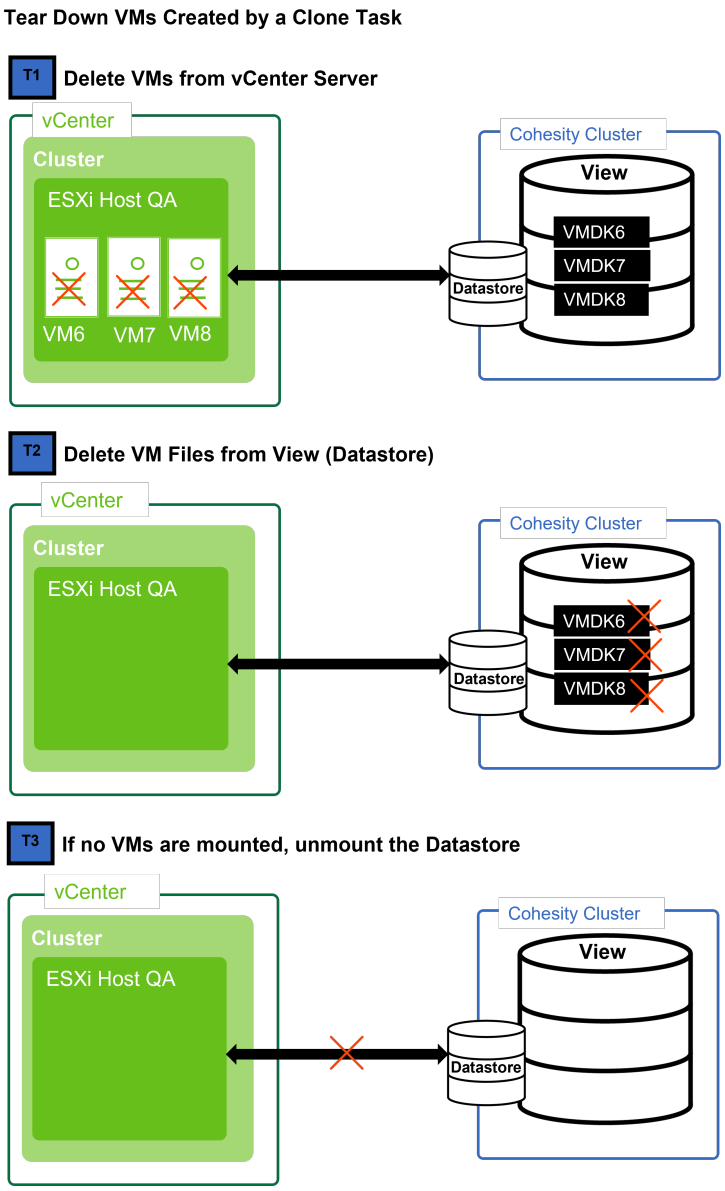

If a cloned VM has been altered after being created by a clone task, the VM may not be automatically cleaned up during a clone teardown. For more information, see VM Not Deleted if Cloned VM is Altered.

Tear Down Clones

Cohesity provides the ability to clone objects from snapshots created by a Protection Group. The cloned objects are copied to a new location, rehydrated and restored to a running state. For example when a VMware object is cloned in an on-premises Cohesity cluster, new VM files (such as VMDKs) are created from snapshots and stored in a view on the Cohesity cluster. This view becomes the datastore for these VM files. The Cohesity cluster creates new VMs in the selected resource pool by mounting the view containing the new VMDKs as the datastore on the ESXi hosts. During cloning, VMware is responsible to select the ESXi hosts that have access to the required Datastore based on load balancing.

Cohesity also provides the ability to tear down (delete) the set of VMs created by a clone task. To learn more, see About Tearing Down Cloned VMs.

> **Note:** The teardown of an MS SQL clone database removes it from the SQL server; the user does not have to do this beforehand. However, in the rare instance you created other additional databases with files on the mounted drive that the cloned database is on, those additional databases should be deleted first through the SQL Server Management Studio.

Tear Down a Clone

1. Navigate to **Test & Dev**.
2. Click a Task Name link in the table.
3. Click **Tear Down Clone** located at the top right of the page.

The clone teardown process of a VMware VM clone completes the following steps:

1. The Cohesity cluster deletes the compute instances of the cloned VMs in the ESXi host.
2. The VM files (such as VMDK and VMX) are deleted on the view acting as a datastore.
3. If no VMs are using the view (datastore), the view is umounted from all ESXi hosts.

After the teardown, the clone task remains visible in the Cohesity UI. If a clone teardown fails and does not delete all the VMs or VM files, you can restart the clone teardown by repeating the procedure above. For example, if the ESXi host that contains the cloned VMs is not available on the network when the clone teardown first runs, you can run the clone teardown again when the ESXi host becomes available and the VMs will be deleted.

The Test & Dev workflow for VMware VMs provides two options when objects are cloned:

1. Use an existing view
2. Create a new view

If an existing view is used, it may contain other data. If a new view is created, it may be used for other purposes, or for future Test & Dev clones. So in either case, the views are not deleted on clone teardown. Further, it does not hinder any other process or action. The user can still clean up the view manually.

VM Not Deleted if Cloned VM is Altered

If a cloned VMware VM has been altered after being created by a clone task, the VM may not be automatically cleaned up during a clone teardown. These special cases are described below:

1. The VM is using files found in an additional external datastore besides the original datastore (view) that is storing the cloned VM. Here are some examples:

- A directory has been NFS mounted onto the VM and the VM is using files in that directory.
- The original VM files were moved to an additional external datastore using Storage vMotion.

2. The VM was converted to a VMware template.
3. The Fault Tolerant (FT) setting was enabled for the VM.

To delete the altered VMs, you can either manually delete the VMs or revert the changes made to the VMs and run the clone teardown again.

Clone Databases

This chapter includes the following topics:

Clone Microsoft SQL Server Databases

Cohesity now provides the ability to clone an MS SQL database to a Standalone SQL Server, FCI, and AG. A major advantage is that the cloned MS SQL database files reside on the Storage Domain and require no additional storage from the MS SQL host. Multiple clones from a single source of MS SQL Database backup require little or no additional storage space in the Cohesity cluster.

> **Note:** Cloning an MS SQL database is supported with Volume-based and file-based backup methods.

Clone an MS SQL Database

1. Navigate to **Test & Dev** > **Clone** > **Databases**.
2. Search and select the MS SQL database. You can optionally perform a search by specifying the wildcard character *. You can also optionally narrow the search results by specifying filter criteria.
3. Provide the following information listed in the table, and then click **Clone**.

Options	Description
Task Name	(*Optional*) Change the default name for the Clone task.
Objects Being Cloned	By default, the database you selected for cloning is displayed. You can click the change icon to select a different database.

Options	Description
Clone Point	By default, the latest snapshot is cloned. Optionally click the Change icon to select a different snapshot to clone. The selected Restore Point is listed
SQL Host	Select the host to which you want to clone the database. By default, you can select only the physical servers as the host. If you want to select MS SQL FCI as the SQL host, Contact Cohesity Support. Cloning the MS SQL database to an MS SQL FCI is supported only if the following criteria are met: • Microsoft SQL Server version must be 2012 and above, supported on SMB file server. • The MS SQL database to be cloned is backed up with a file-based protection group.
SQL Instance	Select the instance for the cloned database.
Rename Database?	Enable this option if you want to rename the cloned database created by this task.
Cohesity Network interface	Select one of the following options: • **Auto Select**: By default, the **Auto Select** option is enabled and the clone task automatically selects the correct VLAN. • **Interface Group**: select a configured interface group from the **Interface Group** drop-down.

Cancel a Database Clone

You can cancel a running clone task for an individual database. Canceling the task may leave the database in an unusable state. You can delete the database manually and then start another clone task as needed.

1. Navigate to **Test & Dev**.
2. Click a Task Name link in the table to display the details page.
3. Click the **Cancel** button. (The Cancel button is available only if the task is running.)
4. Respond to the confirmation prompt.

Tear Down a Database Clone

To delete an MS SQL database clone from the Cohesity Storage Domain, perform the following steps:

1. Navigate to **Test & Dev**.
2. Click the cloned task or hover your mouse over the actions menu located next to the cloned task you want to delete.
3. Click **Tear Down Clone** > **Yes, Tear Down**.

The MS SQL database clone will be deleted from the Cohesity cluster Storage Domain.

> **Note:** The teardown of an MS SQL clone database removes it from the SQL server; the user does not have to do this beforehand. However, in the rare instance you created other additional databases with files on the mounted drive that the cloned database is on, those additional databases should be deleted first through the SQL Server Management Studio.

Clone Oracle Database

Cohesity provides the ability to clone Oracle databases (Standalone, RAC, Multi-container Database (CDB), and VCS). The Oracle database is cloned at Cohesity Storage Domains and requires no additional storage from the Oracle Host. Multiple clones from a single source of Oracle Database backup may require little to no additional storage space

in the Cohesity cluster.

Supported Oracle Release for TDM Cloning

Cohesity supports TDM cloning in the following Oracle releases:

Oracle Release	Clone Support for Oracle Database
11gR2	Supported
12cR1	Supported
12cR2	Supported
18c	Supported
19c	Supported

Prerequisites

- Install the latest Cohesity Agent on the system running the Oracle Database.
- Capture Snapshots of the Oracle database to be cloned. To create Snapshots, see create and run a Protection Group.

Clone an Oracle Database

1. Navigate to **Test & Dev** > **Clone** > **Databases**.
2. Search and select the Oracle database. (You can optionally specify the wildcard character *in the search.) You can also optionally narrow the search results by specifying filter criteria.
3. Provide the following information listed in the table, and then click **Clone**.

CLI Option	Description
Task Name	(Optional) Change the default name for the Clone task.
Objects Being Cloned	The Oracle database is being cloned. You can click the change icon to select a different database.
Clone Point	The latest clone point. You click the change icon and select a different clone point. Select the Snapshot to clone. By default, the latest Snapshot is cloned and the selected Restore Point is listed. You can optionally click **Edit** and choose the previous Snapshot to clone.
Oracle Host	The host to which you clone the database.
Oracle Home	The home directory for the cloned database.
Base Directory	The base directory for the cloned database.
Database Name	The name for the cloned database. The maximum number of characters for the database name is eight. The database name should be unique for the given Oracle Host.

CLI Option	Description

Pfile Parameters	The Pfile parameters option allows you to see the target database parameters before the clone job is initiated. The Pfile inherits some source database parameters to the target database and allows you to edit the target database parameters based on your requirement by specifying the additional parameters as a `key=value` pair.
	
	The Pfile lists the following parameters:
	
	• **Default Parameters**: Default parameters of the target database and you can edit the default
	• **Inherited Parameters**: Parameters from the source database and you can edit the inherited parameters, if required.
	• **Restricted Parameters**: You cannot edit the restricted parameters in the Pfile. The following are the restricted parameters:
	  • INSTANCE_NAME
	  • DB_NAME
	  • DB_UNIQUE_NAME
	  • DB_BLOCK_SIZE
	  • MAX_STRING_SIZE
	  • COMPATIBLE
	
	If you edit the following values in the Clone Oracle GUI, You must click the Reset button on the top-right corner to reset the Pfile and get the updated parameter values in the Pfile box.
	  • Oracle Home
	  • Oracle Base
	  • Database Name
	
	Target database will be created with the specified parameters in the Pfile. For example, if `SGA_TARGET` is specified as a part of the pfile parameters, the clone operation will use the value specified in the pfile during the cloning of the source database.

SGA (System Global Area)	The SGA is an area of memory (RAM) allocated to an Oracle Instance when it starts up. Ideally, SGA's value varies from 1024M to 102400M.

CLI Option	Description
Shell Environment	User can configure their shell environment that executes the Cohesity clone workflow using this setting. An example would be, define your TNS_ADMIN shell variable here to point to a different sqlnet.ora file to be used for your clone target database. For TDE support, the wallet location for your clone target database may depend on a shell variable. Use this to specify your wallet location for cloning a backup taken from a TDE database. For more information, see Clone Oracle TDE Database.
NFS Protocol	Select the NFS protocol version for the clone job. Options include NFSv3 and NFSv4.1. **Note:** NFS protocol is not supported for Windows and AIX databases.
Change DBID	If this option is disabled, the cloned database UUID will be the same as the source database and it will not be displayed on the Source page. Later, If the cloned databases are required to be discovered by Cohesity, then follow the steps below: 1. Start the database in MOUNT mode. 2. Run the following query to get the DBID: “select DBID from GV$database”. 3. Change the database DBID using the NID command. 4. Run the following query “select DBID from GV$database” and confirm that the DBID is different from the one we obtained in step 2. 5. Once confirmed, start the database in OPEN mode. 6. Add the database entry in the oratab file. 7. Do a source refresh. 8. The database should be listed there now. If none of the other databases on the Oracle host have DBID the same as that of the cloned database then, skip steps 1 to 5.

You will be redirected to the status details page. Once the clone task that you initiated completes, the status will be displayed as “Success”.

You can cancel a clone task for an individual database. Canceling the clone task leaves the database in an unusable state. You can delete this database manually at the Oracle host and then restart the clone task as needed.

Cancel an Oracle Database Clone Task

You can cancel a clone task for an individual database. Canceling the clone task leaves the database in an unusable state. You can delete this database manually at the Oracle host and

then restart the clone task as needed.

1. Navigate to **Test & Dev**.
2. Click a Task Name link in the table to display the details page.
3. Click the **Cancel** button. (The Cancel button is available only if the task is running.)
4. Respond to the confirmation prompt.

Test & Dev Using CloudSpin

You can use Cohesity's CloudSpin for Test & Dev. For more information, see Clone VMs to AWS Cloud Using CloudSpin.

Clone vCloud Director (vCD) VMs

Cloning vCloud Director (vCD) VMs using Cohesity involves integrating Cohesity's data protection platform with your vCD environment. Cohesity allows users to clone vCD VMs and vApps from backups for Test & Dev use cases.

Considerations

- Clone vCD VMs should only use the "Detach network" option for all users, tenants, and administrators.

  > **Note:** Administrators can use attach options until they select a network that is part of the vApp being cloned.

- Cloning templates as templates is not supported for all tenants and administrators.
- The option "Detach Network" will prevent the original VM NIC cards from being reattached after cloning.
- If you cannot attach NICs/Org vDC networks during the cloning workflow, please contact Cohesity support for assistance.
- While cloning VMs to the target, the original IP address and IP mode are not recovered to the NIC cards. By default, the IP mode is set to "none" and the IP address remains empty. Contact Cohesity support to recover the original IP addresses for the clone operation.

Chapter 32: Monitoring

This chapter includes the following topics:

- Performance 4592
- Monitor Performance of Workloads 4595
- Configure and Test SNMP 4602
- Storage Statistics 4608
- Health Dashboard 4621
- Reports 4820
- Audit Logs 4868
- Advanced Diagnostics 4893

Performance

The performance statistics page helps you to view the performance of the Cohesity cluster, Storage Domains, and Nodes with respect to the attributes such as CPU, memory, IOPS, throughput, and latency. By default, the page provides performance statistics for the entire cluster.

You can view performance statistics by selecting **System** > **Performance**.

Cluster Performance Statistics

To view Cluster performance statistics, select **Cluster** from the top navigation pane of the performance page.

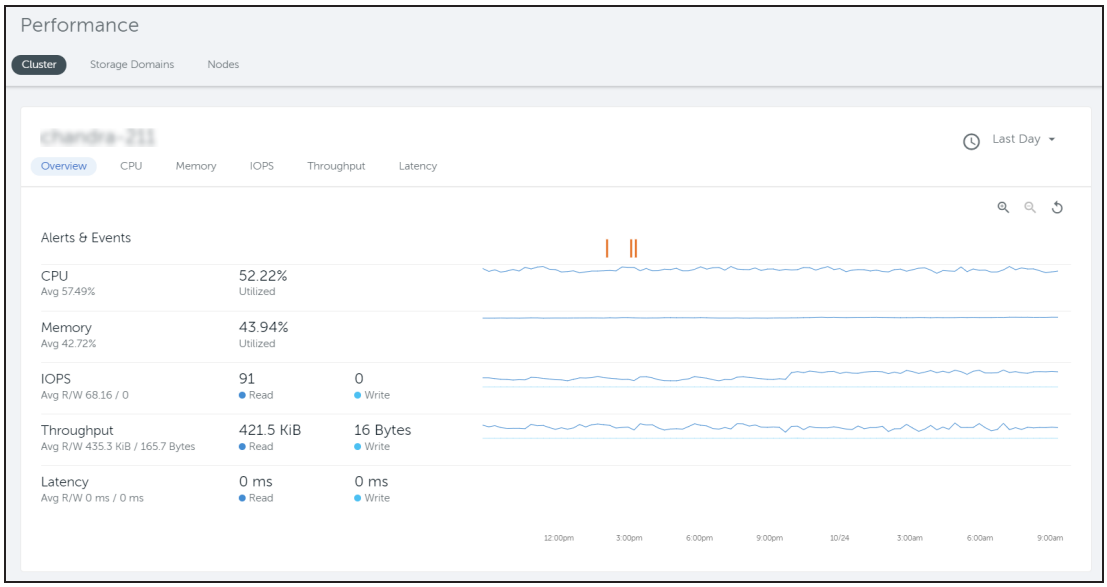

The **overview** tab of the Cluster performance screen shows the performance graph of the following attributes:

- CPU—The total CPU utilization of the Cluster.
- Memory— The total memory utilization of the Cluster.
- IOPS—The read and write I/O operations of the Cluster.
- Throughput—The throughput of the Cluster.
- Latency—The average response time for read and write operations for Views configured with **TestAndDev High** QoS policy.

Following are the characteristics of the performance statistics screen of a cluster:

- You can customize the time frame of each statistic by selecting **Last Day**, **Last Week**, or **Last Month**.
- You can click each attribute to view the activities or jobs that are running in the cluster during the specified time frame. For example, you can click the **CPU** tab to view the jobs that are running on the cluster during the specified time frame.
- You can zoom in on part of the graph by dragging the mouse across a section of the graph or by clicking the **Zoom** icon.
- From each attribute performance page, you can use the **Adapters** and **Tasks** drop-down lists or use the search option and filter the jobs or activities that are displayed.

Storage Domain Performance Statistics

To view performance statistics at the Storage Domain level, select **Storage Domain** from the top navigation pane of the performance page.

By default, the performance statistics of the default Storage Domain are displayed. You can select a Storage Domain from the drop-down menu to view the performance statistics of the selected Storage Domain.

The **overview** tab of the Storage Domain performance screen shows the performance graph of the following attributes:

- IOPS—The IO read and write operations of the selected Storage Domain.
- Throughput—The throughput of the selected Storage Domain.
- Latency—The average response time for read and write operations for Views configured with **TestAndDev High** QoS policy.

Following are characteristics of the performance statistics screen of a Storage Domain :

- You can customize the time frame of each statistic by selecting **Last Day**, **Last Week**, or **Last Month**.

- You can click each attribute to view the activities or jobs that are running in the cluster for that storage domain during the specified time frame. For example, you can click the **IOPS** tab to view the jobs that are running on the cluster during the specified time frame.
- You can zoom in on part of the graph by dragging the mouse across a section of the graph or by clicking the **Zoom** icon.
- From each attribute performance page, you can use the **Adapters** and **Tasks** drop-down lists or use the search option and filter the jobs or activities that are displayed.

Nodes performance Statistics

To view performance statistics per node, select **Nodes** from the top navigation pane of the performance page.

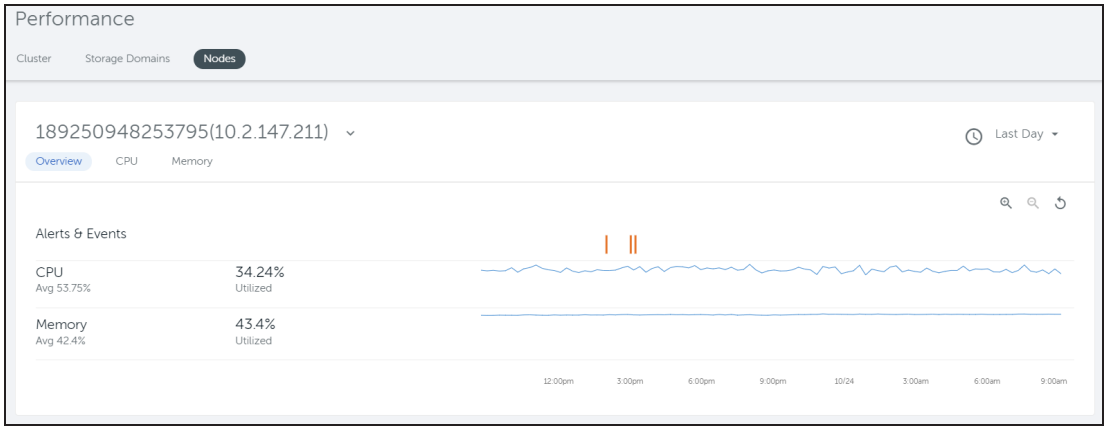

You can select a node from the drop-down menu to view the performance statistics of the selected node.

The overview tab of the Nodes performance screen shows the performance graph of the following attributes:

- CPU—The performance graph of the total CPU utilization of the services that are running on the selected Node.
- Memory—The performance graph of the total memory utilization of the services that are running on the selected Node.

Following are the characteristics of the performance statistics screen of a Node:

- You can customize the time frame of each statistic by selecting **Last Day**, **Last Week**, or **Last Month**.
- You can click each attribute to view the performance graph of the following Cohesity services that are running on the cluster during the specified time frame.
    - Background System Activity—Collects garbage and fixes replication and deduplication.

- File System—Stores and retrieves data.
- Data Protection—Backs up data.
- Metadata—Stores and maintain metadata.
- Indexing—Indexes the backed-up data.

For example, you can click the **CPU** tab to view the performance graph of the services that are running on the node during the specified time frame. The tab displays the percentage of the total CPU utilized by all the services and the percentage of the CPU utilized by each service.

- You can zoom in on part of the graph by dragging the mouse across a section of the graph or by clicking the Zoom icon.

Monitor Performance of Workloads

Cohesity now provides component-level logical statistics and granular-level statistics that enable you to inspect resource usage and workload distribution across the I/O Operations service. Statistics available on the **System > Performance > Workloads** tab help you:

- Check the overall performance of the Cohesity cluster
- Inspect the resource usage across various components
- Detect whether the cluster is overloaded
- Debug workload-level performance issues
- Debug replication performance issues
- Debug intermittent performance issues

Historical data is available for up to 30 days. This enables you to go back in time and debug issues. For example, if the SmartFiles workload was experiencing high latency at a certain time, you can go back to that particular point in time and view the granular statistics to identify the root cause. You can select one of the following time ranges:

- Past hour
- Past 12 hours
- Past 24 hours
- Past 7 days
- Past 30 days

Add Data to the Charts

You can add or update the data displayed on the charts.

To add or update data displayed on the charts:

1. Navigate to **System > Performance** and click the **Workloads** tab.
2. Select a performance metric:
   - Throughput
   - IOPS
   - Latency
3. Click **Add Data**.
4. Add or update data.

   > **Important:** Depending on the chart, data can be one of the following—workloads, adapters, protection groups, or sources and you can add only up to five workloads, adapters, protection groups, or sources.

5. Click **Apply**.

   The following image illustrates how you can add or update workloads. If five workloads are already added, clear the check box corresponding to the workload that you want to remove and select a new workload:

View Tooltips

You can view tooltips by hovering your mouse over the chart. Hover over a specific date and time to view tooltip data that is automatically generated based on the underlying read and write data. For example, the following image illustrates the tooltip data for the selected workloads on December 6, 2022 at 3:27 PM:

Zoom In and Zoom Out of Charts

You can zoom in on the charts to increase the visible time range. To zoom in, click and drag the mouse pointer over the chart. To zoom out, click **Reset zoom**.

The following image illustrates how you can zoom in and zoom out of the charts:

Components and Supported Filtering Options

The following table lists the components and the available filtering options for each component. You can select a filter to drill down and view granular details. As an example, for the **Replication** component, you can filter based on adapters or protection groups and get granular-level statistics.

Component	Filtering Options (Granular-level statistics)
Replication	• Adapters • Protection Groups
Indexing	• Adapters • Protection Groups
Apps	Not applicable
SmartFiles	• Protocol
Archive	• Adapters • Protection Groups
Background	Not applicable
Backup	• Adapters • Protection Groups • Sources
Restore	• Adapters • Recoveries
Other	Not applicable

Monitor Throughput of Workloads

To monitor the throughput of the workloads:

1. Navigate to **System > Performance** and click the **Workloads** tab.
2. Select **Throughput** from the drop-down menu:

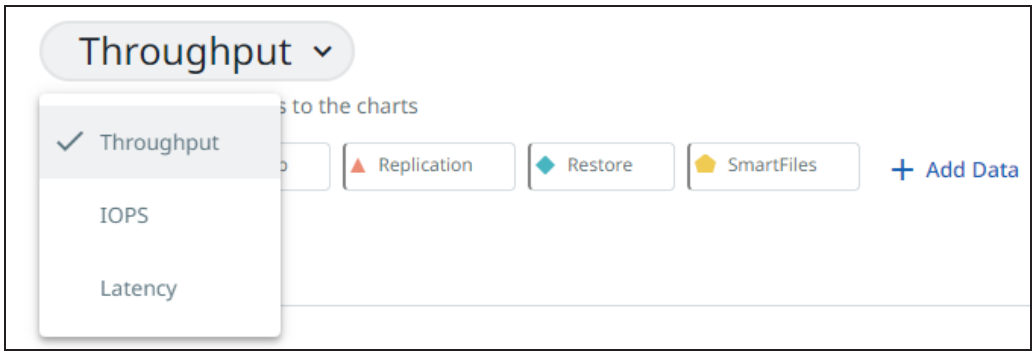

3. Select the workloads.

   The read and write throughput of the workloads are displayed. The following image illustrates the throughput of the selected workloads for the past 7 days:

4. After the charts load, you can click on a workload to drill down further. For example, click **Backup** to view the backup details:

5. You can further slice and dice the metrics based on additional filters. For more information about the available filters, see Components and Supported Filtering Options.

Monitor IOPS of Workloads

To monitor the IOPS of the workloads:

1. Navigate to **System > Performance** and click the **Workloads** tab.
2. Select **IOPS** from the drop-down menu:

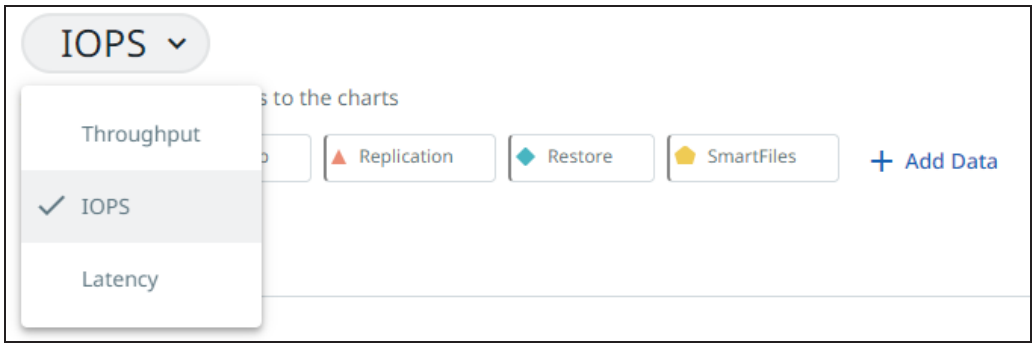

3. Select the workloads.

   The read and write IOPS of the workloads are displayed. The following image illustrates the IOPS of the selected workloads for the past 12 hours:

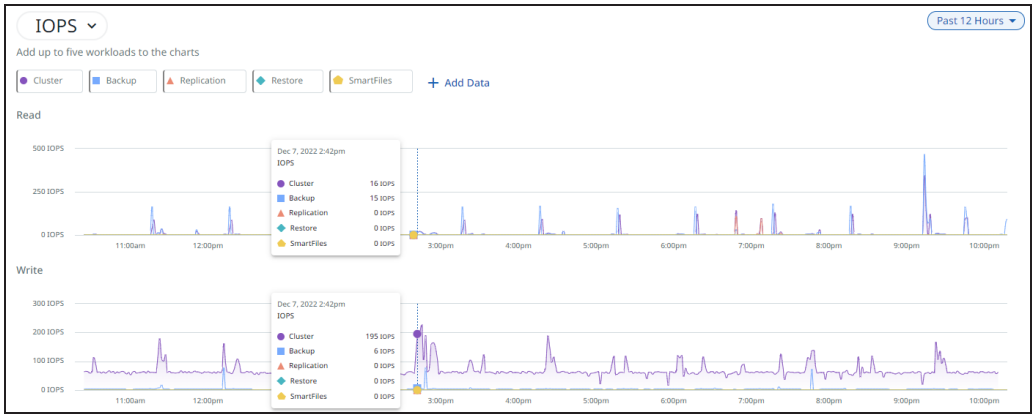

4. After the charts load, you can click on a workload to drill down further. For example, click **Backup** to view the backup details:

5. You can further slice and dice the metrics based on additional filters. For more information about the available filters, see Components and Supported Filtering Options.

Monitor Latency of Workloads

To monitor the latency of the workloads:

1. Navigate to **System > Performance** and click the **Workloads** tab.
2. Select **Latency** from the drop-down menu:

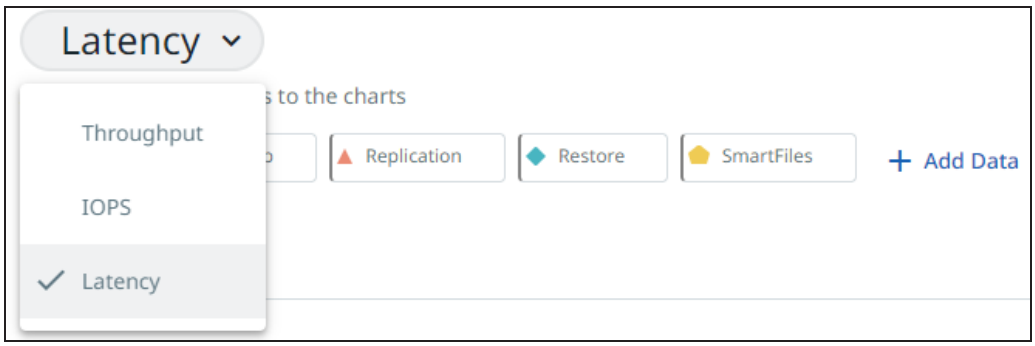

3. Select the workloads.
4. The read and write latency of the workloads are displayed. The following image illustrates the latency of the selected workloads for the past 24 hours:

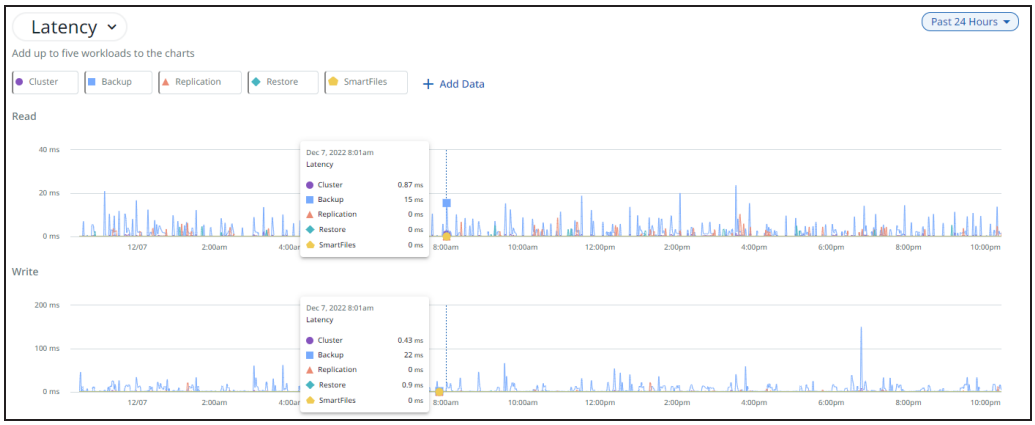

5. After the charts load, you can click on a workload to drill down further. For example, click **Backup** to view the backup details:

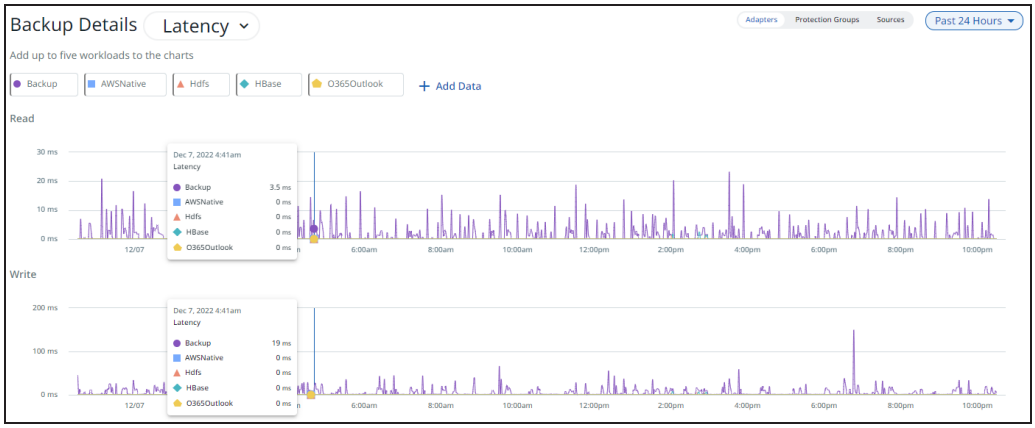

6. You can further slice and dice the metrics based on additional filters. For more information about the available filters, see Components and Supported Filtering Options.

Configure and Test SNMP

You can enable Simple Network Management Protocol (SNMP) on the Cohesity cluster to further integrate Cohesity into your data center and enterprise management solutions. Whenever events occur on the Cohesity cluster, SNMP notifications are sent to your network management system.

Prerequisites and Considerations

- Cohesity supports SNMP versions 2 and 3.
- The Cohesity cluster supports SNMP polling and notifications.
- Ports 161 (UDP, incoming) and 162 (UDP, outgoing) must be open.

- The SNMP network management system must be accessible from the Cohesity cluster.
- Disabling SNMP stops the notification broadcasts but does not clear the existing SNMP configuration information.
- For a list of SNMP Object Identifiers (OIDs), see Alerts Reference.

Configure SNMP

1. In the Cohesity Dashboard, select **Settings** > **Summary** and click **SNMP** > **Configure SNMP**.

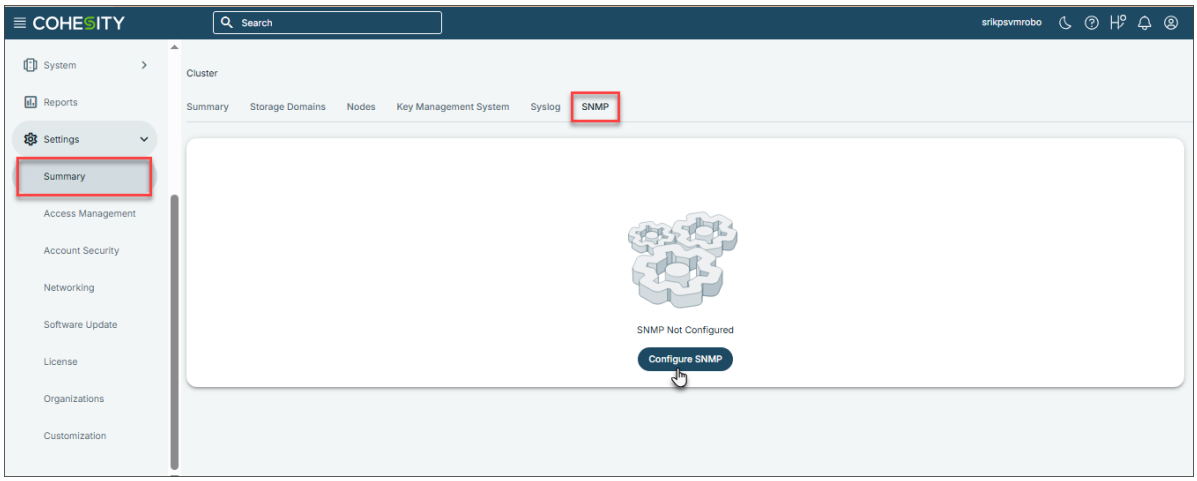

The **SNMP Configuration** dialog appears.

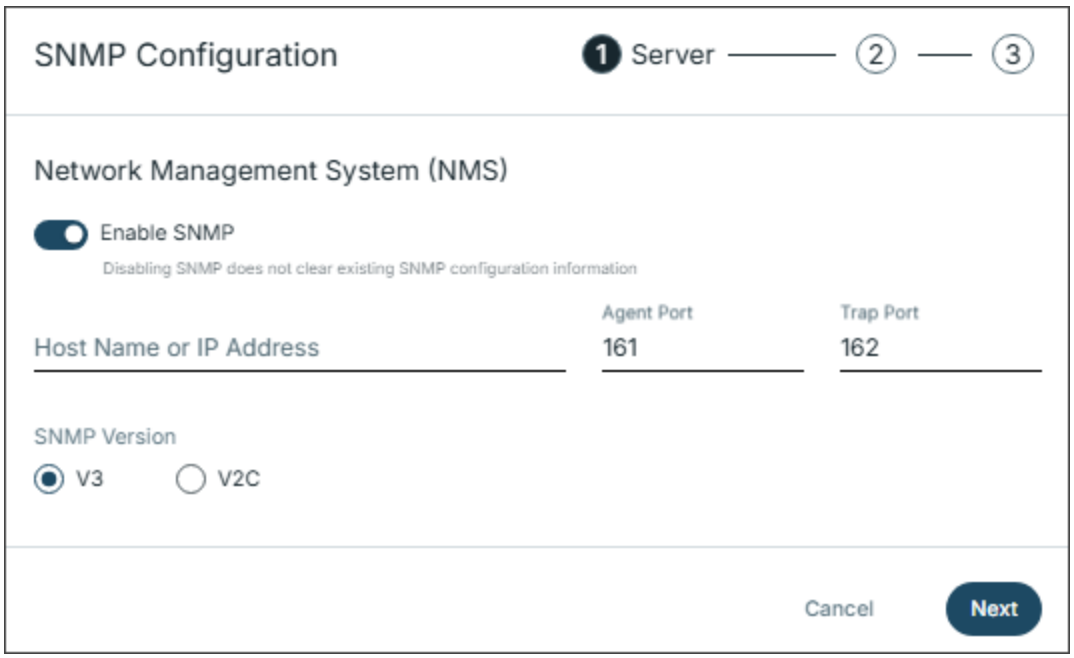

2. Provide the necessary information in the following fields and click **Next**.

Field	Description
Enable SNMP	Use the toggle to enable or disable SNMP. It is enabled by default.
Host Name or IP Address	Enter the IP address or host name of the SNMP trap destination to which you want to send SNMP notifications.
Agent Port	This is the port the SNMP agent listens on. The default is 161.
Trap Port	This is the SNMP trap destination port. The default is 162.
SNMP Version	• V3 - More secure (supports encryption and authentication) • V2C - Basic security with community strings.

3. If you select V3, provide the necessary information in the following fields and click **Next**:

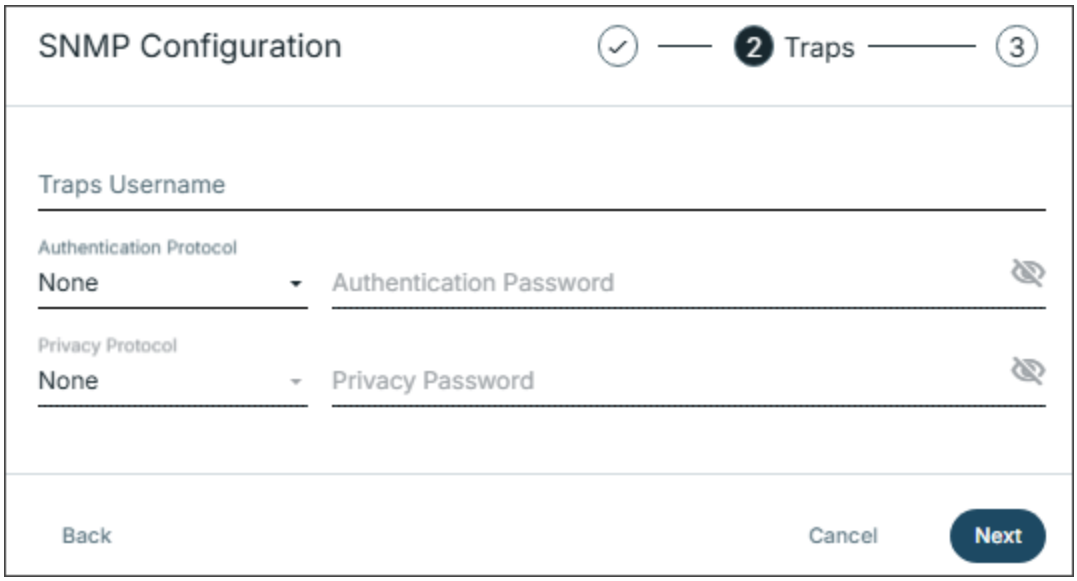

Field	Description
Traps Username	The default SNMP V3 username for the Cohesity cluster is **cohesityV3Trap**. You can change this string.

Field	Description
Authentication Protocol	• **None** - No security • **SHA-224** - Uses HMAC-SHA-224 for authentication • **SHA-256** - Uses HMAC-SHA-256 for authentication • **SHA-384** - Uses HMAC-SHA-384 for authentication • **SHA-512** - Uses HMAC-SHA-512 for authentication
Authentication Password	Required only if you select a SHA protocol.
Privacy Protocol	• **None** - No encryption • **AES-192** - AES-192 (CFB mode) encryption • **AES-256** - AES-256 (CFB mode) encryption
Privacy Password	Required only if you select an AES protocol.

The **Configure User Access** dialog appears.

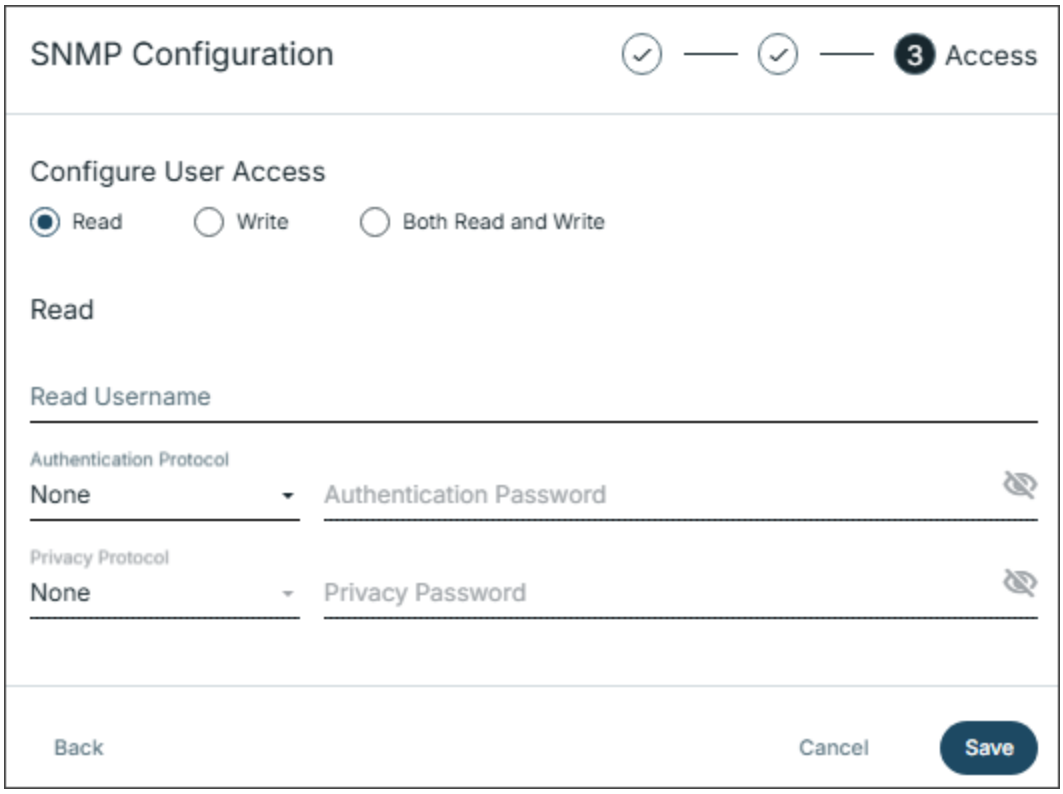

Field	Description
Configure User Access	• **Read** • **Write** • **Both Read and Write**
Read Username	The default read username is **cohesityV3Public**.
Write Username	The default write username is **cohesityV3Private**.
Authentication Protocol	• **None** - No security • **SHA-224** - Uses HMAC-SHA-224 for authentication • **SHA-256** - Uses HMAC-SHA-256 for authentication • **SHA-384** - Uses HMAC-SHA-384 for authentication • **SHA-512** - Uses HMAC-SHA-512 for authentication
Authentication Password	Required only if you select a SHA protocol.
Privacy Protocol	• **None** - No encryption • **AES-192** - AES-192 (CFB mode) encryption • **AES-256** - AES-256 (CFB mode) encryption
Privacy Password	Required only if you select an AES protocol.

4. If you select V2C, provide the necessary information in the following field and click **Next**:

   - **Traps Community String**: The default traps community string for the Cohesity cluster is **cohesityV2Trap**. You can change this string.

The **Configure Community String Access** dialog appears.

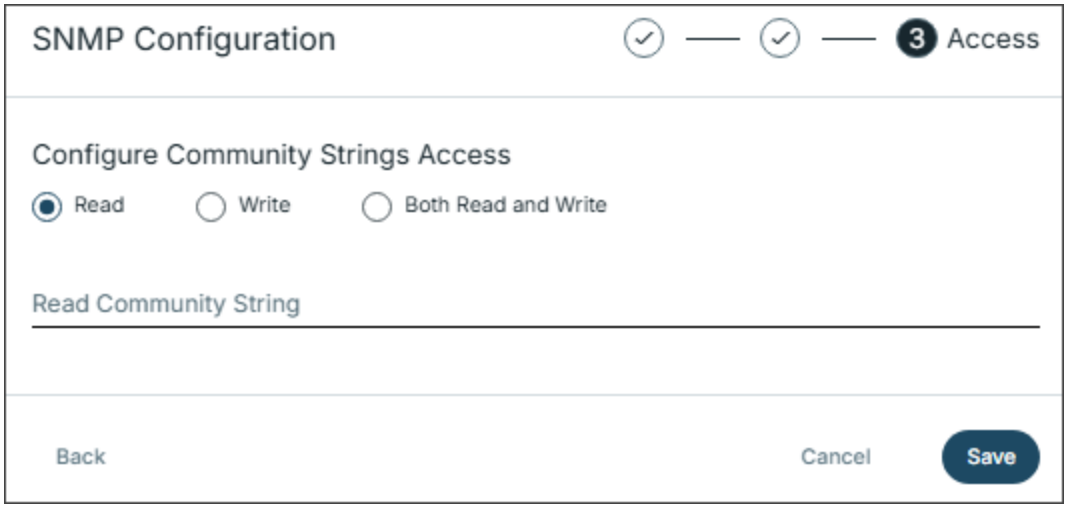

Field	Description
Configure Community Strings Access	• **Read** • **Write** • **Both Read and Write**
Read Community String	The default read community string is **cohesityV2Public**.
Write Community String	The default write community string is **cohesityV2Private**.

5. Click **Save** to save the SNMP configuration.

The page shows the updated SNMP configuration.

To edit the SNMP configuration, go to the SNMP Configuration page, click **Edit Configuration**, and update the necessary fields.

Verify SNMP Transmission to Network Management Systems

1. Locate the Network Management System hostname or IP address by accessing the Cohesity Dashboard and selecting **Settings > SNMP**.
2. Run the following command to find the alerts master node:

```
iris_cli cluster ls-masters
```

3. Run the following `tcpdump` command from alerts master node:

```
sudo tcpdump -ni any port 161 or 162 -vv
```

> **Note:** Port 161 is used for SNMP packets and port 162 is used for SNMP traps.

4. Use the following object ID to generate a non-impacting event:

   LogsDiskSpaceLow 1.3.6.1.4.1.47421.0.6

   This event is triggered when the logs partition reaches over 95% full. Use the following steps to determine available space in the logs partition.

```
df -h /home/cohesity/logs
logs=/home/cohesity/logs ; df -h $logs ; fallocate -l 38g $logs/foo
df -h $logs ; sleep 300 ; rm $logs/foo ; df -h $logs
```

5. Observe the SNMP trap confirmation displayed on the wire with `tcpdump` running on alerts' master node.

```
16:06:30.212028 IP (tos 0x0, ttl 64, id 41099, offset 0, flags [DF],
proto UDP (17), length 210) <source_host_hostname> > <NMS_server_host_
hostname>.snmptrap: { SNMPv2c C=YLX0 { V2Trap(167) R=1766176774
system.sysUpTime.0=198748646 S:1.1.4.1.0=E:47421.0.6
E:47421.1.6="<source_host_hostname> : INFO | CE00401007 -
LogsDiskSpaceLow | Logs partition disk space is low" } } E.....@.@...
..* ....r......0.......YLX0.....iE........0..0...+.......C.....0..
+.......... +......=..0j. +......=...\<source_host_hostname> : INFO |
CE00401007 - LogsDiskSpaceLow | Logs partition disk space is low 6.
ctrl-c out tcpdump
```

> **Note:** The example contents of tcpdump given above are not encrypted. The actual contents may be encrypted based on the SNMP configuration.

Test SNMP Reporting for Cohesity Alerts

Run the following command to test SNMP reporting:

```
iris_cli snmp send-test-trap
```

Storage Statistics

You can view storage statistics by selecting **System > Storage**.

Cluster Level Data Statistics

In the **Cluster** tab:

- You can view the storage statistics for the entire cluster. The details are provided under two panes, namely **Consumption** and **Trend**. Click ⛶ on each pane for a more detailed view.
- You can customize the timeframe by selecting **Last Day**, **Last Week** or **Last Month** in the **Trend** pane.
- You can also customize the data presented in the chart. To remove or add chart data, click ⛶ in the **Trend** pane and toggle any combination of the following fields: **Storage Saved** or **Resiliency Impact**. The selected data type is added to or removed from the chart.

This page reports on the following Cohesity cluster data statistics:

Metrics	Description
Logical Data	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type. • **VMs**: The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For examples, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. • **All Other Workloads**: The data size reported is the actual front end data residing on the server. If a server with 1TB capacity contains 100GB of data, the server reports 100GB. **Note:** • These metrics do not include unprotected objects. • Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000GB (1TB). **SmartFiles**: Size of total unique logical data of active filesystem as residing on Cohesity. This does not count the space consumed by the Data Protection runs (i.e. snapshots) that may be protecting the active filesystem.

Metrics	Description
Data Read	Data Read is the size of the set of protected objects as read by Cohesity for a single backup run. This number is a per protection run statistic and is not additive across backup runs. > **Note:** The Data Read column displays 0 bytes for view-based backups as the protected data resides within a Cohesity view on the cluster and is not read from an external source.
Data In	Total unique logical data residing on the Cohesity cluster. For Data Protection, this is the sum of your first full backup and subsequent change data captured by Cohesity. For SmartFiles, some customers, such as SPs and enterprise customers with multiple tenants, may want to use this as the Logical Data so that snapshot capacity is included in the chargeback.
Data Written	Data written on the Cohesity cluster after the unique logical data has been reduced by data deduplication and data compression. For SmartFiles, this metric value carries the name of Physical Data. > **Note:** This number reflects unique data written, before resiliency operations.
Resiliency Impact	The space required to honor resiliency settings (Erasure Coding / Replication Factor) in addition to the data written. For example, if the Data Written is 100GB and resiliency setting is RF3, the **Resiliency Impact** is 200GB. That is, to honor the resiliency setting, you will need an additional 200GB on top of the 100GB for the **Data Written**.

Metrics	Description
Storage Consumed	The total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for. This data can be stored either all locally or as a combination of local storage and in the Cloud Tier (if Cloud Tier is enabled). For example, if the **Data Written** is 100GB and resiliency is set to RF3, leading to the need for an additional 200GB of space, the resulting **Storage Consumed** is 300GB. If the **Data Written** is 100GB and resiliency is set to EC4:2, leading to the need for an additional 50GB (as for every 4 data fragments, 2 parity fragments are added), the resulting **Storage Consumed** is 150GB. **Note:** If you have enabled Cloud Tiering, then **Storage Consumed** is the sum of space consumed on local storage and space consumed on the cloud tier. In such cases, the sum of **Storage Consumed** value and **Capacity Available** value for the cluster can exceed the cluster capacity.
Storage Reduction	The ratio of **Logical Data** to **Storage Consumed**. For Data Protection, this reflects all operations performed by Cohesity on source data, which include snapshotting, change tracking and ingest, data deduplication, data compression, and data resiliency. For SmartFiles, this reflects all operations performed on the Cohesity hosted active filesystem, which include logical data writes, change tracking, data deduplication, data compression, and data resiliency.
Storage Consumed Local and Cloud	If you have enabled Cloud Tiering, then this metric is the sum of space consumed on local storage and space consumed on the cloud tier.
Storage Saved	The space owing to data deduplication and data compression. For example, if **Data In** is 100GB, deduplication ratio is 4x, compression is 2x, then **Data Written** is 12.5GB. Therefore, the **Storage Saved** is 100GB - 12.5GB = 87.5GB.
Total Capacity	The total physical capacity of all storage devices after subtracting the space reserved for Cohesity software operations.
Available	The amount of space still available on the Cohesity cluster, calculated by subtracting the **Storage Consumed** (Locally and not on Cloud Tier) from the **Total Capacity**.

Storage Domain Level Data Statistics

In the **Storage Domains** tab:

- You can view the storage statistics of the storage domains. To view the storage statistics for a single storage domain, from the **Storage Domain** drop-down, select a storage domain.
- You can view the storage domain assigned to Organizations. From the **Organizations** drop-down, select the organization.
- You can view the data growth rate for the **Past 24 Hours**, **Past 7 Days**, **Past 30 Days** or **Past 90 Days**. From the **Growth** drop-down, select the required option.

Click a storage domain to view more details. The **Consumption** tab shows the current data consumed and the space saved owing to data deduplication and data compression. The **Trend** tab shows the data growth trend. The data is presented in a chart. You can customize the timeframe by selecting **Last Day**, **Last Week** or **Last Month**. Turn on or off the **Storage Saved** or **Resiliency Impact** toggle to add or remove data from the chart.

To sort the list in the table by a Column type, click the Column header. To reverse the sorting, click the Column header again.

The page reports on the following Storage Domain data statistics.

> **Note:** All the statistics that are displayed are specific to the selected Storage Domain. If there is more than one Storage Domain, the Storage Domain statistics will be different from the Cluster level statistics.

Metrics	Description
Quota	The total amount of data that can be written to the Storage Domain (value set for **Storage Domain Physical Quota**). After it has been reduced by deduplication and compression, including the space required for honoring the resiliency settings (EC/RF). **Note:** This field is displayed as **None** if **Storage Domain Physical Quota** is not set.

Metrics	Description
Storage Consumed	The total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for. This data can be stored either all locally or as a combination of local storage and in the Cloud Tier (if Cloud Tier is enabled). For example, if the **Data Written** is 100GB and resiliency is set to RF3, leading to the need for an additional 200GB of space, the resulting **Storage Consumed** is 300GB. If the **Data Written** is 100GB and resiliency is set to EC4:2, leading to the need for an additional 50GB (as for every 4 data fragments, 2 parity fragments are added), the resulting **Storage Consumed** is 150GB. **Note:** If you have enabled Cloud Tiering, then **Storage Consumed** is the sum of space consumed on local storage and space consumed on the cloud tier. In such cases, the sum of **Storage Consumed** value and **Capacity Available** value for the cluster can exceed the cluster capacity.
Storage Consumed (Local)	Data stored locally including the space required for honoring the resiliency settings (EC/RF).
Storage Consumed (Cloud Tier)	Data stored on the Cloud Tier including the space required for honoring the resiliency settings (EC/RF).
Storage Reduction	The ratio of **Logical Data** to **Storage Consumed**. For Data Protection, this reflects all operations performed by Cohesity on source data, which include snapshotting, change tracking and ingest, data deduplication, data compression, and data resiliency. For SmartFiles, this reflects all operations performed on the Cohesity hosted active filesystem, which include logical data writes, change tracking, data deduplication, data compression, and data resiliency.
Storage Saved	The space owing to data deduplication and data compression. For example, if **Data In** is 100GB, deduplication ratio is 4x, compression is 2x, then **Data Written** is 12.5GB. Therefore, the **Storage Saved** is 100GB - 12.5GB = 87.5GB.

Metrics	Description
Logical Data	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type. • **VMs**: The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For examples, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. • **All Other Workloads**: The data size reported is the actual front end data residing on the server. If a server with 1TB capacity contains 100GB of data, the server reports 100GB. **Note:** • These metrics do not include unprotected objects. • Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000GB (1TB). **SmartFiles**: Size of total unique logical data of active filesystem as residing on Cohesity. This does not count the space consumed by the Data Protection runs (i.e. snapshots) that may be protecting the active filesystem.
Data In	Total unique logical data residing on the Cohesity cluster. For Data Protection, this is the sum of your first full backup and subsequent change data captured by Cohesity. For SmartFiles, some customers, such as SPs and enterprise customers with multiple tenants, may want to use this as the Logical Data so that snapshot capacity is included in the chargeback.
Data Written	Data written on the Cohesity cluster after the unique logical data has been reduced by data deduplication and data compression. For SmartFiles, this metric value carries the name of Physical Data. **Note:** This number reflects unique data written, before resiliency operations.

Metrics	Description
Resiliency Impact	The space required to honor resiliency settings (Erasure Coding / Replication Factor) in addition to the data written. For example, if the Data Written is 100GB and resiliency setting is RF3, the **Resiliency Impact** is 200GB. That is, to honor the resiliency setting, you will need an additional 200GB on top of the 100GB for the **Data Written**.
Storage Consumed Local and Cloud	If you have enabled Cloud Tiering, then this metric is the sum of space consumed on local storage and space consumed on the cloud tier.
Available	The amount of space still available on the Cohesity cluster, calculated by subtracting the **Storage Consumed** (Locally and not on Cloud Tier) from the **Total Capacity**.
Target	Name of the **External Target** used for Cloud Tier.

Organizations Level Statistics

In the **Organizations** tab:

- You can view the storage statistics of organizations.
- You can view the storage domain assigned to Organizations. From the **Storage Domain** drop-down, select the storage domain.
- You can view the data growth rate for **Past 24 Hours**, **Past 7 Days**, **Past 30 Days** or **Past 90 Days**. From the **Growth** drop-down, select the required option.

Click an organization to view more details. The **Consumption** tab shows the current data consumed and the space saved owing to data deduplication and data compression. The **Trend** tab shows the data growth trend. The data is presented in a chart. You can customize the timeframe by selecting **Last Day**, **Last Week** or **Last Month**. Turn on or off the **Storage Saved** or **Resiliency Impact** toggle to add or remove data from the chart.

To sort the list in the table by a Column type, click the Column header. To reverse the sorting, click the Column header again.

The page reports on the following organization data statistics.

Metrics	Description
Organization	The name specified for the organization when added to the cluster.

Metrics	Description
Storage Consumed	The total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for. This data can be stored either all locally or as a combination of local storage and in the Cloud Tier (if Cloud Tier is enabled). For example, if the **Data Written** is 100GB and resiliency is set to RF3, leading to the need for an additional 200GB of space, the resulting **Storage Consumed** is 300GB. If the **Data Written** is 100GB and resiliency is set to EC4:2, leading to the need for an additional 50GB (as for every 4 data fragments, 2 parity fragments are added), the resulting **Storage Consumed** is 150GB.
Unique Data	Data that is unique to the objects across the system.
Data Written	Data written on the Cohesity cluster after the unique logical data has been reduced by data deduplication and data compression. For SmartFiles, this metric value carries the name of Physical Data. > **Note:** This number reflects unique data written, before resiliency operations.
Resiliency Impact	The space required to honor resiliency settings (Erasure Coding / Replication Factor) in addition to the data written. For example, if the Data Written is 100GB and resiliency setting is RF3, the **Resiliency Impact** is 200GB. That is, to honor the resiliency setting, you will need an additional 200GB on top of the 100GB for the **Data Written**.
Storage Reduction	The ratio of **Logical Data** to **Storage Consumed**. For Data Protection, this reflects all operations performed by Cohesity on source data, which include snapshotting, change tracking and ingest, data deduplication, data compression, and data resiliency. For SmartFiles, this reflects all operations performed on the Cohesity hosted active filesystem, which include logical data writes, change tracking, data deduplication, data compression, and data resiliency.
Storage Saved	The space owing to data deduplication and data compression. For example, if **Data In** is 100GB, deduplication ratio is 4x, compression is 2x, then **Data Written** is 12.5GB. Therefore, the **Storage Saved** is 100GB - 12.5GB = 87.5GB.

Metrics	Description
Logical Data	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type. • **VMs**: The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. • **All Other Workloads**: The data size reported is the actual front end data residing on the server. If a server with 1TB capacity contains 100GB of data, the server reports 100GB. **Note:** • These metrics do not include unprotected objects. • Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000GB (1TB). **SmartFiles**: Size of total unique logical data of active filesystem as residing on Cohesity. This does not count the space consumed by the Data Protection runs (i.e. snapshots) that may be protecting the active filesystem.
Data In	Total unique logical data residing on the Cohesity cluster. For Data Protection, this is the sum of your first full backup and subsequent change data captured by Cohesity. For SmartFiles, some customers, such as SPs and enterprise customers with multiple tenants, may want to use this as the Logical Data so that snapshot capacity is included in the chargeback.
Data Reduction	The ratio of **Data In** to **Data Written**. This ratio reflects the data deduplication and data compression operations.
Storage Consumed Local and Cloud	If you have enabled Cloud Tiering, then this metric is the sum of space consumed on local storage and space consumed on the cloud tier.
Growth	The physical data written to the storage domain over the number of days specified in the **Growth** field.

Metrics	Description
Growth Rate	The percentage of physical data growth on the storage domain over the number of days specified in the **Growth** field.

Consumers Level Statistics

In the **Consumers** tab:

- You can view the storage statistics of **Views**, **Backups**, **View Backups** and **Replication In**. From the drop-down, select the required option.
- You can view the storage domain assigned to Organizations. From the **Storage Domain** drop-down, select the storage domain.
- You can view the storage domain assigned to Organizations. From the **Organizations** drop-down, select the organization.
- You can view the data growth rate for **Past 24 Hours**, **Past 7 Days**, **Past 30 Days** or **Past 90 Days**. From the **Growth** drop-down, select the required option.

Click a consumer name to view more details. The **Consumption** tab shows the current data consumed and the space saved owing to data deduplication and data compression. The **Trend** tab shows the data growth trend. The data is presented in a chart. You can customize the timeframe by selecting **Last Day, Last Week** or **Last Month**. Turn on or off the **Storage Saved** or **Resiliency Impact** toggle to add or remove data from the chart.

To sort the list in the table by a Column type, click the Column header. To reverse the sorting, click the Column header again.

The page reports on the following consumer data statistics.

Metrics	Description
Consumer	The name of the consumer.
Storage Consumed	The total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for. This data can be stored either all locally or as a combination of local storage and in the Cloud Tier (if Cloud Tier is enabled). For example, if the **Data Written** is 100GB and resiliency is set to RF3, leading to the need for an additional 200GB of space, the resulting **Storage Consumed** is 300GB. If the **Data Written** is 100GB and resiliency is set to EC4:2, leading to the need for an additional 50GB (as for every 4 data fragments, 2 parity fragments are added), the resulting **Storage Consumed** is 150GB.

Metrics	Description
Physical	Data written on the Cohesity cluster after the unique logical data has been reduced by data deduplication and data compression.
Resiliency Impact	The space required to honor resiliency settings (Erasure Coding / Replication Factor) in addition to the data written. For example, if the Data Written is 100GB and resiliency setting is RF3, the Resiliency Impact is 200GB. That is, to honor the resiliency setting, you will need an additional 200GB on top of the 100GB for the Data Written.
Logical Data	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type. • **VMs**: The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. • **All Other Workloads**: The data size reported is the actual front end data residing on the server. If a server with 1TB capacity contains 100GB of data, the server reports 100GB. **Note:** • These metrics do not include unprotected objects. • Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000GB (1TB). **SmartFiles**: Size of total unique logical data of active filesystem as residing on Cohesity. This does not count the space consumed by the Data Protection runs (i.e. snapshots) that may be protecting the active filesystem.
Quota	The total amount of data that can be written to the Storage Domain (value set for **Storage Domain Physical Quota**). After it has been reduced by deduplication and compression, including the space required for honoring the resiliency settings (EC/RF). **Note:** This field is displayed as **None** if **Storage Domain Physical Quota** is not set.

Metrics	Description
Data In	Total unique logical data residing on the Cohesity cluster. For Data Protection, this is the sum of your first full backup and subsequent change data captured by Cohesity. For SmartFiles, some customers, such as SPs and enterprise customers with multiple tenants, may want to use this as the Logical Data so that snapshot capacity is included in the chargeback.
Data Written	Data written on the Cohesity cluster after the unique logical data has been reduced by data deduplication and data compression. For SmartFiles, this metric value carries the name of Physical Data. **Note:** This number reflects unique data written, before resiliency operations.
Unique Data	Data that is unique to the objects across the system.
Data Reduction	The ratio of Data In to Data Written. This ratio reflects the data deduplication and data compression operations.
Storage Reduction	The ratio of **Logical Data** to **Storage Consumed**. For Data Protection, this reflects all operations performed by Cohesity on source data, which include snapshotting, change tracking and ingest, data deduplication, data compression, and data resiliency. For SmartFiles, this reflects all operations performed on the Cohesity hosted active filesystem, which include logical data writes, change tracking, data deduplication, data compression, and data resiliency.
Storage Saved	The space owing to data deduplication and data compression. For example, if **Data In** is 100GB, deduplication ratio is 4x, compression is 2x, then **Data Written** is 12.5GB. Therefore, the **Storage Saved** is 100GB - 12.5GB = 87.5GB.
Storage Consumed Local and Cloud	If you have enabled Cloud Tiering, then this metric is the sum of space consumed on local storage and space consumed on the cloud tier.
Growth	The physical data written to the storage domain over the number of days specified in the **Growth** field.
Growth Rate	The percentage of physical data growth on the storage domain over the number of days specified in the **Growth** field.

Health Dashboard

The Health dashboard provides detailed information about the health of your Cohesity cluster. You can use this information to identify your Cohesity cluster's potential problems and address them before becoming severe.

The dashboard provides an overview of the number of Critical, Warning, and Healthy (Info) alerts triggered by the Cohesity cluster. It collates all the alerts triggered by your cluster, categorizes them based on the alert category, and then displays them as the following components:

- Hardware
- Software
- Maintenance
- Data Service

In each of these components, the alert categories are displayed as individual cards and contain the number of alerts triggered for the category. The alert categories are color-coded based on the severity of the alerts:

- **Red** - Critical
- **Yellow** - Warning
- **Green** - Informational (Info)

Hardware

In the **Hardware** section, the dashboard provides an overview of the alert categories triggered by the Cohesity hardware. You can click the alert category to view the associated alerts in the **Alerts** tab and take action, respectively.

The following are the alert categories triggered by the Cohesity hardware:

- Memory
- CPU
- NIC
- Firmware
- Disk
- Power Supply
- Fan
- Temperature
- Chassis
- Node
- Cluster

For more information on the alert categories, see Alerts Reference.

Software

In the **Software** section, the dashboard provides an overview of the alert categories triggered by the Cohesity software component. You can click the alert category to view the associated alerts in the **Alerts** tab and take action, respectively.

The following are the alert categories triggered by the Cohesity software component:

- Operating System
- App MarketPlace
- Node Health
- Metadata
- Indexing
- Data Path

For more information on the alert categories, see Alerts Reference.

Maintenance

In the **Maintenance** section, the dashboard provides an overview of the alert categories triggered by other components apart from the Cohesity hardware, software, and the data services running on the Cohesity cluster.

You can click the alert category to view the associated alerts in the **Alerts** tab and take action, respectively.

The following are the alert categories triggered by other components apart from the Cohesity hardware, software, and the data services running on the Cohesity cluster:

- Cluster Management
- Networking
- Configuration
- Fault Tolerance
- License
- Storage Usage
- Upgrade
- Security

For more information on the alert categories, see Alerts Reference.

Data Service

In the **Data Service** section, the dashboard provides an overview of the alert categories triggered by the data services running on the Cohesity cluster. You can click the alert category to view the associated alerts in the **Alerts** tab and take action, respectively.

The following are the alert categories triggered by the services running on the Cohesity cluster:

- Remote Replication
- Quota
- Archival & Restore
- Backup & Restore

For more information on the alert categories, see Alerts Reference.

Alerts

The Cohesity cluster creates an alert when:

- It finds a potential problem or when it exceeds the defined threshold.
- A Protection Group is configured to trigger an alert indicating the success or failure of the job.

You can configure general alert email notifications in the **System > Health** > **Settings** tab and job-specific email notifications in a Protection Group.

Each alert has a severity rating that indicates the seriousness of the problem:

- **Critical**—Immediate action is required because it detects a severe problem that might be imminent or major functionality is not working, such as a disk failure.
- **Warning**—Action is required, but the affected functionality is still working, such as the disk space on the Cohesity cluster is 80% to 90% full.
- **Informational**—Immediate action is not required, and the alert provides an informational message.

The **Alerts** tab displays all the alerts that are generated by the Cohesity cluster. Alerts are grouped into one alert (with multiple occurrences) if the alerts are of the same type, and they are reporting the same Object and occur within the specified time-frame of the alert.

Analyze the Alert

From the **Alerts** tab, you can click on an alert and view the alert details on the **Details for <Alert_Name>** page.

The **Details for <Alert_Name>** page includes a timeline view that shows the date and time when the alert was triggered. The page also provides the following details of the alert:

Details	Description
Alert Code	The alert code. You can click on the alert code for detailed information about the alert.

Details	Description
Severity	The severity rating of the alert.
Type	The alert type. It defines the Cohesity component that triggered the alert.
Category	The alert category.
Status	The status of the alert. It can be Open, Resolved, or Note.
Description	A brief description of the problem that triggered the alert.
Cause	A brief description of the cause of the problem.

Alert Notification

You can configure general alert email notifications in the **System** > **Health** > **Settings** tab and job-specific email notifications in a Protection Group. For more information, see Configure Alert Notification Settings.

Alert Dedup

The **Alerts** tab displays all the alerts generated by the Cohesity cluster. Alerts are grouped into one alert (with multiple occurrences) if the alerts are of the same type, reporting the same Object, and occur within the specified time frame of the alert. For information on the deduplication interval, see Alerts Reference.

Resolve Alerts

In case you are aware of the problem and confirm that the issue has been resolved or if the issue does not require further attention, from the **Alerts** tab, you can manually resolve those alert(s).

You can either create a new resolution of the alert(s) or attach an existing resolution to the alert(s).

> **Note:** If you want to view the details of an alert and resolve it, see Resolve an alert in the Details for <Alert_Name> page.

Create a new Resolution

To create a new resolution:

1. In the **Alerts** tab, select an alert or multiple alerts that you plan to resolve and click **Resolve Selected Alerts**.

2. On the **Resolution** page, do the following:
    1. Select Create new resolution.
    2. In the Resolution Summary field, add a resolution summary for the alert.
    3. In the Resolution Description field, add a brief description of the resolution.
    4. Click Resolve.

The resolution is added to the selected alerts, and the status of the alert(s) is marked as **Resolved**.

Attach an Existing Resolution

To attach an existing resolution to the alert(s):

1. In the **Alerts** tab, select an alert or multiple alerts that you plan to resolve and click **Resolve Selected Alerts**.
2. On the **Resolution** page, do the following:
    1. Select **Associate with existing resolution**.
    2. From the **Resolution Summary** drop-down, you can search and select the resolution that you plan to attach to the alert.
    3. Click **Resolve**.

The existing resolution is attached to the selected alerts, and the status of the alert(s) are marked as **Resolved**.

Resolve an alert in the Details for <Alert_Name> page

Once you have reviewed the alert, you can resolve the alert using the page's **Resolution** section. In the **Resolution** section, you can either create a new resolution for the alert or attach an existing resolution.

To create a new resolution:

1. In the **Resolution** section, select **Create new resolution**.
2. In the **Resolution Summary** field, add a resolution summary for the alert.
3. In the **Resolution Description** field, add a brief description of the resolution.
4. Click **Resolve**.

To attach an existing resolution:

1. In the **Resolution** section, select **Associate with existing resolution**.
2. From the **Resolution Summary** drop-down, you can search and select the resolution that you plan to attach to the alert.
3. Click **Resolve**.

Auto-Resolve an Alert

Some alerts are automatically resolved when the dependent success alert is triggered. For example, `CreateSnapshotFailedAlert` is auto-resolved when the

`CreateSnapshotSucceededAlert` is triggered. The complete list of such dependencies is available in the Alerts Reference.

Alert Retention

For **Critical** and **Warning** alerts, the retention period is six months. On the other hand, informational alerts are stored only for 14 days; after that, the alerts are not accessible.

> **Note:** You cannot customize the retention period for alerts.

View Alert Resolution Summary

From the **Resolution Summary** tab on the Health dashboard, you can summarize the resolutions added to the alerts.

View the Resolution Summary

To view the resolution summary:

1. Navigate to **System > Health** and click the **Resolution Summary** tab.
2. In the **Resolution Summary** page, you can click on the Resolution summary to view the alerts that are resolved using the resolution. Also, you can click on the alert to view the alert details.

Configure Alert Notification Settings

From the **Health** dashboard, you can configure general alert notification rules in the **Settings** tab. You can configure email, Syslog, SNMP, and webhook as the notification output for the alert notification.

> **Important:**
> In multi-tenant deployments, the rules for alert notifications work as follows:
>
> - If the notification rules are set up at the organization level, only alerts raised at that level will trigger notifications.
> - If the notification rules are set up at the cluster level, and alerts are raised at the organization level, all organizations' alerts will trigger notifications.
>
> Additionally, notifications will be sent if the alert's category, severity, and name match the specified filters.

> **Note:** Alert notifications are sent once per active alert within the dedup interval.

Add Alert Notification Rule

You can add different alert notification rules that send emails, SNMP, Syslog, and/or cURL HTTP POST requests to a webhook URL based on the alert categories, severities, and names. Rules can be configured at both the service provider and the tenancy level.

Prerequisites

Before you add an alert notification rule, make sure that you:

- Configure an SMTP Server on the Cohesity cluster, if you plan to send email notifications when an alert is triggered. For details on configuring the SMTP server, see Configure Settings.
- Configure the SNMP setting on the Cohesity cluster, if you plan to send SNMP notifications to your network management system when an alert is triggered. For details on configuring SNMP settings, see Configure and Test SNMP.
- Add a Syslog server to the Cohesity cluster, if you plan to send Syslog notifications to the server when an alert is triggered. For details on configuring the Syslog server, see Add a Syslog Server.

To add an alert notification rule:

1. Navigate to **System > Health** and click the **Settings** tab.
2. Click **Add Alert Notification Rule**.
3. On the **Add Alert Notification Rule** page, do the following:
    1. **Rule Name**: Type a descriptive name for the rule, for example, Cluster Disk Offline or Backup Job Failure.
    2. **Alert Category**: Optional. Select one or more categories from the drop-down. Otherwise, all alerts in any category will trigger the rule.
    3. **Alert Severities**: Optional. Select one or more severities from the drop-down. Otherwise, all alerts with any severity will trigger the rule.
    4. **Alert Name**: Optional. Select one or more names from the drop-down. Otherwise, any Alert name will trigger the rule. If you selected any categories, the list includes only alerts in those categories.
    5. **Email**: Optional. Click **Add** and from the drop-down menu, perform the following:
        1. Select **To** and type an email address or distribution list of the recipients to whom you plan to send the email notification.
        2. Select **CC** and type an email address or distribution list of the recipients to whom you plan to send a copy of the email notification.
    6. **SNMP**: Optional. Toggle on if you want the rule to send SNMP notifications when the alert is triggered.
    7. **Syslog**: Optional. Toggle on if you want the rule to send Syslog notifications when the alert is triggered.

8. **Webhook**: Optional. Toggle on and provide the URL and cURL details if you want the rule to send a web callback to a server. Type the URL for the server and type any cURL options you want to use. When an alert is triggered, the Cohesity cluster sends an HTTP request to the server. Your application can then interpret the request. For example, the webhook might notify the server about a critical disk alert, and your application might open a trouble ticket to track the problem. For more information, see Configuring Webhooks.

4. Click **Save**.

Edit Alert Notification Rule

You can edit an existing alert notification rule and update the following:

Change the rule name.

- Select or discard alert categories from the **Alert Categories** list.
- Select or discard alert severities from the **Alert Severities** list.
- Select or discard alert names from the **Alert Name** list.
- Add or remove email addresses
- Enable or disable SNMP notification
- Enable or disable Syslog notification
- Enable or disable Webhook.

The updated alert notification rule is applied from the next alert that is triggered.

To edit an existing alert notification rule:

1. Navigate to **System** > **Health** and click the **Settings** tab.
2. Click the edit icon next to the rule.
3. On the **Edit Alert Notification Rule** page, update the settings as per your requirement.
4. Click **Save**.

Delete Alert Notification Rule

You can delete the alert notification rules that you do not want to use for alert notifications.

To delete alert notification rule:

1. Navigate to **System** > **Health** and click the **Settings** tab.
2. Click the delete icon next to the rule.
3. In the **Delete Alert Notification Rule** dialog box, click **Delete**.

Configuring Webhooks

Webhooks are HTTP callbacks which are usually triggered by some event. Webhooks are configured on one website and when an event occurs on this website, an HTTP request is made to the configured URL which invokes action on the other website.

You can enable webhooks for Cohesity alerts by creating an alert notification rule. When the alert is triggered and meets the criteria in the rule, the cluster sends an HTTP request to the specified website. Your application can then interpret the request. For example, the webhook might notify the website about a critical disk alert and your application might open a trouble ticket to track the problem.

To create an alert notification rule and enable a webhook:

1. Navigate to **System > Health** and select the **Alerts** tab.

   The **Alerts** summary page is displayed.
2. Select the **Settings** tab, and click **Add Alert Notification Rule**.
3. Configure the following details:

Parameter	Description
Rule Name	Alert rule name.
Alert Category/Severity/Name	Choose the alert category/severity and name from the list. You can specify one or multiple values.
Send Alert Notification Via Email	The configured email addresses will receive alert email notifications.
Send Alert Notification via Webhook	The URL to which the system will connect when the above alert notification rule is matched. That is the alerts will be notified to the configured URL.

4. Toggle on **Webhook** and provide the URL and cURL options.

> **Note:** Ensure that the Webhook configuration follows the rfc3986 standards.

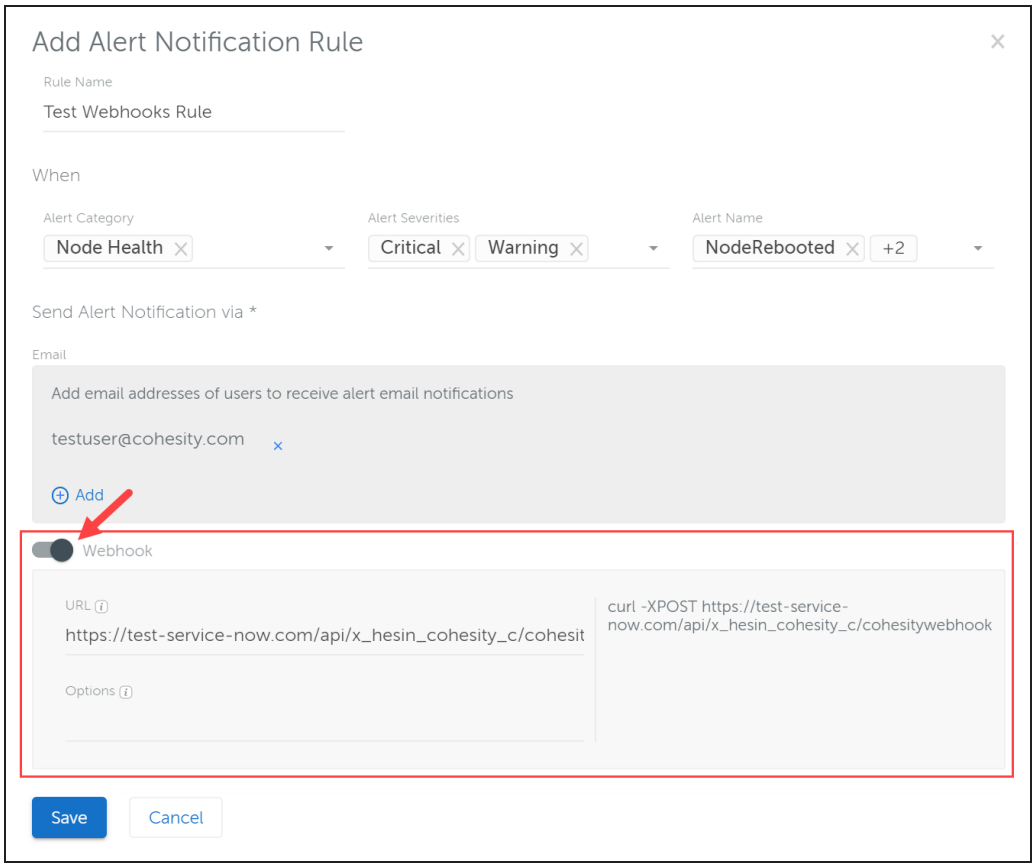

5. Click **Save**.

Alert Request

When an alert is triggered, a sample payload as shown below will be available at the configured URL:

CURL Request:

```
Curl -XPOST https://test-service-now.com/api/x_hesin_cohesity_c/co-
hesitywebhook
```

The Payload sent to the above URL:

```
{
       "alertType": "1060",
       "alertId": "437441:1665640208440874",
       "alertName": "DnsServerUnReachable",
       "alertDescription": "DNS server is not reachable.",
       "alertSeverity": "WARNING",
       "alertCategory": "Networking",
       "alertCause": "DNS servers [10.2.0.1] unreachable. Please update DNS
Server.",
       "alertHelpText": "Check and provide alternate DNS server.",
```

```
       "languageCode": "en-us",
       "alertUrl": "https://cluster-ip/health/alert-info/437441:1665640208440874",
       "clusterName": "cohesity",
       "clusterId": "4398906079065976",
       "alertCode": "CE03601060",
       "alertProperties":
       {
              "cluster_id": "4398906079065976",
              "alert_message": "DNS servers [10.2.0.1] unreachable. Please update DNS
Server."
       },
       "firstOccurrenceUsecs": "1665640208440874",
       "lastOccurrenceUsecs": "1665640208440874"
}
```

> **Note:** *prod-cluster.eng.cohesity.com* refers to the IP address of the cluster.

Customized Template for URL Request and Data Payload

Webhook also supports template options for both URL requests and data payload.

Data payload can be sent in any customized format as a template with placeholders. These placeholders are replaced with the actual values of the generated alert.

The placeholders can be any of the following:

```
"$alertType", "$alertId", "$alertName", "$alertDescription", "$aler-
tSeverity", "$alertCategory", "$alertCause", "$alertHelpText", "$lan-
guageCode", "$alertUrl", "$clusterName", "$clusterId", "$alertCode",
"$alertProperties", "$firstOccurrenceUsecs", "$lastOccurrenceUsecs”
```

The following is an example of a template for a Slack channel:

```
--template {"blocks": [{ "type": "section", "fields": [{"type": "plain_
text", "text": "alert name" }, {"type": "mrkdwn", "text": "*$alertName*"},
{"type": "plain_text", "text": "alertSeverity"}, {"type": "mrkdwn", "text":
"*$alertSeverity*"}, {"type": "plain_text", "text": "alertCategory"},
{"type": "mrkdwn", "text": "*$alertCategory*"}, {"type": "plain_text",
"text": "alertDescription" }, {"type": "mrkdwn", "text": "*$aler-
tDescription*" }, {"type": "plain_text", "text": "alertCause" }, {"type":
"mrkdwn", "text":"*$alertCause*"}]}]}
```

Webhook URL path can also be customized with placeholders.

Example:

`"http://example.webhooksite.com?alertName=$alertName"`, where `alertName=$alertName` is the placeholder.

[content type]

If the content-type is not specified, by default, it is application/json. If a content-type is specified it must be acceptable by the webhook server.

Example:

```
-H content-type:text/html
```

[curl options]

For security reasons, valid webhook options can only be `-H( or --header)`, `-d(or --data)` and `--template`. If `--template` is provided then `-d(or --data)` is ignored. If `--template` is not provided then the payload in `--data` is merged with the default alert JSON content. Ensure that you provide a valid JSON entry for the `-d(or --data)` option.

Example:

```
-d {"accountId": 12345, "accountName": "test" }
```

Resolving Alerts

When an alert is triggered at the Cohesity cluster side, the above data as shown in the sample payload must be notified to the configured URL. Using the above payload, the alerts can be resolved at the target based on the **alertID**. This eliminates the need to switch back to the Cohesity cluster to manage and resolve alerts.

To resolve alerts, the alertURL from the sample payload must be extracted and used in the API request as shown below:

URL Request:

```
https://prod-cluster-
.eng.cohesity.com/irisservices/api/v1/public/alertResolutions
```

Method : POST

Payload :

```
{"resolutionDetails":{"resolutionSummary":"Resolved at 11/1, high pri-
ority"},"alertIdList":["157373326:1577990357283489"]}
```

> **Note:** *alertIdList* refers to the alertID in the payload. *Resolved at 11/1, high priority* is a user defined string.

Alerts Reference

Download a spreadsheet containing 7.4.1 alerts.

Hardware Alerts

Software Alerts

Data Service Alerts

Maintenance Alerts

Hardware Alerts

Download a spreadsheet containing **Hardware** alerts.

Disk Alerts

Node Alerts

Power Supply Alerts

CPU Alerts

Chassis Alerts

Memory Alerts

Temperature Alerts

FAN Alerts

NIC Alerts

Firmware Alerts

Disk Alerts

CE00101009 MissingDisk

Alert Description: The node is missing disks.

Reason: This alert is generated when the number of disks assigned to a node does not match the number of disks supported by the hardware model of the node. The alert is triggered once per hour as long as a disk is missing. This alert can be caused by one of the following conditions.

- The Cohesity cluster was created with fewer disks than supported by the hardware model of the node.
- A node with fewer disks than supported by the hardware model of the node, is added to the Cohesity cluster. The disk(s) may be pulled out or not completely inserted.

Action: Find the node with the issue and add enough disks as supported by the hardware model of the node. This may involve pushing in a partially pulled-out disk or inserting a new disk if a slot in the node is empty and following the instructions in the hardware replacement procedure.

Severity: Warning

Related Error Code: Not Applicable

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.8

Platform: All

CE00101034 BootDiskHealth

Alert Description: The boot disk may be unhealthy on the node.

Reason: This alert is triggered when a boot disk might be unhealthy.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.168

Platform: All

CE00101071 DiskAddedToBlacklist

Alert Description: A disk is added to the denylist on the node.

Reason: This alert is triggered when a disk is added to the denylist on a node, either manually or through auto-detection of a bad disk.

Action: Contact Cohesity Support.

Related Error Code:

Error Codes	Description
kMissingStaticRoutes	The alert with the error code is raised when there are missing routes in the routing table than what is configured in the Cohesity cluster configuration.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.119

Platform: All

CE00101073 DiskSlotAvailable

Alert Description: The node has unused disk slots available

Reason: This alert is triggered when there is periodic monitoring of the system. The alert can be caused if the current number of disks is less than the expected number of disks on a platform.

Action: No action needed.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.120

Platform: All

CE00102203 DiskAwaitHigh

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The wait time for disk ID is too high.

Reason: This alert with a Critical severity, is triggered when the average time for I/O requests issued to a disk exceeds 500ms for 5 minutes.

Action: Contact Cohesity Support to investigate the disk problem and determine if the disk needs to be replaced.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.159

Platform: All

CE00102203 DiskAwaitHigh

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The wait time for disk ID is too high.

Reason: This alert with a Warning severity, is triggered when the average time for I/O requests issued to a disk exceeds 400ms for 5 minutes.

Action: Contact Cohesity Support to investigate the disk problem and determine if the disk needs to be replaced.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.159

Platform: All

CE00109015 AccessToExternalTargetFailed

Alert Description: Access to external target failed.

Reason: This alert is triggered when access to an External Target fails due to a permission denied error or network connectivity error. This alert is triggered once a day.

Action: Verify the login credentials or network connectivity.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.91

Platform: All

CE00113003 WriteLimit

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The disk with the disk ID is reaching its write limit on the node.

Reason: This alert, with a Critical severity, is triggered when the SSD write cycle exceeds 95%. The SSD is exceeding write thresholds, so it may need to be replaced soon.

Action: Contact Cohesity Support to investigate the problem with the SSD and determine if the SSD needs to be replaced.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.9

Platform: All

CE00113003 WriteLimit

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The disk with the disk ID is reaching its write limit on the node.

Reason: This alert, with a Warning severity, is triggered when the SSD write cycle exceeds 75%. The SSD exceeds write thresholds, so it may need to be replaced soon.

Action: Contact Cohesity Support to investigate the problem with the SSD and determine if the SSD needs to be replaced.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.9

Platform: All

CE00113012 DriveRemoved

Alert Description: The drive was disconnected or removed.

Reason: This alert is triggered when the removal of a Hard Disk Drive (HDD) is detected.

Action: No action is required if you intentionally removed the HDD. If it was unintentional, check if the HDD is correctly inserted into its slot.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.17

Platform: C2000 Series Only

CE00113014 DriveFault

Reason: This alert is triggered when a Hard Disk Drive (HDD) fault is detected by monitoring the IPMI SEL log.

Action: Contact Cohesity Support to investigate the problem with the HDD and determine if the HDD needs to be replaced.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.19

Platform: C2000 Series Only

CE00101001 MarkDiskForRemoval

Alert Description: The disk is marked for removal.

Reason: This alert is triggered when a disk previously removed from the Cohesity cluster is inserted back in the node. The alert is triggered for both types of disk removals:

- A graceful removal of a disk where the disk is marked offline before being pulled as described in hardware replacement.
- An ungraceful removal of a disk. The disk is not marked offline before it is pulled.

Action: If you unintentionally removed a disk, just reinsert the disk back into the slot. If you intentionally removed the disk but did not follow the graceful removal process, reinsert the disk. For more information, see hardware replacement. You can ignore this alert if you intentionally removed the disk and followed the hardware replacement procedure.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.3

Platform: All

CE00101005 NewDiskFound

Alert Description: A new disk is physically detected on the node.

Reason: This alert is triggered when a new disk is inserted into a node. This alert is triggered every 10 minutes until the new disk is added to the node.

Action: Add a new disk by following the instructions in hardware replacement.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.5

Platform: All

CE00101026 DiskAddedToCluster

Alert Description: A disk on the node is logically added to the cluster.

Reason: This alert is triggered if the Cohesity cluster detects that a new disk is added.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.30

Platform: All

CE00101037 UsbDiskFound

Alert Description: USB disk is detected on the node.

Reason: The cluster has detected a USB disk attached to the node. This alert is triggered when the cluster has detected a USB disk attached to the node. It is triggered daily until resolved.

Action: If this is an expected event, for example, the cluster is being upgraded with software on a USB, no action is required. Verify that the USB was not inadvertently left in the node after the upgrade was completed.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.167

Platform: All

CE00101037 UsbDiskFound

Alert Description: USB disk detected on the node.

Reason: This alert is triggered when the cluster has detected a USB disk attached to the node. It is triggered daily until resolved.

Action: If this is an expected event, for example, the cluster is being upgraded with software on a USB, no action is required. Verify that the USB was not inadvertently left in the node after the upgrade was completed.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.167

Platform: All

CE00109001 DiskMarkedOffline

Alert Description: The disk is marked offline.

Reason: This alert is triggered when read/write errors occur during disk I/O operations. The alert can be caused by one of the following conditions:

- Disk read/write fails with an error.
- Disk read/write access is hung and does not complete in a reasonable time.

Action: Contact Cohesity Support to investigate the disk problem and determine if the disk needs to be replaced.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.101

Platform: All

CE00101027 ClusterPartitionCreateFailed

Alert Description: Unable to create cluster partition on the cluster

Reason: This alert is triggered when a user populates the cluster Partition fields (VIPs and Hostname) during cluster Create, but the creation fails.

Action: Create the cluster Partition by using the Cohesity Dashboard.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.182

Platform: All

CE00101092 DiskRemovalCancelled

Alert Description: Removing a disk on the node is cancelled.

Reason: This alert is triggered when a disk removal is cancelled.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.239

Platform: All

CE00101083 DiskRemoved

Alert Description: The node bond mode has been migrated.

Reason: This alert is triggered when the disk is removed from the Cohesity cluster.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.222

Platform: All

CE00101100 NfsShareNotHealthy

Alert Description: The NFS share mounted to the node is not healthy.

Reason: This alert is triggered when the Cohesity cluster cannot access the NFS share from Pure FlashRecover. This alert can be caused by one of the following conditions:

- The NFS share IP address might be down.
- Export rule issues.
- The Pure Flash Array is down.

Action: Check Pure Flash Array.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.249

Platform: PXG1

CE00113035 SsdMediaError

Alert Description: Ssd Media Error

Reason: This alert is triggered when there is a media error on the SSD.

Action: Update the SSD firmware.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.314

Platform: All

CE00113035 SsdMediaError

Alert Description: Ssd Media Error

Reason: This alert is triggered when there are more than 200 media errors on the SSD.

Action: Update the SSD firmware.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.314

Platform: All

CE00113036 DiskNeedFirmwareUpgrade

Alert Description: The disk needs a firmware upgrade on the node.

Reason: This alert is triggered if the SSD firmware version has a known issue related to an SSD media error. When there is no media error, the alert is info. When there is any media error, the alert is critical.

Action: Update the SSD firmware.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.315

Platform: All

CE00113036 DiskNeedFirmwareUpgrade

Alert Description: The disk needs a firmware upgrade on the node.

Reason: This alert is triggered if the SSD firmware version has a known issue unrelated to an SSD media error.

Action: Update the SSD firmware.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.315

Platform: All

CE00101115 DiskOkToRemove

Alert Description: A disk on the node is ready to be removed.

Reason: This alert is triggered when a disk is ready to be physically removed from the node in the Cohesity cluster. The alert indicates that data healing is completed on the disk.

Action: You can physically remove the disk from the node for replacement.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.293

Platform: All

CE00113001 DiskIsBad

Alert Description: Disk is bad.

Reason: This alert is triggered if uncorrected errors are detected on the disk or if the disk temperature exceeds the set threshold.

Action: Contact Cohesity Support to investigate the problem with the disk and determine if the disk needs to be replaced.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.66

Platform: All

CE00113013 DrivePresent

Reason: This alert is triggered when the upgrade or link has issues after the node rebooted due to the upgrade.

Action: Verify if the link issue is caused by the system reboot or upgrade.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.214

Platform: All

CE00101029 DiskNotHealthy

Alert Description: A disk on the node is not healthy.

Reason: The disk is detecting read/write errors. This alert is triggered when one of the following conditions occur:

- When uncorrectable errors are detected on a disk.
- When the disk temperature exceeds the threshold set by the disk vendor.

Action: Contact Cohesity Support to investigate the disk problem and determine if the disk needs to be replaced.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.1

Platform: All

CE00109034 DiskMarkedForRemoval

Alert Description: Disk marked for removal.

Reason: This alert is triggered when the disk initialization process fails.

Action: Contact Cohesity Support

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.219

Platform: All

CE00110028 AzureCEHostCachingEnabled

Alert Description: In the case of Azure Cloud Edition clusters, if Host Caching is not set to 'None' on Azure nodes/VMs, it might lead to potential data loss.

Reason: This alert is triggered when the Azure disk has HostCaching enabled. One alert is triggered every 24 hours.

Action: Refer to the Host caching configuration for Azure Cloud Edition cluster KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.338

Platform: Azure Cloud Cluster

CE00101134 DiskRemovalStuck

Alert Description: The disk removal process is stuck.

Reason: This alert is triggered if there are any services stuck for giving an ACK for removal when a disk removal is initiated and the threshold limit (default 24 hrs) has exceeded. This alert can be caused by a backend component like ChunkRepo, Apollo or Yoda not ACKing removal due to an internal failure.

Action: Contact Cohesity Support to investigate why the component is not ACKing removal. The component that failed to ACK removal needs to be triaged and directed to the appropriate squad/team.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.366

Platform: All

Node Alerts

CE00201023 RaidDegraded

Alert Description: The node's RAID is degraded.

Reason: This alert is triggered when the Cohesity cluster determines that a node in the cluster does not have a redundant system disk. It is typically triggered when a boot SSD or NVMe SSD is offline or failed.

Action: Check whether there is an offline SSD. If yes, fix or replace the offline SSD. Otherwise, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.29

Platform: All

CE00201038 RaidInactive

Alert Description: The node has inactive RAID disks.

Reason: This alert is triggered to ensure that new disks are used, or existing disks are cleaned before RAID rebuild.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.160

Platform: All

CE00201042 LessCPUCore

Alert Description: The node does not meet minimum CPU core count requirements.

Reason: This alert is triggered when the node does not meet the minimum CPU core count requirements. It is triggered every hour.

Action: Review the CPU available for VM and increase the vCPU.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.205

Platform: Virtual Edition

CE00213018 NodeRemoved

Alert Description: The node was disconnected or removed

Reason: This alert is triggered when a node is removed from a chassis.

Action: If the node was intentionally removed, no action is required. If this alert was triggered for an unknown reason, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.87

Platform: C4000 Series Only

CE00213019 NodeInserted

Alert Description: The node was inserted

Reason: This alert is triggered when a node is inserted into a chassis.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.88

Platform: C4000 Series Only

CE00213027 SystemLedAmber

Alert Description: Node system LED.

Reason: This alert is triggered when a hardware issue is detected.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.31

Platform: C2000 Series Only

CE00201019 ProductModelIncompatible

Alert Description: Product models are not compatible on the node.

Reason: This alert is triggered when the current combination of product models for the nodes in a Cohesity cluster is not supported.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.69

Platform: All

CE00201024 VirtualApplianceResourcesUnexpected

Alert Description: The resources for the virtual appliance on the node are not matching Cohesity recommendations

Reason: This alert is triggered when a Virtual Machine hosting a Cohesity cluster does not have the recommended values for vCPUs, memory, SSD etc. For details, see the appropriate Virtual Edition Setup Guide from Initial Virtual Edition Cluster Setup.

Action: Refer to the appropriate Virtual Edition Setup Guide for the recommended resources and make corrections.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 31536000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.181

Platform: Virtual Edition

CE00201079 NodeVirtualizationSupportDisabled

Alert Description: Virtualization support is disabled on the node.

Reason: This alert is triggered when any configuration change has been checked on. Virtualization is disabled in the platform firmware. Other possible reasons are kernel or driver errors.

Action: Verify if the virtualization support option is enabled in the platform firmware.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.60

Platform: All

CE00201091 SilicomNicFound

Alert Description: Found Silicom NIC in the node

Reason: This alert is triggered when the Silicom Network Interface Card (NIC) is detected.

Action: Verify if the port numbering is in reverse order, similar to other NICs.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.238

Platform: All

CE00201135 NodeRemovalStuck

Alert Description: The node removal process is stuck.

Reason: This alert is triggered if there are any services stuck for giving an ACK for removal when a node removal is initiated and the threshold limit (default 24 hrs) has exceeded. This alert can be caused by a backend component like ChunkRepo, Apollo or Yoda not ACKing removal due to an internal failure.

Action: Contact Cohesity Support to investigate why the component is not ACKing removal. The component that failed to ACK removal needs to be triaged and directed to the appropriate squad/team.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.365

Platform: All

Cluster Alerts

CE00301039 RemoveNodeSuccess

Alert Description: Removal of the node is successful

Reason: This alert is triggered daily.

Action: If a node is intentionally being removed from the cluster, no action is required. Otherwise, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00301040 ForcefulRemoveNodeWarning

Alert Description: The node is forcefully being removed from the Custer.

Reason: This alert is triggered daily.

Action: If a node is intentionally being removed from the cluster, no action is required. Otherwise, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: Not Applicable

Platform: All

CE00301041 AddNodeSuccess

Alert Description: The node has successfully joined the cluster.

Reason: This alert is triggered every hour.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00301053 VerifyRemoteClusterConnection

Alert Description: The cluster needs to confirm a connection to a remote Cluster. This alert is triggered yearly until resolved.

Action: Do one or more of the following:

- Verify the remote cluster connection on both Clusters: Log in to the Cohesity Dashboard, select **Infrastructure > Remote Clusters** and verify the remote cluster connection and details.
- Verify the remote cluster connection by using the CLI on both Clusters. Run the command: `iris_cli remote-cluster ls`.
- Verify the static routes on both Clusters by using the CLI. Run the command: `iris_cli route ls`.
- Verify the static route file, for example:

```
[cohesity@node-1 ~]$ cat /etc/sysconfig/network-scripts/route-
bond1198.16.0.0/20 via 10.2.32.1 dev bond1198.16.10.179/32 via
10.2.32.1 dev bond1198.16.10.180/32 via 10.2.32.1 dev
bond1198.16.10.181/32 via 10.2.32.1 dev bond1
```

Severity: Warning

SNMP OID: Not Applicable

Platform: All

CE00301059 NtpServerUnReachable

Alert Description: The NTP server is not reachable for the cluster.

Reason: This alert is triggered when the NTP server is not reachable.

Action: Check the state of the NTP server and provide an alternate NTP server if necessary.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: Not Applicable

Platform: All

CE00301060 DnsServerUnReachable

Alert Description: The DNS server is not reachable for the cluster.

Reason: This alert is triggered when a DNS server is not reachable. Incorrect input was provided for creating the cluster or there is a networking issue in the current node.

Action: Check the state of the DNS server and provide an alternate DNS server if necessary.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00301060 DnsServerUnReachable

Alert Description: The DNS server is not reachable for the cluster.

Reason: This alert is triggered when a DNS server is not reachable. Incorrect input was provided for creating the cluster or there is a networking issue in the current node.

Action: Check the state of the DNS server and provide an alternate DNS server if necessary.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: Not Applicable

Platform: All

CE00301061 ClusterGatewayUnReachable

Alert Description: The cluster gateway is not reachable from the node.

Reason: This alert is triggered when a cluster gateway is not reachable. Incorrect input was provided for creating the cluster or there is a networking issue in the current node.

Action: Check the state of the cluster gateway and provide an alternate gateway if necessary.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00301062 IpmiNetworkConfigFailed

Alert Description: IPMI network config failed on the node.

Reason: This alert is triggered when an IPMI network configuration fails on a node. Incorrect input was provided for creating the cluster or there was a networking issue in the current node.

Action: Check and reconfigure the IPMI on the node.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: Not Applicable

Platform: All

CE00301063 IpmiUserUpdateFailed

Alert Description: IPMI user update failed on the node.

Reason: This alert is triggered when an IPMI user update fails on a node. Incorrect input was provided for creating the cluster or there is a networking issue in the current node.

Action: Check and reconfigure the IPMI user.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: Not Applicable

Platform: All

CE00301064 LowFaultTolerantViewBox

Alert Description: Low fault tolerant storage domain found for the cluster.

Reason: This alert is triggered when a low fault tolerant Storage Domain is found. The metadata failure tolerance factor for the cluster was increased to a value higher than the failure tolerance factor of the Storage Domain mentioned in the alert.

Action: Consider increasing the fault tolerance/EC factor of the Storage Domain to match the metadata failure tolerance for the cluster.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00301065 VerifyVlanHostName

Alert Description: Error while verifying hostname entry in DNS server for the interface group.

Reason: This alert is triggered during verification of the hostname entry on the DNS server for the specified interface group.

Description: Incorrect input was provided for creating the cluster or there is a networking issue in the current node.

Action: Change the hostname in the DNS server or enter the hostname that matches the hostname in the DNS server.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00301074 IpmiConfigAbsent

Alert Description: IPMI configuration is absent on the cluster.

Reason: IPMI is not configured on the physical cluster.

Action: Configure IPMI.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00302122 StorageDomainSpaceUsageHigh

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The Storage Domain is full.

Reason: This alert with a Critical severity is triggered when a Storage Domain is over 90% full. After the alert is initially raised, if the Storage Domain continues to be over 90% full, this alert is raised every hour.T

Action: Do one or more of the following:

- Add more nodes to your Cohesity cluster.
- Clean up Snapshots that are no longer required.
- Archive Snapshots to an External Target.
- Configure Cloud Tier to hold rarely used data.
- Consider changing Storage Domains to use Erasure Coding. Capacity utilization will be reduced over time in the background.
- See the Alert: CE03802121/CE03802122 Cluster Space Usage High KB article.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.105

Platform: All

CE00302122 StorageDomainSpaceUsageHigh

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The Storage Domain is full.

Reason: This alert with a Warning severity is triggered when a Storage Domain is 80%-90% full. After the alert is initially raised, if the Storage Domain continues to be 80%-90% full, this alert is raised every ho

Action: Do one or more of the following:

- Add more nodes to your Cohesity cluster.
- Clean up Snapshots that are no longer required.
- Archive Snapshots to an External Target.
- Configure Cloud Tier to hold rarely used data.
- Consider changing Storage Domains to use Erasure Coding. Capacity utilization will be reduced over time in the background.
- See the Alert: CE03802121/CE03802122 Cluster Space Usage High KB article.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.105

Platform: All

CE00302123 NodeFailureIsNotTolerated

Alert Description: Effective usage space in the cluster is high, the cluster may not have enough space to heal if a node failure happens.

Reason: This alert is triggered when the Cohesity cluster detects that a *potential* node failure could result in an effective Cohesity cluster disk usage of 95% or greater. This alert Cause field indicates the current *actual* disk usage, which is typically less than 95%. If a node in the Cohesity cluster fails, its contents need to be copied to the other nodes. If the Cohesity cluster detects that the contents of a potentially failed node cannot be copied to the remaining nodes due a lack of disk space, this alert is triggered to indicate that node failures cannot be tolerated.

Action: Do one or more of the following:

- Add more nodes to your Cohesity cluster.
- Clean up Snapshots that are no longer required.
- Archive Snapshots to an External Target.
- Configure Cloud Tier to hold rarely used data.
- Consider changing Storage Domains to use Erasure Coding. Capacity utilization will be reduced over time in the background.
- See the Alert: CE03802121/CE03802122 Cluster Space Usage High KB article.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.106

Platform: All

CE00313030 SupportChannelAboutToClose

Alert Description: The Support Channel is about to be closed after being opened temporarily.

Reason: This alert is triggered every five minutes until the Support Channel is closed.

Action: Reopen the Support Channel if necessary. For more information, see Manage the Support Channel.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: Not Applicable

Platform: All

CE00316002 ClusterSpaceUsageHigh

Alert Description: The cluster is full.

Reason: Cohesity Proactive Support predicts that the specified cluster will reach 90% disk utilization in <nn> days.

Action: For information on reclaiming space on the cluster, see the Cohesity KB article ALERT: CE00502121 Cluster Space Usage High.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: Not Applicable

Platform: All

CE00301099 RemoteStorageFileSystemUsageHigh

Alert Description: Pure Flash Array file system usage is high

Reason: This alert is triggered when the file system usage of Pure FlashRecover exceeds 80%.

Action: Increase the Pure Flash Array capacity.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 248

Platform: PXG1

CE00301098 RemoteStorageFileSystemDeleted

Alert Description: File system on Pure Flash Array is deleted.

Reason: This alert is triggered when the file system on Pure FlashRecover is deleted.

Action: Check the Pure Flash Array.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 247

Platform: PXG1

CE00301097 RemoteStorageFileSystemMountFailed

Alert Description: Failed to mount the NFS file system.

Reason: This alert is triggered when the Cohesity cluster has failed to mount the NFS file systems from Pure FlashRecover.

Action: Check the Pure Flash Array.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 246

Platform: PXG1

CE00301096 RemoteStorageManagementIpUnreachable

Alert Description: Remote storage Management IP is unreachable.

Reason: This alert is triggered when the Pure FlashRecover data or management IP address is not reachable.

Action: Check the Pure Flash Array network configuration.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 245

Platform: PXG1

CE00301094 RemoteStorageUsageHigh

Alert Description: Remote storage usage is high

Reason: This alert is triggered when the total capacity usage of Pure FlashRecover exceeds 80%.

Action: Increase the Pure Flash Array capacity.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 242

Platform: PXG1

CE00301093 RemoteStorageCapacityChanged

Alert Description: Remote storage capacity has changed.

Reason: This alert is triggered when the capacity of Pure FlashRecover has changed.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 243

Platform: PXG1

CE00301106 RemoteStorageDataIpUnreachable

Alert Description: Pure Flash Array is not reachable.

Reason: This alert is triggered when the IP address of Pure FlashRecover is not reachable.

Action: Verify the Pure Flash Array network.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 261

Platform: PXG1

CE00301118 LowReplicationFactor

Alert Description: Cluster is on Replication Factor 1. There is no resiliency to handle disk failure.

Reason: If you are using a single node Cohesity cluster, the alert is triggered to notify the risks of low replication factor. This alert is triggered once, until resolved by the customer.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: N/A

Platform: ROBO Edition

CE00301118 LowReplicationFactor

Alert Description: Cluster is on Replication Factor 1. There is no resiliency to handle disk failure.

Reason: If you are using a single node Cohesity cluster, the alert is triggered to notify the risks of low replication factor. This alert is triggered once, until resolved by the customer.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: N/A

Platform: ROBO Edition

CE00301110 RemoteStorageUnexpectedNfsExportRules

Alert Description: NFS export rule mismatch on Pure FlashBlade.

Reason: The alert is triggered when an unexpected file system export rule is detected, such as the export rule being manually modified.

Action: No action required

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 269

Platform: PXG1

CE00301147 RemoteStorageRegistrationInvalid

Alert Description: The remote storage registration information is invalid.

Reason: The registration information for remote storage (management IP address or API token) stored in the cluster is invalid.

Action: Verify and update the remote storage registration information.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.418

Platform: PXG1/PXG2

CE00301132 RemoteStorageConfigureFailed

Alert Description: Failed to configure remote storage.

Reason: This alert is triggered when the allocated capacity is less than the minimum required capacity.

Action: Increase the allocated capacity using the Cohesity UI.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.352

Platform: All

Power Supply Alerts

CE02013010 VoltOutOfHighRange

Alert Description: The voltage is above the operating threshold.

Reason: This alert is triggered when the current voltage exceeds the IPMI High voltage threshold setting.

Action: Contact Cohesity Support to investigate this issue and determine if any action should be taken.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.15

Platform: C2000 Series and C4000 Series

CE02013011 VoltOutOfLowRange

Alert Description: The voltage is below the operating threshold.

Reason: This alert is triggered when the current voltage is lower than the IPMI Low voltage threshold setting.

Action: Contact Cohesity Support to investigate this issue and determine if any action should be taken.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.16

Platform: C2000 Series and C4000 Series

CE02013015 PowerSupplyRemoved

Alert Description: The Power Supply was disconnected or removed.

Reason: This alert is triggered when the removal of a Power Supply is detected.

Action: No action is required if you intentionally removed the Power Supply. If not intentional, check the Power Supply.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.20

Platform: C2000 Series and C4000 Series

CE02013016 PowerSupplyInserted

Alert Description: The Power Supply was inserted.

Reason: This alert is triggered when the addition of a Power Supply is detected.

Action: No action is required if you intentionally added a Power Supply.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.21

Platform: C2000 Series and C4000 Series

CE02013020 NodePoweredOff

Alert Description: The node was powered off.

Reason: This alert is triggered when the node is powered off.

Action: No action is required if the node was intentionally powered off. If the node was powered off for an unknown reason, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.99

Platform: C2000 Series and C4000 Series

CE02013026 PowerSupplyAcLost

Alert Description: The AC power supply is lost.

Reason: This alert is triggered when one of the power supplies fails or is pulled out of the chassis.

Action: If the power supply is pulled out of the chassis, push it back in. If the power supply has failed, contact Cohesity Support. For more information, see the Alert: CE00413026 ACPowerSupplyLost KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.27

Platform: C2000 Series and C4000 Series

CE02013034 PowerSupplyAcRestored

Alert Description: The AC power supply is restored.

Reason: This alert is triggered when the AC power supply is restored.

Description: There will be no functional impact if the A/C power cable connection is recovered. However, if the cable connection recovery is not made intentionally, ensure that the A/C power cable is connected to the firmware.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.267

Platform: All

CPU Alerts

CE02113028 CpuThrottled

Alert Description: The CPU throttled.

Reason: The alert is triggered when a node's CPU is throttled.

Description: The alert might be caused by lower than expected CPU frequency because of a possible BIOS issue. You can power cycle the node to recover the CPU throttle; if the CPU throttle issue is persistent or reoccurring, then the possible causes might be:

- Firmware issue.
- Thermal issues in the node.
- Power supply issues.

Action: Perform the following:

- Verify that the node is running the latest firmware version.
- Thermal issues in the node might cause the CPU to overheat. Check the node environment for issues such as whether the data center is too hot or issues with air circulation in the server rack.
- Power supply issue might cause the CPU to throttle to reduce the power consumption. Verify if there are power supply failures in the IPMI SEL logs and replace the insufficient power supply.

  If the issue still persists, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.32

Platform: All

CE02101042 LessCPUCore

Alert Description: The node does not meet the minimum CPU core count requirement.

Reason: This alert is triggered when the node does not meet the minimum CPU core count requirements. It is triggered every hour.

Action: Review the CPU available for VM and increase the vCPU.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.205

Platform: Virtual Edition

Chassis Alerts

CE01901010 NodeChassisChanged

Alert Description: A node is detected on a different chassis.

Reason: This alert is triggered when the Cohesity cluster detects that a node has been moved to a different chassis than the chassis registered in the Cohesity cluster configuration. This alert may be triggered whenever the cluster services start on a node, which is typically when a node first boots up.

Action: Verify if the node was deliberately moved from one chassis to another.

Related Error Code:

Error Codes	Description
`kChassisChangedError`	The alert with the error code is raised when the node in a Cohesity cluster is moved from one chassis to another.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.4

Platform: All

Memory Alerts

CE02213005 MemUncorrectableEcc

Alert Description: Memory Uncorrectable ECC is being reported.

Reason: This alert is triggered when uncorrectable memory is detected in an ECC DIMM.

Action: Cohesity Support to investigate if this is a genuine DIMM failure and, therefore the DIMM needs to be replaced.

Related Error Code:

Error Codes	Description
kMemUncorrectableEcc	The alert with this error code is raised when Memory Uncorrectable ECC is reported.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.10

Platform: C2000 Series and C4000 Series

CE02201043 SmallMemorySize

Alert Description: The node does not meet minimum memory size requirements.

Reason: The node does not meet the minimum memory size requirement. This alert is triggered daily.

Action: If the alert is triggered for a Virtual Edition, review the VM RAM and increase the RAM. If the alert is triggered for a Physical Edition, see the Alert: CE00201043 SmallMemorySize KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.170

Platform: All

CE02213006 MemCorrectableEcc

Alert Description: Memory Correctable ECC is being reported.

Reason: This alert is triggered when correctable memory is detected in an ECC DIMM.

Action: If a resolution is not found, see the ALERT: CE00413006 MEMCORRECTABLEECC KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.11

Platform: C2000 Series and C4000 Series

Temperature Alerts

CE02313025 ThermalTripTriggered

Alert Description: The CPU is overheating.

Reason: This alert is triggered when a temperature threshold has been reached.

Description: The node has detected a thermal issue and the detected temperature is out of range.

Action: Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.26

Platform: All

CE02313008 TempOutOfHighRange

Alert Description: Temperature is above the threshold.

Reason: This alert is triggered when the sensed temperature is higher than the IPMI High temperature setting.

Action: Check the temperature of the data center and correct if necessary.

> **Note:** If the temperature gets too high, the system can automatically shut down.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.13

Platform: All

CE02313009 TempOutOfLowRange

Alert Description: Temperature is below the operating threshold.

Reason: This alert is triggered when the sensed temperature is lower than the IPMI Low temperature setting.

Action: Cohesity Support to investigate this issue and determine if any action should be taken.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.14

Platform: All

FAN Alerts

CE02413031 FanRemoved

Alert Description: The fan is removed.

Reason: A fan has been removed from the chassis.

Action: No action is required if the fan was removed for maintenance reasons. Verify that the removal of the fan is intentional.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.33

Platform: C4000 Series Only

CE02413032 FanInserted

Alert Description: The fan is inserted.

Reason: A fan has been inserted in the chassis.

Action: No action is required if the fan was inserted for maintenance reasons.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.34

Platform: C4000 Series Only

CE02413033 FanSpeedOutOfLowRange

Alert Description: Fan Speed out of low range.

Reason: The fan speed has dropped below an acceptable threshold, indicating a fan has failed and requires maintenance.

Action: Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.35

Platform: C2000 Series and C4000 Series

NIC Alerts

CE02513022 LinkIsUp

Alert Description: Node Network Interface is up

Reason: This alert is triggered when a network interface changes state from down to up.

Action: No action unless you see a large number of these alerts.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.23

Platform: All

CE02513023 LinkIsDown

Alert Description: The network interface of the node in the Cohesity cluster is down.

Reason: This alert is triggered when a network interface changes state from up to down.

Action: No action is required unless you see a large number of these alerts. Check cabling and the switch settings if you see a large number of these alerts.

Related Error Code:

Error Codes	Description
kLinkDown	The alert with the error code is raised when the link is down.

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.22

Platform: All

CE02501057 ProductModelInterfacesInconsistency

Alert Description: Inconsistent network interfaces found on the nodes in the same product model.

Reason: This alert is triggered when there are inconsistent network interfaces on the nodes in the same product model.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.196

Platform: All

CE02501058 ClusterInterfacesInconsistency

Alert Description: Inconsistent network interfaces found in the heterogeneous cluster.

Reason: This alert is triggered when there are inconsistent network interfaces on different product models in a heterogeneous cluster.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.197

Platform: All

Firmware Alerts

CE02601022 FirmwareIncompatible

Alert Description: Node firmware is incompatible.

Reason: This alert is triggered when the firmware detected on the node is not supported by Cohesity.

Action: Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.185

Platform: All

CE02601045 BIOS/BMCVersionMismatch

Reason: This alert is triggered daily.

Description: Nodes in the cluster have incompatible BIOS/BMC versions.

Action: Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.208

Platform: All

CE02601081 BMCModuleNotAvailable

Alert Description: The BMC module is not available on the node.

Reason: This alert is triggered when the BMC module is not available. The alert can be caused due to a BMC watchdog error, or the version might not be updated.

Action: Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.224

Platform: Physical

Software Alerts

Download a spreadsheet containing **Software** alerts.

Node Health Alerts

Operating System Alerts

Data Path Alerts

Metadata Alerts

Indexing Alerts

Apps Infrastructure Alerts

App MarketPlace Alerts

Node Health Alerts

CE00401033 ServiceDown

Alert Description: A service on the node is down.

Reason: This alert is triggered whenever a critical Cohesity service is down. It is triggered every hour until resolved. The alert can be caused when a critical Cohesity service is down for more than 6 hours.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.163

Platform: All

CE00413007 WatchdogTriggered

Alert Description: The last reboot of the node was caused by the IPMI Watchdog.

Reason: The alert is triggered when the Watchdog module reboots the node. This mechanism is used to detect whether the system software has become unresponsive. When the watchdog timer (5 min) expires, the watchdog module forces the node to reboot to recover the unresponsive software.

Action: Contact Cohesity Support to investigate this issue and determine if any action should be taken.

Related Error Code:

Error Codes	Description
kWatchdogTriggered	The alert with this error code is raised when the IPMI Watchdog causes the last reboot of the node.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.12

Platform: C2000 Series and C4000 Series

CE00413007 WatchdogTriggered

Alert Description: The last reboot of the node was caused by the IPMI Watchdog.

Reason: The alert is triggered when the Watchdog module reboots the node. This mechanism is used to detect whether the system software has become unresponsive. When the watchdog timer (5 min) expires, the watchdog module forces the node to reboot to recover the unresponsive software.

Action: Contact Cohesity Support to investigate this issue and determine if any action should be taken.

Related Error Code:

Error Codes	Description
kWatchdogTriggered	The alert with this error code is raised when the IPMI Watchdog causes the last reboot of the node.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.12

Platform: C2000 Series and C4000 Series

CE00413021 NodeRebooted

Alert Description: The Node rebooted.

Reason: This alert is triggered when a node is rebooted.

Action: No action if you intentionally rebooted a node. If not intentional, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.25

Platform: C2000 Series and C4000 Series

CE00402301 TaskProgressHung

Alert Description: The root task has some hung sub-tasks

Reason: This alert is triggered when the Cohesity cluster detects that Pulse Task is hung and not making any progress.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.97

Platform: All

CE00401008 TimeService

Alert Description: The node cannot synchronize time.

Reason: This alert is triggered when a node is unable to synchronize time from a NTP server to the Cohesity cluster time service. This alert can be triggered by one of the following conditions:

- The configured NTP server is invalid or down.
- The configured NTP server only implements a subset of the NTP protocol.

> **Note:** There are known time synchronization issues with Windows NTP servers and `time.nist.gov`.

Action: Verify if the configured NTP server is reachable.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.7

Platform: All

CE00401008 TimeService

Alert Description: The node cannot synchronize time.

Reason: This alert is triggered when a node is unable to synchronize time from a NTP server to the Cohesity cluster time service. This alert can be triggered by one of the following conditions:

- The configured NTP server is invalid or down.
- The configured NTP server only implements a subset of the NTP protocol.

> **Note:** There are known time synchronization issues with Windows NTP servers and `time.nist.gov`.

Action: Verify if the configured NTP server is reachable.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.7

Platform: All

CE00401068 ClusterBringupError

Alert Description: An error is found during cluster start-up.

Reason: This alert is triggered when an error occurs during cluster start-up.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.75

Platform: All

CE00401070 ScriptRunError

Alert Description: The health-check script has errors.

Reason: This alert is triggered when a health-check script is running and detects an error.

Action: Check if there are any syntax or logic errors in the scripts. If not, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 1800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.65

Platform: All

CE00411001 FrequentProcessRestarts

Alert Description: A process has restarted frequently on the node.

Reason: This alert is triggered when a Cohesity cluster process is restarted. A process can restart for one or more of the following reasons:

- An unexpected condition (assert) was encountered.
- An operation within the process is hung and this triggered the failfast mechanism to restart the process.

Action: Refer to the following KB articles:

- Understanding Alert CE00411001 and frequent service restarts on a Cohesity Cluster
- Alert: CE00411001 Process stats_exec is restarting frequently
- Alert: CE00411001 Process iris_exec is restarting frequently
- Alert: CE00411001 Process rtclient_exec is restarting frequently
- Alert: CE00411001 Process yoda_exec is restarting frequently
- Alert: CE00411001 Process librarian_exec is restarting frequently
- Alert: CE00411001 Process nexus_exec is restarting frequently
- Alert: CE00411001 Process magneto_exec is restarting frequently
- Alert: CE00411001 Process apollo_exec is restarting frequently
- Alert: CE00411001 Process gandalf_exec is restarting frequently
- Alert: CE00411001 Process nexus_proxy_exec is restarting frequently

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.138

Platform: All

CE00411003 OutOfMemoryRestart

Alert Description: Process restarted due to out of memory.

Reason: This alert is triggered when any process gets killed due to out of memory.

Action: Contact Cohesity Support.

Severity: Warning

Dedup Interval: 43200 seconds

SNMP OID: Not Applicable

Platform: All

CE00401089 AppNetworkNodeUnhealthy

Alert Description: App Network Node status is unhealthy.

Reason: This alert is triggered when the node-specific axon error checks fail. The node might have network virtualization issues.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.235

Platform: All

CE00411002 UnresponsiveSubProcessFailure

Alert Description: The process is unresponsive.

Reason: This alert is triggered when the parent process detects that the child process is unhealthy and unresponsive.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.151

Platform: All

CE00401108 DuplicateIpDetected

Alert Description: A duplicate IP address is detected in the node.

Reason: This alert is triggered when duplicate node IP addresses are detected in the Cohesity cluster. Duplicate node IP addresses in the Cohesity cluster cause network issues.

Action: Verify if the same IP address is assigned to multiple nodes in the Cohesity cluster or malicious actors are trying to steal the IP address on the nodes in the Cohesity cluster.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.266

Platform: All

CE00401107 NodeTcpOrphanConnectionsExceeded

Alert Description: The number of TCP Orphan connections exceeds the limit.

Reason: This alert is triggered when TCP FIN WAIT connections exceed the maximum threshold (100). It is typically triggered if there are unclosed TCP connections due to the following reasons:

- The TCP connections are closed due to a network issue.
- The end-user removes the network connection without closing the application.
- The Cohesity applications are closing the connections without waiting for the ACK from the server. If there are many such open connections, the new connections are denied. The Cohesity cluster detects this condition and proactively closes the connections.

Action: No action is required if you are using a physical cluster. If you are using Cohesity Cohesity Cloud Edition for AWS and the network usage exceeds the AWS allowed limit (for M4X limit is 2Gbps/100kpkts/sec), then Cohesity recommends upgrading the EC2 instance to M5N.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.262

Platform: All

CE00401109 RemoteClusterNotReachable

Alert Description: The remote cluster is not reachable.

Reason: This alert is triggered when the Cohesity cluster cannot connect to the remote UI or REST API port.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.268

Platform: All

CE00401117 ServiceGflagsChanged

Alert Description: Service Gflags are changed in the cluster.

Reason: The alert is triggered when you have applied a gflag in the Cohesity cluster.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.302

Platform: All

CE00401113 UnexpectedIpAddress

Alert Description: Unexpected IP Address in an interface.

Reason: This alert is triggered when an unexpected IPv6 address is present in an interface. The alert might be caused when SLAAC assigns an IP address.

Action: Restart the network or contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.289

Platform: All

CE00401116 HostShellAccessExpiry

Alert Description: Host shell access is about to expire.

Reason: This alert is triggered to notify that the host shell access will expire in an hour.

Action: Contact Cohesity Support if you want to extend the duration to access the host shell.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.297

Platform: All

CE00401148 NodeTcpOpenConnectionsExceeded

Alert Description: The number of TCP open connections is exceeding the limit.

Reason: This alert is triggered when the number of TCP open connections surpasses 90% of the total connection count, as defined by the configured GFLAG threshold.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.1.441

Platform: All

CE00401148 NodeTcpOpenConnectionsExceeded

Alert Description: The number of TCP open connections is exceeding the limit.

Reason: This alert is triggered when the number of TCP open connections surpasses 85% of the total connection count, as defined by the configured GFLAG threshold.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.1.441

Platform: All

CE00401149 NetworkInterfaceFlaky

Alert Description: The number of network interface flaps is exceeding the limit.

Reason: This alert is triggered when a network interface exhibits "flaky" behavior. A physical interface is considered flaky when it repeatedly transitions between an up and down state within a specific time period, exceeding a defined threshold.
Definition of Flaky Behavior
The network interface transitions up or down more than n times within a t-second window.
These thresholds are defined by default as, n = 10 transitions and t = 60 seconds. Both values are configurable using a GFLAG.
Example of Flaky Behavior
If the interface undergoes transitions like this: [U D U D U D U D U D U]. This results in 11 counts (each up and down counts individually) within 60 seconds, which exceeds the default threshold of 10. Consequently, the alert is raised to signal potential instability or hardware issues.

Action: If you encounter this alert, follow these steps to identify and address the issue:
- **Inspect the Physical Link**: Verify the physical connection of the network interface.

- **Examine the Network Interface Card (NIC)**: Confirm that the NIC is functioning correctly.
- **Check the Peer Router or Switch**: Investigate the device connected to the interface (peer router or switch).
If the issue persists after the initial checks, contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.444

Platform: All

CE00422001 NetworkRPCTimeoutThresholdExceeded

Alert Description: I/O Operations Services has detected that RPC calls to the node are timing out and exceeding the allowed threshold.

Reason: This alert is triggered when unhealthy or degraded network conditions are detected on the node, leading to RPC timeouts.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.1.459

Platform: All

CE00422002 NetworkRPCTimeoutRecovered

Alert Description: I/O Operations Services has detected that RPC timeouts to the node have returned to normal and are now within the allowed threshold.

Action: No action required

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.1.460

Platform: All

CE00422003 ComponentStateFlagged

Alert Description: A node is flagged with a potential network issue based on elevated network errors relative to normal behavior.

Reason: This alert is triggered when the network error signals for a node are elevated relative to healthy peers.

Action: No action required.

Related Error Code: Not applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.509

Platform: All

CE00422004 ComponentStateRecovering

Alert Description: A degraded network behavior has been detected on a node.

Reason: This alert is triggered when the node's measured network performance is significantly worse than that of healthy peers.

Action: No immediate action is required. If the node remains degraded, contact Cohesity Support.

Related Error Code: Not applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.510

Platform: All

CE00422005 ComponentStateDegraded

Alert Description: The network performance on a node remains degraded despite automated recovery attempts.

Reason: This alert is triggered when automated recovery does not restore healthy network behavior for the node.

Action: Inspect network configuration and upstream connectivity for the affected node. If you need help interpreting the alert context, contact Cohesity Support.

Related Error Code: Not applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.511

Platform: All

CE00422006 ComponentStateRecovered

Alert Description: Network behavior on a node has returned to a healthy state.

Reason: This alert is triggered when the node's network path returns to normal.

Action: No action required.

Related Error Code: Not applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.512

Platform: Platform: All

Operating System Alerts

CE02713004 IpmiEvent

Reason: This alert is triggered when an unusual IPMI event is detected such as the Power supply is no longer redundant.

Action: Contact Cohesity Support to investigate the problem and determine if any action should be taken.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: C2000 Series and C4000 Series

CE02713024 IpmiSelCleared

Alert Description: The node IPMI SEL log file is cleared.

Reason: This alert is triggered when the size of the IPMI SEL Log exceeds 80% of its maximum size and the IPMI SEL Log file is cleared.

Action: No action is required. This alert is logged for auditing.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.24

Platform: C2000 Series and C4000 Series

CE02701004 KernelPanic

Alert Description: The node crashed from a kernel panic.

Reason: This alert is triggered when a kernel panic occurs on the node. This alert is triggered once per kernel panic.

Action: See the Alert: CE00401004 node crashed from a kernel panic KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.103

Platform: All

CE02701054 SysSrvcErrors

Alert Description: The node bond mode has been migrated.

Reason: This alert is triggered after the Cohesity cluster is upgraded to 6.2.0. This alert is triggered if error logs are found for any of the Cohesity core services on the cluster nodes.

Action: See the ALERT: CE00201054 FOUND ERRORS IN SYSTEM SERVICES LOG KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.193

Platform: All

CE02701067 SecurityPolicyDenial

Alert Description: Security policy denial found on the node.

Reason: This alert is triggered when a security policy denies access to an app on a node.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.84

Platform: All

CE02701160 MicrosoftSmtpDeprecated

Alert Description: The cluster email configuration is using a Microsoft SMTP server that relies on basic authentication with username and password. Microsoft is deprecating this method, and email delivery through this server may stop working after the deprecation.

Reason: This alert is triggered when the cluster email configuration uses a Microsoft SMTP server that matches Microsoft infrastructure (for example, hostnames ending in .office365.com, .outlook.com, .live.com, or .hotmail.com) and relies on basic authentication. The system raises this warning so you can migrate before basic authentication is no longer supported.

Action: Migrate cluster email delivery from Microsoft SMTP with basic authentication to Microsoft Graph API with OAuth 2.0. Update the cluster email configuration accordingly. If you use multi-factor authentication (MFA) or email one-time password (OTP) with the current Microsoft SMTP server, update those settings after migration. On supported releases, new Microsoft SMTP configurations may be blocked; use Microsoft Graph API for email instead. If you no longer use Microsoft SMTP, disable SMTP or configure a non-Microsoft SMTP server so the cluster no longer matches this condition. Contact Contact Cohesity Support if you need help planning the migration.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.492

Platform: All

CE02701161 MicrosoftGraphAccessTokenFetchFailed

Alert Description: The cluster failed to fetch an OAuth 2.0 access token for the Microsoft Graph API used for email delivery.

Reason: This alert is triggered when the cluster cannot obtain an OAuth 2.0 access token from Azure Entra Active Directory for the configured Microsoft Graph API application. The failure prevents the cluster from sending email notifications through Microsoft Graph.

Action: Verify that the client ID, client secret, and tenant ID in the cluster email configuration are correct. Ensure the client secret has not expired and that the Azure Entra app registration has the required API permissions for Microsoft Graph. After correcting the configuration, the cluster will automatically retry fetching the access token. Contact Cohesity Support if the issue persists.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.493

Platform: All

CE02701162 MicrosoftGraphClientSecretExpiry

Alert Description: The Microsoft Graph API client secret configured for email delivery is nearing its expiration date or has already expired.

Reason: This alert is triggered when the client secret for the configured Microsoft Graph API application is going to expire within the warning threshold. An expired client secret causes email delivery failures.

Action: Generate a new client secret for the Azure Entra app registration, then update the cluster email configuration with the new secret value and its expiration date. Update the configuration before the secret expires to prevent email delivery failures. Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.494

Platform: All

CE02701162 MicrosoftGraphClientSecretExpiry

Alert Description: The Microsoft Graph API client secret configured for email delivery is nearing its expiration date or has already expired.

Reason: This alert is triggered when the client secret for the configured Microsoft Graph API application expires on that day or has already expired. An expired client secret causes email delivery failures.

Action: Generate a new client secret for the Azure Entra app registration, then update the cluster email configuration with the new secret value and its expiration date. Update the configuration before the secret expires to prevent email delivery failures. Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.494

Platform: All

Data Path Alerts

CE02806001 MetadataInconsistencyAlert

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: Internal error detected on the cluster.

Reason: This alert is triggered when the Cohesity cluster detects an internal issue during scan. Periodically the Cohesity cluster runs a background scan looking for potential internal issues.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.100

Platform: All

CE02806002 InternalSchedulingWarning

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: Internal scheduling issue.

Reason: Periodically the Cohesity cluster runs a background scan looking for potential internal issues. This alert is triggered if the background scan cannot complete on schedule. A scan may not complete on schedule for many reasons. The issue may be benign and fix itself.

Action: Contact Cohesity Support if the issue persists after a few days.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.100

Platform: All

CE02806003 CloudTierWarning

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: Cloud tier warning.

Reason: This alert is triggered when the disk space usage of the Cohesity cluster is over 80% and the Storage Domains do not have sufficient cold data to spill the Cloud Tier to reduce the disk space usage. This alert is triggered once per day.

Action: Contact Cohesity Support or use the Cohesity CLI to change the cold data threshold.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.100

Platform: All

CE02806012 AlgorithmBacklogAlert

Alert Description: A background system algorithm is backlogged.

Reason: This alert is triggered when the background system algorithm has been backlogged for too long. A background system algorithm is backlogged if the algorithm has not run for 2 times the average run time of that algorithm.

Action: If the system activity schedule is not set or enough resources are available but still the alert is raised, then contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE02806013 ChunkFileBadState

Alert Description: The chunk file in the Storage domain is in a bad state for a long time.

Reason: This alert is triggered when a chunk file with a bad replica is not fixed for 10 days. The alert is caused when a disk fails.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.76

Platform: All

CE02806025 GCStalledAlert

Alert Description: The garbage collection has stalled on the cluster.

Reason: This alert is triggered when garbage collection gets stalled on the cluster.

Action: Contact Cohesity Support. If corrective action is not taken, the cluster will run out of space causing an outage.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.360

Platform: All

CE02809003 ActiveDirectoryNotReachable

Alert Description: The active directory domain is not reachable.

Reason: This alert is triggered when the registered Active Directory server is not reachable. The alert can be caused by one of the following conditions:

- The Active Directory server is down.
- Incorrect network settings on the Cohesity cluster.
- Incorrect DNS settings on the Cohesity cluster.

Action: Check the following:

- If the Active Directory server is online.
- The Cohesity cluster can establish a network connection to the Active Directory server (ping test is successful).
- The network settings on the Cohesity cluster.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.102

Platform: All

CE02809014 ChunkFileInconsistency

Alert Description: Potential data access issue for chunk file.

Reason: This alert is triggered once a day when the Cohesity cluster detects a inconsistency in a data segment stored on the Cohesity cluster.

Action: Contact Cohesity Support to resolve this data segment inconsistency.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.90

Platform: All

CE02809020 ViewLatencyHigh

Alert Description: The view read/write latency exceeds the specified threshold.

Reason: This alert can be caused due to hardware contention (not enough Nodes are provisioned), interference from other workloads or background operations, or the access pattern of the workload.

Action: Consider setting the Alert threshold and taking the appropriate action for your environment.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.177

Platform: All

CE02809023 AdDomainControllersNotReachable

Alert Description: The active directory domain controller is not reachable.

Reason: This alert is triggered when one or more of the AD domain controllers is not reachable, when the cluster is registered to an AD domain

Action: Check the status of the Domain Controllers in their AD environment.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.67

Platform: All

CE02809024 AdPreferredDomainControllerNotReachable

Alert Description: The preferred domain controller is not reachable.

Reason: This alert is triggered when the preferred Domain Controller is not reachable, when the cluster is registered to AD and a preferred DC is set for the AD domain.

Action: Check the status of the Preferred Domain Controllers in their AD environment.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.175

Platform: All

CE02820009 UnregisterExternalTargetFailed

Alert Description: Unregistration for external target failed.

Reason: This alert is triggered when an error occurs while deleting an external target.

Action: Check the error details on the external target page in the Cohesity Dashboard.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE02820010 ExternalTargetNotReachable

Alert Description: External target not reachable from node.

Reason: This alert is triggered when the vault is not reachable from nodes for a certain period of time. One or more nodes cannot reach the external target.

Action: Check the connection status of the external targets, and also verify the health of the individual cluster nodes.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE02809002 DiskWithTooManyDeltaRecords

Alert Description: The disk has too many internal metadata records.

Reason: This alert is triggered if any disks in the Cohesity cluster have more than 500K of delta records in its WAL (Write-Ahead Logging). The clean-up of the delta records is delayed due to the following conditions:

- The NFS/SMB server process (called the Bridge) on the Cohesity cluster is very busy.
- The disks are very busy.

Action: See the Alert: CE00409002 DiskWithTooManyDeltaRecords KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 21600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.150

Platform: All

CE02806014 ManualBackgroundSystemThrottlingApplied

Alert Description: Manual throttling of background system algorithms applied.

Reason: This alert is triggered when manual throttling of background system algorithms is applied. This alert can be caused if manual throttling is applied to background system algorithms, and the throttling schedule is skipped.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE02809032 PotentialMinorClusterError

Alert Description: The cluster might have an internal error that requires action by Cohesity Support.

Reason: This alert is triggered when a minor inconsistent chunk of data is found in the Cohesity cluster. The alert can be caused due to memory corruption.

Action: No action required.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.178

Platform: All

CE02809021 BlobJournalInconsistency

Alert Description: The blob journal in the view box is inconsistent.

Reason: This alert is triggered when the blob journal for a storage domain is inconsistent. The alert can be caused due to memory corruption.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.176

Platform: All

CE02809033 PotentialClusterError

Alert Description: The cluster might have an error requiring Cohesity Support action.

Reason: This alert is triggered when a minor inconsistent chunk of data is found in the Cohesity cluster.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.79

Platform: All

CE02809047 ViewSlowProtocolOp

Alert Description: Slow operation detected from the view.

Reason: This alert is triggered if there is performance degradation due to the duration of the internal operations associated with any NFS or SMB operations taking more than 5 minutes. The alert is raised every 10 minutes after the initial occurrence of the incident.

Action: All the necessary logs are stored under the logs folder. Contact Cohesity Support to download the required files for analysis.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.274

Platform: All

CE02820011 NASVaultIsFull

Reason: This alert is triggered when the space usage of the external targets of type NAS is at 95%.

Action: Increase the space on the external target or modify the retention policy for archives.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: N/A

Platform: All

CE02820012 VaultConfigValidationFailed

Reason: The alert is triggered when a lifecycle policy is associated with the vault or due to permission issues.

Description: The alert is caused when the configuration validation fails for the vault used as the external target.

Action:

- If a lifecycle policy is associated with the vault, then remove the policy.
- If there are permission issues, then fix the issue.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: N/A

Platform: All

CE02820012 VaultConfigValidationFailed

Reason: The alert is triggered when a lifecycle policy is associated with the vault or due to permission issues.

Description: The alert is caused when the configuration validation fails for the vault used as the external target.

Action:

- If a lifecycle policy is associated with the vault, then remove the policy.
- If there are permission issues, then fix the issue.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: N/A

Platform: All

CE02806015 AlgorithmFrequentlyTimeOutAlert

Reason: The alert is triggered when Healer and Space reclamation services algorithms do not start on schedule or time out repeatedly.

> **Note:** Once an alert is triggered, a new alert gets generated whenever an algorithm finishes.

Description: The alert is caused when background algorithms are timing out repeatedly or if any background algorithm has not started on the regular schedule.

Action: Contact Cohesity Support.

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

CE02809054 ObjectStorePortalInitializationFailure

Alert Description: Bridge S3 HTTP Server failed to initialize.

Reason: The alert is triggered when the ObjectStore (S3) Portal fails to start. Once the alert is raised, I/O Operations Services continues to work without any crash, but any S3 requests are denied. If I/O Operations Services are restarted for some reason, the alert is triggered if the error still exists. The alert can be caused by one of the following conditions:

- Invalid TLS certificate.
- Invalid network settings.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 60 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.326

Platform: All

CE02809053 CloudDataOpFailure

Alert Description: Cloud operation failures detected on external targets.

Reason: This alert is triggered when 50% or more I/O operations on the cloud storage tier fail in the last one hour. The alert might be caused due to degraded cloud connectivity or unavailability.

Action: Verify the cloud connectivity and availability.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.325

Platform: All

CE02806016 RepeatedlyFailingShardAlert

Alert Description: Background tasks have hit a slowness issue because Apollo shards are failing more often than expected.

Reason: This alert is triggered when a Healer and Space Reclamation Services algorithm frequently time out and is behind schedule for a long time.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.340

Platform: All

CE02806017 FragmentActionsRepeatedlyFailingAlert

Alert Description: Many background actions in Apollo are hitting unexpected errors.

Reason: This alert is triggered when many background actions in Healer and Space Reclamation Services encounter unexpected errors.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: TBD

Platform: All

CE02806018 BacklogDuringBlackOut

Alert Description: Apollo is in backlog state for more than a threshold number of days during the blackout window auto-scheduling.

Reason: This alert is triggered when a backlog is detected during the black out window auto scheduling, and it is present for more than a threshold number of days. This alert can be caused when some background services are not functioning and a backlog action item is found for the service as a whole.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 345600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.341

Platform: All

CE02806020 TimeConsumingShardAlert

Alert Description: Apollo shard runtime is greater than the average value by a threshold.

Reason: This alert is triggered when a shard of the Healer and Space Reclamation Services algorithm runtime exceeds the average value of the remaining shards corresponding to the same fragment by a threshold.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.342

Platform: All

CE02806019 HighClusterUtilization

Alert Description: Cluster utilization is high.

Reason: This alert is triggered when the cluster utilization is high (>=85%).

Action: Contact Cohesity Support or refer to the KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 345600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.321

Platform: All

CE02806021 HighNonDedupToDedupPercent

Alert Description: High non-dedup to dedup percent observed in inline dedup Storage domain.

Reason: This alert is triggered when the non-dedup to dedup percent is observed above a threshold for a storage domain.

Action: Contact Cohesity Support

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 259200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.343

Platform: All

CE02809052 ApolloIssuedActionVersionMismatch

Alert Description: Apollo is issuing actions with incorrect API versions.

Reason: This alert is triggered if the cluster upgrade is not in progress and there is a version mismatch between the bridge and the MapReduce and Space Reclamation service executions.

Action: You can ignore occasional alerts. If the alert is persistent, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: N/A

Platform: All

CE02809055 ViewCloneMetadataBloat

Alert Description: There are many unflushed records in the metadata journal for the view during cloning.

Reason: This alert is triggered when the number of unflushed records in the metadata journal at the time of cloning exceeds the threshold. The alert can be caused when the unflushed records in the metadata journal at the time of cloning is high.

Action: Contact Cohesity Support if the alert is persistent.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.317

Platform: All

CE02809056 ChunkFileContainerInconsistency

Alert Description: Chunk File Container inconsistency, cluster internal error that requires action by Cohesity Support.

Reason: This alert is triggered when too many stripes for chunk file containers are corrupt or not existent. The alert can be caused when all the stripes for the chunk file container are inconsistent.

Action: Contact Cohesity Support if the alert is persistent.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.318

Platform: All

CE02820013 ArchiveObjectGCFailedAlert

Alert Description: Archive object garbage collection in external target failed.

Reason: This alert is triggered when the icebox tries to GC archive objects in the external target but fails.

Action: Assess if the objects are protected by data lock on purpose. If yes, ignore it. Check if the required permissions are provided. If not, contact Cohesity Support

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

CE02809046 EntityWithTooManyOpenHandles

Alert Description: There are too many open handles for the same entity.

Reason: The SMB protocol triggers the alert when the open handle (file descriptor) number for a specific file exceeds the configured threshold. By default, this alert is disabled. The alert can be caused when some rogue applications create file handles and

never close these handles. These open file handles cause the metadata bloat and cause the File Server to be in the crash loop.

Action: You can list the open file handles for the file for which the File server has raised the alert with the `iris_cli` command: `iris_cli view ls-sfo`.
The command provides a list that details the client IP address, user name, and so on that created an open handle on the file. In addition, this list provides the application details that create these many open handles on the file. You must contact the application vendor to check whether the application is leaking these open handles.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.273

Platform: All

CE02806024 ECConversionBacklogged

Alert Description: Old un-encoded chunk file observed in the Storage domain. Erasure coding algorithm is backlogged.

Reason: This alert is triggered when the wide stripe EC container enabled view box has an erasure coding backlog.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

CE02809062 HeavyFileOperationContention

Alert Description: Operation on the same file has been too heavy and rejected.

Reason: This alert is triggered when extreme contention or overload is detected in Cohesity's SMB service stack. These overload conditions result in users experiencing errors accessing files. This alert is caused when hundreds of users access the same file simultaneously, or many infected client machines try to open the same file repeatedly.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.356

Platform: All

CE02809048 BucketMetadataModified

Alert Description: Metadata of the view is modified.

Reason: This alert is triggered once per minute per bucket (view) whenever a S3/Swift view metadata (except ACL) is modified.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 60 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.298

Platform: All

CE02809049 BucketACLUpdated

Alert Description: Access Control List of S3 view is updated

Reason: This alert is triggered once per minute per bucket (view) whenever a S3/Swift bucket ACL is updated.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 60 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.299

Platform: All

CE02809050 BucketCreated

Alert Description: New view is created.

Reason: This alert is triggered on the creation of every S3 or Swift view.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 60 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.300

Platform: All

CE02809051 BucketDeleted

Alert Description: The view is deleted.

Reason: This alert is triggered on the deletion of every S3 or Swift view.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 60 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.301

Platform: All

CE02806023 kBrickMetadataInconsistentState

Alert Description: Chunk mapping for chunk in brick with leaf node is in inconsistent state.

Reason: This alert is triggered for every Healer run, if Healer detects some brick with metadata inconsistency. Some files in the cluster are in an inconsistent state, so read/write to that file may cause issues.

Action: Contact Cohesity Support and file a field issue with the MapReduce and Space Reclamation service team.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.349

Platform: All

CE02806028 ImbalanceThresholdExceededAlert

Alert Description: Imbalance threshold exceeded for chunkfiles. This can cause increase in the cluster usage.

Reason: There is sub-optimal erasure coding overhead of greater than 100GB in an erasure coding enabled view box. When the minimum data stripe size and maximum data stripe size of the erasure coded chunk files differ a lot, the erasure coding involves space overhead when compared to an ideal case where all data stripes are of the same size. Such EC imbalance space overheads are automatically reclaimed in the background by the MapReduce and Space Reclamation service. This alert is caused when the overheads are not reclaimed in a timely manner.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.387

Platform: All

CE02806029 MorphedGarbageThresholdExceededAlert

Alert Description: Morphed garbage threshold exceeded for chunkfiles in cluster.

Reason: Some view boxes contain more than 100GB of garbage that are not reclaimed from the disks. This accounts for greater than 20% of the view box utilization. This alert is caused when the garbage is not reclaimed in a timely manner. Cohesity software optimally reclaims garbage from the disks by holding off until a threshold value is reached. After that, the MapReduce and Space Reclamation service reclaims the garbage periodically.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.385

Platform: All

CE02806026 OverwriteThrottlingScheduleAlert

Alert Description: A backlog in the background clean-up task was observed, once cluster utilization becomes high, its priority will be increased which could create delays in the SmartFiles workloads.

Reason: If the usage of the Cohesity cluster is high and a backlog in the background clean-up task is observed, then the throttling schedule is overridden.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.370

Platform: All

CE02809072 EnableObjectServicesViewAllowlistFeature

Alert Description: Unauthorized access to S3 views using the S3 protocol has been detected due to the S3 allowlist feature being disabled.

Reason: This alert is triggered when S3 views are accessed from an IP address that is not included in the allowlist.

Action: Contact Cohesity Support to enable the allowlist feature for S3 views. For more information, see How To Enable SmartFiles S3 View Allowlisting KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 432000 seconds

SNMP OID: Not Applicable

Platform: All

CE02806033 HighUniqueUntrackedYodaPhysicalUsageAlert

Alert Description: A high amount of untracked physical usage by Indexing Services has been detected on the cluster.

Reason: This alert is triggered when the Indexing Service's untracked physical usage surpasses a given threshold on the cluster.

Action: Contact Cohesity Support.

Related Error Code: NA

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.422

Platform: All

CE02809061 AVHighRequestQueueSize

Alert Description: The antivirus scan queue is too long

Reason: This alert is triggered when the request queue size for an antivirus group exceeds a specified threshold. The alert indicates that the antivirus server is not well provisioned.

Action: Increase the number of antivirus servers.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.353

Platform: All

CE02809060 AVHighResponseTime

Alert Description: Antivirus scan slowness is detected.

Reason: This alert is triggered when the response time for an antivirus group exceeds a specified threshold. The alert indicates that the antivirus scans are taking longer than the configured threshold, indicating potential performance issues with the antivirus service.

Action: Increase the number of antivirus servers.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.354

Platform: All

CE02809071 ExcessiveSMBAccessAttempt

Alert Description: Excessing SMB requests from the client to the node.

Reason: The alert is triggered when an excessive number of SMB mount requests are made within a short period of time from the same client by the same user.

Action: Investigate the client generating the excessive SMB mount requests to identify and resolve the underlying issue.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.423

Platform: All

CE02820014 UpgradeExternalTarget

Alert Description: Upgrade external target to Incremental forever format.

Reason: This alert is triggered when existing external targets use the Incremental with Periodic Full (v1) archival format.

Action: Upgrade the external target to the Incremental Forever format.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.431

Platform: All

CE02820015 SwitchLegacyGlacierToGlacierTier

Alert Description: The external target is currently using legacy Glacier. Switch to S3 Glacier tiers.

Reason: This alert is triggered when existing external targets use Legacy Glacier. Cohesity recommends switching to the S3 Glacier tier for better compatibility and performance.

Action: Migrate to the S3 Glacier tier.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.432

Platform: All

CE02810045 KubernetesUnquiesceFailure

Alert Description: Unquiesce operation failed for the backup job.

Reason: This alert is triggered when an unquiescing operation fails for a specific Kubernetes protection job.

Action: In the Cohesity Dashboard, drill down into the protection run to view detailed information about the failure, including the affected namespace. Ensure that the applications within the namespace are in a healthy state to resolve the issue.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.440

Platform: All

CE02810042 MissedBackupRuns

Alert Description: Failed to schedule backup.

Reason: This alert is triggered when backup runs fail to execute according to the configured backup schedule.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.11

Platform: All

CE02810043 OutOfComplianceBackupRuns

Alert Description: Protection runs are out of compliance.

Reason: This alert is triggered when backup runs do not complete successfully as per the defined backup schedule.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.421

Platform: All

CE02810048 ProgressMonitorFailure

Alert Description: Progress monitor is inactive.

Reason: This alert is triggered when the Cohesity cluster fails to create the Progress Monitor itself—not just because it stopped reporting.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

Severity: 1.3.6.1.4.1.47421.0.456

Platform: All

CE02809075 ObjectServiceLocalIngressPaused

Alert Description: Data ingress for the S3 session is paused.

Reason: This alert is triggered when data ingress for an S3 session remains paused and does not resume within the specified time window (30 minutes). Internal flow-control mechanisms may pause data ingress if memory usage exceeds a threshold, and it resumes automatically once memory is freed.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.474

Platform: All

CE02809076 ObjectServiceClientTimeouts

Alert Description: S3 client disconnected while the server was still processing the request.

Reason: This alert is triggered when the S3 client times out or closes the connection before the server completes processing. This may occur due to client-side configuration issues or network interruptions.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.476

Platform: All

CE02809073 ObjectServiceServerTimeouts

Alert Description: Server timeout encountered during an S3 operation.

Reason: This alert is triggered when an S3 server takes too long to respond to a request, resulting in a timeout.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.464

Platform: All

CE02809074 ObjectServiceTooManyRequests

Alert Description: S3 requests exceeded threshold.

Reason: This alert is triggered when the number of active S3 requests exceeds the maximum outstanding threshold (512). Once the threshold is reached, additional requests are throttled.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.473

Platform: All

CE02820016 LambdaValidationFailed

Alert Description: Validation of the Lambda function for the external target failed.

Reason: This alert is triggered when the garbage collection to avoid egress charges by the hot tier targets is affected because of failure in the lambda functions.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.463

Platform: External targets that require lambda functions

CE02806034 BackgroundOperationTargetAlert

Alert Description: An external target required for background operations is not configured or not eligible for use.

Reason: The external target does not meet the requirements for system background services.

Action: Verify the target configuration and refer to the Knowledge Base for supported configurations.

Related Error Code: Not Applicable.

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.499

Platform: Cohesity Cloud Edition

CE02809078 ObjectServiceObjectVersionsLimitReached

Alert Description: An S3 object is nearing the maximum supported number of versions.

Reason: Versioning is enabled and the number of stored versions is approaching the configured limit.

Action: Monitor the object's version count and remove unnecessary versions to avoid reaching the limit.

Related Error Code: Not Applicable.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.483

Platform: All

CE02809065 PossibleMetadataInconsistency

Alert Description: Verification operation found possible metadata inconsistency.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 28800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.388

Platform: All

CE02809068 ReplicationHandshakeStuck

Alert Description: The handshake for a replication job to the remote cluster is stuck.

Reason: The alert is raised if the remote cluster does not respond to the handshake request for a given job in a given time window.

Action: Check the network connection to see if there's any issue. If not, contact Cohesity Support.

Related Error Code:

Error Codes	Description
`kTimeout`	The handshake request from the local cluster to the remote cluster did not get a response in a given time.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.402

Platform: All

CE02809069 ClockSkewBetweenClusterAndDC

Alert Description: A big clock skew between the Cohesity cluster and Domain Controller, causing the authentication to fail.

Reason: The clock skew between the Cohesity cluster (or Hyx node if existent) and the DC controller is more significant than a threshold.

> **Important:** Time synchronization is crucial for Kerberos authentication. Kerberos tickets use timestamps to prevent replay attacks by rejecting old or expired tickets. If system clocks are out of sync, attackers can exploit the gap to gain unauthorized access. Keeping all systems synchronized ensures security and prevents authentication vulnerabilities.

Action: Perform the following:
- Verify the NTP server configuration on both the Cohesity cluster and the Domain Controller to ensure they are using the same NTP server.
- Use the `iris_cli cluster sync-ntp` command to synchronize the NTP servers on the Cohesity cluster.
- Check the network connectivity between the Cohesity cluster and the NTP server to ensure there are no issues.

Related Error Code:

Error Codes	Description
`kDomainController`	If the clock skew between the Cohesity cluster (or Hyx node if existent) and the DC controller is more significant than a threshold, authentication will fail through this DC. Thus, the time skew is periodically monitored. If it gets close to the threshold, The alert is raised as a reminder to sync up the clock on the DC.

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.405

Platform: All

CE02806031 HighUnreclaimedGarbageAlert

Alert Description: A critically high amount of unclaimed garbage has been detected in the Cohesity Cluster. This can cause an increase in the cluster usage.

Reason: The garbage percentage concerning the cluster capacity is consistently higher (> 15%) for the last 3 Garbage Collection cycles.

Action: Contact Cohesity Support to examine the garbage backlog in the Cohesity cluster.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.403

Platform: All

Metadata Alerts

CE02915002 MetadataServerUnreachable

Alert Description: Metadata server on disk has been unreachable for a long time.

Reason: This alert is triggered when a metadata server in the Cohesity cluster is unreachable for one hour. After the initial alert is triggered, if the metadata server continues to be unreachable, the alert is triggered every 24 hours. The alert can be caused by one of the following conditions:

- A node is powered off.
- A network connectivity issue due to a node or software/hardware issue is causing a metadata server to fail.

Action: Power on any nodes that are powered off. Resolve any network issues between nodes. Contact Cohesity Support if the problem persists.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.110

Platform: All

CE02915003 MetadataUnavailable

Alert Description: Metadata server failures have exceeded the fault tolerance level.

Reason: This alert is triggered when multiple replicas of a metadata server in the Cohesity cluster are unreachable for 1 hour. Metadata cannot be read or updated due to this condition, so the cluster will be unusable. After the initial alert is triggered, if the

majority of the metadata servers continue to be unreachable, the alert is triggered every 24 hours. This alert can be caused by one of the following conditions:

- A node is powered off.
- A network connectivity issue due to a node or software/hardware issue is causing a metadata server in multiple nodes to fail.

Action: Power on any nodes that are powered off. Resolve any network issues between nodes. Contact Cohesity Support if the problem persists.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.111

Platform: All

CE02915004 MetadataShuffleNotProgressing

Alert Description: Metadata shuffle is not progressing.

Reason: This alert is triggered when the Cohesity cluster metadata shuffle is stuck, possibly due to insufficient space on an SSD or another internal issue. This alert is triggered every 24 hours if the condition continues.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.112

Platform: All

CE02915009 MetadataInstanceHasUnrecoverableError

Alert Description: A corruption has been detected on the disk.

Reason: This alert is triggered when the metadata server instance has encountered an unrecoverable error on the disk or a node on the Cohesity cluster.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

CE02915010 MetadataMigrationsNotProgressing

Alert Description: Metadata is unable to migrate to adapt to configuration changes.

Reason: This alert is triggered when the metadata server migrations do not progress as expected.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.240

Platform: All

CE02915011 ScribeBackgroundOpsStalling

Alert Description: Background operations are running for a longer time than expected.

Reason: The alert is triggered when background operations on metadata storage take too long. The metadata server might lack CPU and disk resources.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.250

Platform: All

CE02915012 TablesLargerThanExpected

Alert Description: The table size is larger than an expected fraction of the disk's metadata capacity.

Reason: This alert is triggered when the Distributed Data Store service table size is greater than 20% of the total metadata capacity.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: Not Applicable

SNMP OID: 1.3.6.1.4.1.47421.0.77

Platform: All

CE02915015 SwitchedToSpaceMode

Alert Description: Due to space constraints, the Distributed Data Store service is unable to operate in the performance mode and automatically shifts to the space mode.

Reason: This alert is triggered when the metadata server switches from performance mode to space mode.

Action: No action required. However, there may be a performance impact on the cluster. Please contact Cohesity Support if there are severe performance regressions.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: Not Applicable

SNMP OID: 1.3.6.1.4.1.47421.0.386

Platform: All

CE02915016 MetadataMigrationUsingSlowerDataTransferMode

Alert Description: The metadata server migration process is using a slower record-based data transfer mode.

Reason: This alert is triggered when the metadata server migration is performed using a slower record-based data transfer mode, which causes the migration to take longer and affects performance.

Action: Contact Cohesity Support to investigate this issue and determine if any action should be taken.

Related Error Code: N/A

Severity: Warning

Dedup Interval: 86400

SNMP OID: N/A

Platform: All

Indexing Alerts

CE03008001 SearchClusterHealth

Alert Description: Search cluster unhealthy.

Reason: This alert is triggered when the Cohesity cluster has detected 5 times that the elastic search is down on a node.

Action: If this alert occurs frequently, contact Cohesity Support to investigate if all nodes are running the elastic search.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 1800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.107

Platform: All

CE03008004 SearchNodeBusy

Alert Description: The search node is using more than 80% of the CPU on the machine.

Reason: This alert is triggered when the Cohesity cluster detects that the elastic search is busy on a node. The cluster attempts to kill the elastic search on the node. This alert can be caused by the following conditions:

- There is a high load on elastic search.
- The elastic search is consuming too much memory.

Action: If this alert occurs frequently, contact Cohesity Support to investigate if all nodes are running the elastic search.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE03008005 NonEmptyUnmountedDirectory

Alert Description: The unmounted directory is not empty, it has some entries.

Reason: This alert is triggered when the Cohesity cluster creates a directory to mount an internal view and the Cohesity cluster tries to garbage collect this directory and it is not empty. This alert can be caused by files and directories being written to the mount point or an internal error.

Action: Contact Cohesity Support to investigate.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE03008006 VMWareMountFailure

Alert Description: Too many VMWare mount failures.

Reason: This alert is triggered when the Cohesity cluster cannot index a VM due to an access problem.

Action: If this alert occurs frequently, contact Cohesity Support to investigate.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE03008007 RhinoUnhealthy

Alert Description: Too many Rhino healthchecks failed.

Reason: This alert is triggered when the Cohesity cluster has detected 3 times that Map Reduce Engine status is unhealthy.

Action: Contact Cohesity Support to investigate.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE03008008 VMVolumeIndexingSkipped

Alert Description: One or more disk files missing for indexing.

Reason: This alert is triggered when the Cohesity cluster fails to mount a volume of a VMDK during indexing.

Action: If this alert occurs frequently, contact Cohesity Support to investigate.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: Not Applicable

Platform: All

CE03008009 IndexingBacklog

Alert Description: Backlog in indexing of VM contents.

Reason: This alert is triggered when VM indexing is pending for more than 48 hours. The Cohesity cluster detects that the indexing of the contents of a VM is taking longer than expected.

Action: See the ALERT: CE00508009 INDEXINGBACKLOG KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.201

Platform: All

CE03008010 IndexingPaused

Alert Description: Indexing of the object contents is paused.

Reason: This alert is triggered when the Cohesity cluster detects that the indexing component has less than 10 MB of disk space available on the cluster. The alert is triggered when the Cohesity cluster pauses the indexing of the contents of objects because the cluster detects that available disk space is too low.

Action: Do the following:

- Add more nodes to your Cohesity cluster.
- Clean up Snapshots that are no longer required.
- Archive Snapshots to an External Target.
- Configure Cloud Tier to hold rarely used data.
- Consider changing Storage Domains to use Erasure Coding. Capacity utilization will be reduced over time in the background.
- See the Alert: CE03802121/CE03802122 Cluster Space Usage High KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 14400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.52

Platform: All

CE03008011 CorruptFileSystemNotIndexed

Alert Description: Indexing of objects with corrupt filesystem is skipped.

Reason: This alert is triggered if the indexing of a VM snapshot causes the indexing service to become unresponsive. The alert is only applicable for xfs VMs.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 14400 seconds

SNMP OID: Not Applicable

Platform: All

CE03008013 VMOptimizedNTFSVolumeIndexingSkipped

Alert Description: Optimized NTFS indexing path is skipped Indexing with unoptimized path.

Reason: This alert is triggered when the optimized NTFS indexing path encounters an error and indexing reverts to the unoptimized path. The Cohesity cluster detected that one or more NTFS volumes were not indexed using the optimized indexing path. The indexing for those volumes is attempted again with the unoptimized path.

Action: See the Alert: CE00608013 VMOptimizedNTFSVolumeIndexingSkipped KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.49

Platform: All

CE03008014 VMOptimizedXFSVolumeIndexingSkipped

Alert Description: Optimized XFS indexing path skipped one or more volumes of a server, hence taking an unoptimized path for indexing.

Reason: The Cohesity cluster detected that one or more XFS volumes were not indexed using the optimized indexing path. This alert is triggered weekly.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Severity: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.48

Platform: All

CE03008015 SearchIndexMigrationBacklog

Alert Description: Items are pending in the search index migration work queue of the cluster for hours.

Reason: This alert is triggered when the search index migration to Librarian is taking longer than expected.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 14400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.47

Platform: All

CE03008016 SearchIndexMigrationTriggered

Alert Description: Search index migration from Librarian is triggered on the cluster.

Reason: This alert is triggered when the search index migration to Librarian has been triggered. During the search index migration, any alerts about an indexing backlog can be ignored.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.46

Platform: All

CE03008017 SearchIndexMigrationFinished

Alert Description: Search index migration from Librarian is completed on the cluster.

Reason: This alert is triggered when the search index migration to Librarian has finished.

Action: In the Cohesity Dashboard, do a recovery-based search for files or folders. If there are any search errors, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.45

Platform: All

CE03008018 SearchIndexManualMigrationNeeded

Alert Description: Manual Migration to Librarian is required.

Reason: This alert is triggered when the automatic migration to Librarian is skipped.

Action: Contact Cohesity Support to investigate.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE03018001 SearchIndexSizeExceedsThreshold

Alert Description: The search index size on the disk is exceeds the threshold.

Reason: The search index size on the disk exceeds the threshold. Writes to the disk will start failing when the disk gets full. This alert is triggered daily.

Action: Contact Cohesity Support for help with freeing the disk space.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.113

Platform: All

CE03018002 SearchIndexShuffleNotProgressing

Alert Description: Data shuffle is not progressing.

Reason: The data shuffle for the search index is not progressing. This alert is triggered daily.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.114

Platform: All

CE03018003 SearchIndexUnavailable

Alert Description: The search index cannot achieve the configured fault tolerance factor.

Reason: The search index cannot achieve the fault tolerance factor configured for the cluster. Multiple nodes might be powered down or there might be a software or hardware issue affecting multiple nodes. This alert is triggered daily.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.117

Platform: All

CE03008019 YodaAgentUnhealthy

Alert Description: Yoda agent health check failed.

Reason: This alert is raises when the Yoda agent status is unhealthy. Either the Nexus proxy VM is down or the Yoda agent is not up or crashing frequently.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE03008020 ClearStaleMountPointsTimeout

Alert Description: Timeout while trying to clear stale mount points.

Reason: This alert is triggered if the indexing service fails to clear the stale mount points after 5 minutes. The alerts are triggered at 1-hour intervals.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE03008021 ESDiskUsageExceedsThreshold

Alert Description: ES indices size on the disk is exceeding the threshold.

Reason: This alert is triggered when ES disk usage is about to exceed the reservation limit. All the indices in ES, such as files, audit logs, and so on, might be exceeding the allocated quota of disk space.

Action: Contact Cohesity Support to increase the disk space.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 251

Platform: All

CE03001028 NodeRebootedForYodaStuck

Alert Description: Node is being rebooted as Yoda is stuck.

Reason: This alert is triggered when there is an issue with XFS indexing. Due to bugs in the Linux kernel, XFS volumes' indexing in VMs that do not use app-consistent snapshots may result in instability.

Action: Restart the nodes.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.186

Platform: All

CE03008022 DiskWithCorruptedPartitionTable

Alert Description: VM with disks has a corrupted MBR/GPT partition table.

Reason: The alert is triggered when one or more volumes of the protected VMs have corrupted MBR partition header.

Action: First, confirm if the VM is consistent and is mountable at the source. If not, contact Cohesity Support for any issue during backup.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

CE03008023 RemoteIndexingFallback

Alert Description: Remote indexing for volume has failed. Falling back to local indexing.

Reason: The alert is triggered if the proxy VM or the services running within the proxy VM are not stable during indexing.

Action: You can ignore occasional alerts. If the alert is persistent, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

CE03008024 TenantIndexingConfigMissing

Alert Description: Tenant Indexing cloud config is missing.

Reason: This alert is triggered if WorkItem snapshot was not indexed because the tenant cloud config is not populated yet.

Action: Contact cloud-ops to create an indexing config for the tenant.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.381

Platform: All

CE03008025 StaleDirsFoundInYodaCloneBackupViews

Alert Description: Manual cleanup of Yoda clone backup views is needed to improve the indexing speed.

Reason: This alert is triggered when the Yoda clone backup view contains at least 5 stale directories (older than 14 days). Once triggered, it is checked again after 7 days. This alert can be caused due to unknown bugs while deleting clone directories once indexing finishes.

Action: Contact Cohesity Support to cleanup Yoda clone backup view using available scripts.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.382

Platform: All

CE03008029 TeamsPublicChannelsIndexingAlert

Alert Description: The indexing of public channels for the team is skipped.

Reason: The maximum number of public channels in a Team supported during team indexing is defined by the flag *yoda_slave_o365_teams_max_public_channels_supported*. If any team has more than this number of public channels, the indexing of all public channels is skipped. Otherwise, it results in the indexing service crash loop due to a stack overflow issue (due to deep recursion of function calls). The alert is raised to inform that the public channels indexing is skipped. Indexing of private channels of the team is continued and is not affected.

Action: Contact Cohesity Support for details.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.407

Platform: All

CE03008026 CloneBackupCleanupIssues

Alert Description: The cleanup for the cloned backup directory has failed.

Reason: The alert is triggered when the temporary data needed for indexing cannot be cleaned up.

Action: If the issue persists and you receive recurring alerts, contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.380

Platform: All

CE03008030 IndexingItemSkippedAlert

Alert Description: Some of the indexing items on the object are skipped.

Reason: This alert is triggered when some of the indexing items on the object are skipped due to the configured threshold values in the index backlog.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.410

Platform: All

CE03008031 VMVolumesNotDiscovered

Alert Description: VM volumes could not be discovered.

Reason: This alert is triggered when VM volumes cannot be discovered from the metadata, typically due to filesystem integrity issues.

Action: Verify the filesystem consistency on the VM by running filesystem integrity tools.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 604800 seconds

Severity: 1.3.6.1.4.1.47421.0.429

Platform: NA

CE03008032 LogicalVolumeNamingFallbackAlert

Alert Description: A fallback logical volume name was used during virtual machine indexing.

Reason: The system could not determine the original logical volume name and applied a default name.

Action: Contact Cohesity Support to investigate and resolve incorrect or missing logical volume metadata.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.495

Platform: All

CE03010065 ProtectionJobAutoPausedCredentialFailure

Alert Description: A protection job paused due to a credential or log-on failure.

Reason: This alert is triggered when the protection job pauses due to errors caused by expired or invalid credentials. The product pauses the job to avoid repeated failures and to give you time to correct the credentials.

Action: Update the source credentials for the protection job. When valid updated credentials are detected, the protection job resumes automatically.

Related Error Code: Not applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.505

Platform: All

CE03010066 ProtectionJobAutoResumedCredentialRefresh

Alert Description: A protection job resumed automatically after the credentials were refreshed.

Reason: This alert is triggered when updated credentials are detected and a protection job that was paused due to credential rotation or refresh scenarios resumes.

Action: No action required.

Related Error Code: Not applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.506

Platform: All

Apps Infrastructure Alerts

CE03219001 AppInstanceUnhealthy

Alert Description: Detected an unhealthy app instance.

Reason: This alert is triggered when an unhealthy app instance is detected. An app instance can be unhealthy if some of its components are unavailable. This can be caused by a node reboot, cluster upgrade or insufficient resources.

Action: If a node is rebooting or the cluster is upgrading, wait for the reboot or upgrade to complete before checking app instance health. If resources are not available, consider pausing other app instances to make resources available to the unhealthy app instance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.200

Platform: All

CE03219002 InsufficientResourcesToLaunchSysSrvPod

Alert Description: Failed to launch system service pos for instance.

Reason: This alert is triggered when a node in the Cohesity cluster does not have enough resources to launch a system service pod for the application.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.253

Platform: All

CE03219011 AppInstanceHealthy

Alert Description: Detected healthy application instance.

Reason: This alert is triggered when an app has been successfully running for more than 5 minutes after installation and launch. If an AppInstanceUnhealthy alert was previously raised for the same app instance, this alert resolves it.

Action: No action is required.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.437

Platform: All

CE03219012 NodePresentInKubernetesCluster

Alert Description: Nodes are part of the Kubernetes cluster.

Reason: This alert is triggered when a previously quarantined or removed node is added back to the Kubernetes cluster. If a NodeNotPresentInKubernetesCluster alert was previously raised for the same node, this alert resolves it.

Action: No action is required.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.438

Platform: All

App MarketPlace Alerts

CE03219005 NodeNotPresentInKubernetesCluster

Alert Description: Node(s) are not part of the Kubernetes cluster.

Reason: This alert is triggered when all the nodes in the cluster are not part of the Kubernetes cluster as well. It is triggered periodically each day.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.328

Platform: All

System Service App Alerts

CE04219003 SystemServiceAppInstallationFailed

Alert Description: Failed to install system service app on a particular node in the cluster.

Reason: This alert is triggered for a specific node in the cluster when the installation of a system service app fails. System service apps, which are internal to the cluster, have elevated privileges and access to resources compared to marketplace apps. They must be installed on all nodes in the cluster, and any failure in the installation process on a particular node results in this alert being raised.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.294

Platform: All

CE04219007 SystemServiceAppInstallationSuccess

Alert Description: The System service app is installed successfully on the node.

Reason: This alert is triggered when a system service app is successfully installed on a specific node in the cluster. System service apps, which have elevated privileges and greater resource access than marketplace apps, must be installed on all nodes. When the installation succeeds, this informational alert is generated. Additionally, if a SystemServiceAppInstallationFailed alert exists for the same service and version on the same node, this alert resolves it.

Action: No action is required.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.433

Platform: All

CE04219006 KubernetesControlPlaneNotUp

Alert Description: The Kubernetes control plane is not running on the Cohesity cluster.

Reason: This alert is triggered when the Kubernetes control plane is in an unhealthy state. It should be addressed with priority only if a marketplace application is deployed.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.425

Platform: All

CE04219008 KubernetesControlPlaneUp

Alert Description: The Kubernetes Control Plane is up on the Cohesity cluster.

Reason: This alert is triggered when the Kubernetes control plane is successfully deployed on the Cohesity cluster, which hosts container-based workloads. Once the control plane is up and running, this alert is generated. Additionally, if a KubernetesControlPlaneNotUp alert exists for the same Kubernetes cluster, this alert resolves it.

Action: No action is required.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.434

Platform: All

Data Service Alerts

Download a spreadsheet containing **Data Service** alerts.

Backup and Restore Alerts

Archival and Restore Alerts

Remote Replication Alerts

Quota Alerts

View Failover Alerts

Backup and Restore Alerts

CE00610001 ProtectionGroupSucceeded

Alert Description: Backup run completed successfully.

Reason: This alert is triggered if the Protection Group Run is successful.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE00610002 ProtectionGroupFailed

Alert Description: The backup run of the protection group failed.

Reason: This alert is triggered if a Protection Group Run fails. A protection run can fail for any of the following reasons:

- There is an issue with primary environment such as a removed VM or a Snapshot failure.

- The primary storage is full. (The primary storage contains the Objects that are backed up by the Cohesity cluster.)
- The Cohesity Agent is unreachable while attempting to back up physical servers.
- The storage on the Cohesity cluster is full.

Action: In the Cohesity Dashboard, drill-down on a protection run to get details about the failure including the objects (such as VMs, databases, physical servers etc) that were not backed up. Read the Cohesity and VMware documentation for a resolution. If the resolution is not found, see the Alert: CE00610002 BackupJobFailed KB article.

Related Error Code:

Job Type	Error Codes	Description
All	`kNonExistent`	If the parent entity cannot be found in the registered sources or there are no eligible sources for the protection run, an alert with the error code is raised.
	`kInvalid`	The alert with this error code is raised in the following scenarios: • When the parameters for the host entity's connector cannot be retrieved. • When the retention policy is missing from the protection policy. • When the backup parameters for the Sfdc environment are not configured.
	`kInternalError`	Failed due to error injection. Should not happen in release builds.
	`kDisabled`	The alert with this error code is raised when the Storage Array Snapshots are disabled.
	`kSchedulingError`	The alert with this error code was raised when the scheduled protection run was supposed to start earlier, but it couldn't get scheduled.
	`kRetry`	For Universal Data Adapter jobs that require an AWS S3 view backup, the Iris user config needed to fetch the S3 details was not found.
	`kRejected`	The alert with this error code is raised when a job exceeds its object limit and needs to be split across multiple jobs.

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.108

Platform: All

CE00610003 ProtectionGroupSlaViolated

Alert Description: The SLA of the protection run is violated.

Reason: This alert is triggered if the amount of time that a protection run takes to complete exceeds the amount of time specified in the SLA. This alert can also be triggered due to a failed protection run. A protection run may take longer than the specified SLA for the following reasons:

- If the primary storage is slow.
- The network is slow.
- The Cohesity cluster is overloaded.
- You specified an SLA that is too short.

Action: Investigate why the protection run took longer than the specified SLA. If appropriate, adjust the time period specified in the SLA. If the alert is triggered due to a failed protection run, you can override the default behavior and disable the alert when there is a protection run failure. For more information, see the ProtectionGroupSlaViolated alert triggered for backup jobs which has failed runs KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.121

Platform: All

CE00610004 ProtectionGroupCancelledBlackoutWindow

Reason: This alert is triggered if a protection run aborts because the protection run started executing during a blackout period specified for the job and the Abort Running Jobs during Blackouts toggle is set on for this Protection Group.

Description: By default this toggle is set off which means after a protection run starts, it continues to execute even when a black period specified for this job starts. However, a new protection run will not start during a blackout period.

Action: No action required.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE00610005 ProtectedObjectFailed

Alert Description: The backup failed .

Reason: This alert is raised when the Cohesity cluster detects that an object (such as a VM) failed to be backed up during a Protection Group Run. One alert is raised for each object (such as a VM) that failed to be backed up. For instructions on how to enable this alert, contact Cohesity Support. A protection run can fail to back up an object for any of the following reasons:

- There is an issue with primary environment such as a removed VM or a Snapshot failure.
- The primary storage is full. (The primary storage contains the Objects that are backed up by the Cohesity cluster.)
- The Cohesity Agent is unreachable while attempting to back up physical servers.
- The storage on the Cohesity cluster is full.

Action: See the CE00610005 | BackupRestore - BackupObjectFailed KB article for a resolution.

Related Error Code:

Job Type	Error Codes	Description
kVMware	kNonExistent	The alert with the error code is raised in the following scenarios: • When the virtual machines protected by a VM protection group are not found in the VM hierarchy. • Due to decommissioned or repaved VM.
	kVSphereError	The alert with this error code is raised in the following scenarios: • When a timeout occurs while connecting to the vCenter from the Cohesity cluster. • Pings to the vCenter server fail due to IP connectivity issues. • The account used to register the source doesn't have proper permissions.
	kVixError	The alert with the error code is raised in the following scenarios: • The Cohesity node cannot resolve the DNS name of the ESX host. • The Cohesity node cannot connect to the ESX host over port 902. • When the Cohesity node connecting to the VMDK to read the changed blocks through the ESX host failed.
	kTransportError	The alert with the error code is raised when there is a communication issue between the Cohesity cluster and the Cohesity Agent on the Host.

Job Type	Error Codes	Description
	`kInvalidRequest`	The alert with the error code is raised in the following scenarios: • The source file system inclusion list contains duplicate paths. • The metadata or input file for the File Directive backup or User-based backup has the same file or folder path listed as duplicate. • The exclude path is not a child path of the include path. • The File Directive backup or User-based backup does not specify the metadata or input file when running the backup scripts or specifies an invalid or non-existent file name. • When the inclusion path is not found, the protection run fails. • During a protection run, the Cohesity cluster requests the source server to retrieve a list of available volumes. The protection run fails when a path that has been explicitly included is missing.

Job Type	Error Codes	Description
kPhysicalFiles	kTransportError	The alert with the error code is raised when there is a communication issue between the Cohesity cluster and the Cohesity Agent on the Host.
	kInvalidRequest	The alert with the error code is raised in the following scenarios: • The source file system inclusion list contains duplicate paths. • The metadata or input file for the File Directive backup or User-based backup has the same file or folder path listed as duplicate. • The exclude path is not a child path of the include path. • The File Directive backup or User-based backup does not specify the metadata or input file when running the backup scripts or specifies an invalid or non-existent file name. • When the inclusion path is not found, the protection run fails. • During a protection run, the Cohesity cluster requests the source server to retrieve a list of available volumes. The protection run fails when a path that has been explicitly included is missing.
	kNonExistent	The alert with the error code is raised when metadata or input file part of the File Directive backup or User-based Backup is nonexistent on the source server.

Job Type	Error Codes	Description
	`kFsError`	The alert with the error code is raised in the following scenarios: • Microsoft Windows permissions issue. • The credentials were changed recently. • The domain account used to register the host with the Cohesity cluster is not in the domain(s) to which the cluster has been joined. • The domain account used to register the host with the Cohesity cluster is not in a domain that is trusted by the domain(s) to which the cluster has been joined. • The domain account used to register the host with the Cohesity cluster is in a domain that is denied in the cluster's Active Directory configuration or trusted.
`kCassandra`	`kTransportError`	The alert with the error code is raised when there is a communication issue between the Cohesity cluster and the Cohesity Agent on the Host.
	`kImanisError`	The alert with the error code is raised in the following scenarios: • Source connectivity issue. • The keyspace being backed up is not present in the data center configuration on the Cassandra cluster. • In a Cassandra cluster, keyspaces and tables are spread across nodes; and are created such that data would be spread across nodes based on the configured replication factor. • If for instance RF is 3, and snapshot creation activity fails more than three times, backup would be expected to fail. • Review yarn application log to determine why snapshot creation failed.

Job Type	Error Codes	Description
kNetapp	kInvalidRequest	The alert with the error code is raised in the following scenarios: • The source file system inclusion list contains duplicate paths. • The metadata or input file for the File Directive backup or User-based backup has the same file or folder path listed as duplicate. • The exclude path is not a child path of the include path. • The File Directive backup or User-based backup does not specify the metadata or input file when running the backup scripts or specifies an invalid or non-existent file name. • When the inclusion path is not found, the protection run fails. • During a protection run, the Cohesity cluster requests the source server to retrieve a list of available volumes. The protection run fails when a path that has been explicitly included is missing.

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.109

Platform: All

CE00610006 PolicyFieldsDeprecated

Alert Description: Policy settings have been deprecated.

Reason: This alert is raised after the Cohesity cluster is upgraded from a 4.x release to a 5.x release and the cluster detects that some policy settings used in the current policies on the cluster have been deprecated.

Action: Open the listed policy in the Cohesity Dashboard, verify the current settings and make any necessary adjustments. See the ALERT: CE00610006 POLICYFIELDSDEPRECATED KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE00610007 ProtectedObjectCancelledBlackoutWindow

Alert Description: Backup was cancelled.

Reason: This alert is triggered when the backup of an object is canceled due to a blackout window. The Protection Group uses a policy that defines a blackout window and the job is configured to abort if it runs during the blackout window.

Action: No action required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.124

Platform: All

CE00610009 ProtectedObjectSlaViolated

Alert Description: SLA in backup run violated.

Reason: This alert is triggered when the service level agreement violation (SLA) occurs for an individual object in a Protection Group run. A Protection Group run may take longer than the specified SLA for the following reasons:

- If the primary storage is slow.
- The network is slow.
- The Cohesity cluster is overloaded.
- You specified an SLA that is too short.

Action: Investigate why the Protection Group run took longer than the specified SLA. If appropriate, adjust the time period specified in the SLA.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.130

Platform: All

CE00610011 ProtectionGroupProgressingSlowly

Alert Description: The protection group is progressing slowly.

Reason: This alert is triggered when throttling is observed from Microsoft Office 365 during backups of OneDrives and SharePoint Online. The alert triggered once a day for a Protection Group. This alert is caused when the backups exceed the download or upload limits defined in the Microsoft Office 365 source.

Action: If you repeatedly observe this issue, contact Microsoft Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.139

Platform: All

CE00610014 VMMigrationIdentified

Alert Description: Migrated VM(s) identified on the vCenter.

Reason: This alert is triggered when the Cohesity cluster identifies a VM in a vCenter that was earlier part of another vCenter registered on the Cohesity cluster.

Action: No action is required if the migrated VMs are mentioned in the alert. If not, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.255

Platform: All

CE00608002 MissingVMBackup

Alert Description: Missing VM backup.

Reason: This Critical Alert is triggered when a descriptor VMDK file has been backed up but the corresponding flat VMDK file is missing from the backup.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE00608003 VMCrackingSkipped

Alert Description: VM contents not indexed.

Reason: This alert is triggered when the Cohesity cluster detects 5 consecutive unsuccessful attempts to index a VM. This alert can be caused by the following conditions:

- The Cohesity cluster is not able to mount the VMDK.
- The VM Snapshot has an issue.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE00608012 MissingDiskFiles

Alert Description: One or more disk files are missing for indexing.

Reason: This alert is triggered when one or more disk files are missing for indexing during a backup job, which could lead to issues with the indexing and restore processes.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.98

Platform: All

CE00610013 ContinuousDataProtectionDisabled

Alert Description: CDP is disabled for VM.

Reason: This alert is triggered when CDP is disabled for an entity in a Protection Group. This alert can be caused when an entity cannot be protected using CDP due to the following reasons:

- CDP is not supported for VMware VMs with ESXi host version less than 6.5.
- The source of the VM is an ESXi host and not a VMware vCenter.
- The VM is present on an ESXi host, which is not part of an ESXi cluster.
- The VM has physical RDM disks.
- Another Cohesity cluster is managing the ESXi cluster on which the VMware VM is running.

Action: Fix the validation errors related to the source provided in the alert information. For other errors, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.233

Platform: All

CE00610017 ProtectionPolicyDeleted

Alert Description: A Protection Policy was deleted.

Reason: The alert is triggered when you delete a Protection Policy.

Action: No action required. Using this alert, you can validate if a valid user deleted the Protection Policy.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.234

Platform: All

CE00610018 ProtectionGroupDeleted

Alert Description: A Protection Group was deleted.

Reason: The alert is triggered when you delete a Protection Group.

Action: No action required. Using this alert, you can validate if a valid user deleted the Protection Group.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.291

Platform: All

CE00610019 PolicyDatalockChanged

Alert Description: The DataLock was changed in a Protection Policy.

Reason: The alert is triggered if you enable or disable DataLock for a Protection Policy.

Action: No action required. Using this alert, you can validate if the DataLock was enabled or disabled by a valid user.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.274

Platform: All

CE00610020 PolicyDatalockDurationChanged

Alert Description: The DataLock retention in the Protection Policy was changed.

Reason: The alert is triggered when you change the DataLock duration for a Protection Policy.

Action: No action required. Using this alert, you can validate if a valid user modified the DataLock configuration.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.275

Platform: All

CE00610025 ContinuousDataProtectionAlert

Alert Description: Alert is raised for CDP VM.

Reason: This alert is triggered when a user action is required for CDP protected entity. This alert can be caused when a user action is required in one of the following conditions for CDP protected entity:

- When a user manually manages the VM Storage Policy
- When a user performs Pause, Deactivation or Delete operation on a Protection group

Action: Perform the required action mentioned in the alert. For more details, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.296

Platform: All

CE00610026 DeferSchedule

Alert Description: Deferred schedule to resolve conflicts among multiple schedules.

Reason: The alert is triggered when there are scheduling conflicts among multiple schedules. For example, when multiple schedules are defined in a Protection Policy that can run simultaneously, a specific priority order is applied based on backup type, retention, etc. The alert lets you know that a specific schedule with given retention has been deferred.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.324

Platform: All

CE00610024 SourceConnectivityFailure

Alert Description: The source cannot be reached.

Reason: This alert is triggered when the Cohesity cluster detects that it is no longer connected with the vCenter. The Cohesity cluster failed to connect with the vCenter due to one of the following conditions:

- The user credentials of the vCenter have been updated recently but are not in sync with the Cohesity cluster's vCenter configuration.
- The vCenter is not reachable.

Action: Do one of the following:

- If you get an "incorrect username and password" alert, update the current vCenter credentials in the Cohesity cluster.
- If you get a "timeout" or "not reachable" alert, try the following to ensure the vCenter is accessible by the Cohesity cluster:
  - Validate network connectivity between the Cohesity cluster and the vCenter.
  - Ensure the vCenter Server is running.
  - Diagnose the error using the vCenter logs.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 7200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.292

Platform: All

CE00610030 ProtectionServiceNotInitialized

Alert Description: Protection Service is not initialized.

Reason: This alert is triggered when the Protection service is not initialized for more than 60 minutes.

Description: The Protection service may not initialize on time because the master node is unresponsive.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.346

Platform: All

CE00610027 ObjectDeletionRejected

Alert Description: Could not delete objects in the run before run expiry. They will be deleted when the run snapshot is expired.

Reason: This alert is triggered when the user deletes a specific snapshot for an object in a Protection Group instead of deleting the entire view.

Description: The deletion is rejected because the Protection Group view is marked immutable and therefore individual object deletion can not be performed.

Action: Evaluate if the user can delete the entire view instead of deleting the individual snapshots.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.329

Platform: All

CE00610021 ProtectionPolicyModified

Alert Description: A Protection Policy was modified.

Reason: This alert is triggered when a Protection Policy is modified. The modification might include any changes apart from DataLock-related changes in the policy.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.276

Platform: All

CE00610022 ProtectionGroupModified

Alert Description: A Protection Group was modified.

Reason: This alert is triggered when a Protection Group is modified. The modification might include changing the Protection Group name, adding or deleting any entities in the Protection Group, etc.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.277

Platform: All

CE00610023 ProtectionRunModified

Alert Description: The Protection run in a Protection Group was modified.

Reason: This alert is triggered when a Protection run is modified. The modification might include deleting a Protection run, enabling a legal hold, etc.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.278

Platform: All

CE00610016 ObjectBackupSlaViolated

Alert Description: SLA for backup of the object is violated.

Reason: This alert is triggered when the SLA of object backup is violated. This alert can be generated when the load on the Cohesity cluster is higher than anticipated, or the primary source is loaded and the Cohesity cluster is not able to back it up fast enough.

Action: Customer can verify if a new workload has recently added to the Cohesity cluster or check if the primary source is throttling Cohesity APIs/calls.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.235

Platform: All

CE00610031 NasSourceSnapshotDeletionFailed

Alert Description: The deletion of a snapshot failed.

Reason: This alert is triggered mostly due to transient network connectivity issues between the cluster and the NAS source.

Action: Delete the snapshot created by Cohesity on the NAS source manually.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.363

Platform: All

CE00610032 IsilonChangelistCleanupFailed

Alert Description: The deletion of a changelist failed.

Reason: This alert is triggered mostly due to transient network connectivity issues between the cluster and the Isilon source.

Action: Delete the changelist created by Cohesity on the Isilon source manually.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.364

Platform: All

CE00610033 OracleRestoreChainBreak

Alert Description: There is a restore chain break in the previous backup.

Reason: This alert is triggered when a gap is discovered in between backups which might cause an unrecoverable time range. This alert can be caused by an unknown backup error or due to some backup being deleted inappropriately (manually or due to the policy).

Action: Run a full/incremental backup to fix the log chain break. If a chain break alert is still triggered, contact Cohesity Support or refer to KB for details/resolution.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.362

Platform: All

CE00610034 RestoreTaskFailed

Alert Description: Restore of the object in the job failed.

Reason: This alert is triggered when there is any failure in a restore task.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.369

Platform: All

CE00610035 RestoreTaskWarning

Alert Description: Restore of the object in a job succeeded with warnings.

Reason: This alert is triggered when there is any warning in a restore task. The alert can be caused when a restore task succeeds but with warnings.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.368

Platform: All

CE00610041AgentCertIsAboutToExpire

Alert Description: The certificate associated with the agent is about to expire

Reason: This alert is triggered when the certificate on the agent is about to expire in 30 days

Action: No action required. Certificate renews automatically

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.401

Platform: All

CE00610036 MagnetoUnableToReachYoda

Alert Description: Protection Service is unable to contact the indexing service.

Reason: This alert is triggered when the protection service hasn't made a successful call to the indexing service for a long time and has a long queue of pending requests built up waiting to be sent. This can happen when the indexing service is down or overloaded.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.392

Platform: All

CE00610037 ProtectionGroupWarning

Alert Description: Backup run of protection group finished with warning(s).

Reason: This alert is triggered when a backup run has finished or was canceled and has warnings attached to it.

Related Error Code:

Job Type	Error Codes	Description
All	`kCancelled`	• Backup run was canceled due to hitting a blackout window. • Backup run was canceled as it timed out. • Failed to back up some of the objects in a protection run as they have already been backed up in another protection run.
	`kSetInodeIdFloorFailed`	• We could not set the inode id floor on the backup view. This can cause the next archival to be a full archival.
	`kScriptError`	• Failed to execute the pre-backup/post-snapshot/post-backup script of the job.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.392

Platform: All

CE00610038 ProtectionRunExpiring

Alert Description: Protection run is close to expiry and the replication or archival task has failed.

Reason: This alert is triggered when the replication or archival task has failed on at least one target, and the primary backup is close to expiry.

Action: Extend the retention period on the primary cluster or retry the failed replication or archival task if the retention is needed beyond the current expiry on the primary cluster.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.395

Platform: All

CE00610039 OracleUnresponsiveProgressQuery

Alert Description: The Oracle database is unresponsive consecutively.

Reason: The Oracle database might be busy or unresponsive due to memory issues.

Action: On the Oracle database host, verify the Oracle database status.

Related Error Code:

Error Codes	Description
`kTimeout`	Oracle database might be busy or unresponsive due to memory issues. Please check the database status.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.397

Platform: All

CE00610040 FailedToChangeAuthForBackupView

Alert Description: Failed to set or unset the service authorization configuration for the backup view(s).

Reason: Due to transient errors, the change in service authorization configuration for the backup view(s) failed. The operation is retried in the next iteration.

Action: Verify whether the Cohesity cluster services are stable.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.400

Platform: All

CE00610015 ObjectBackupFailed

Alert Description: Backup of an object failed.

Reason: This alert is triggered when a backup operation for a specific object fails. The alert can occur in multiple contexts, including AWS native environments, Microsoft 365 applications, and VMware environments.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.236

Platform: All

CE00610047 StuckTask

Alert Description: Identified a backup task that became stuck and failed to progress as expected.

Reason: This alert is triggered when a backup task fails to report progress within the expected timeframe.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.452

Platform: All

CE00610051 ExternallyTriggeredBackupError

Alert Description: Error encountered during externally triggered backup.

Reason: This alert is triggered when an externally triggered backup from the source to view encounters an error

Action: If the alert reason is kTimeout, investigate whether the source can access the Bridge and Magneto services of the cluster. Otherwise, investigate the client-side logs based on the alert description and cause.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.465

Platform: All

CE00610052 ExternallyTriggeredRestoreError

Alert Description: Error encountered during externally triggered restore operation.

Reason: This alert is triggered when an externally triggered restore from the source encounters an error.

Action: If the alert reason is kTimeout, investigate whether the source can access the Bridge and Magneto services of the cluster. Otherwise, investigate the client-side logs based on the alert description and cause.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.466

Platform: All

CE00610057 MagnetoStateDBWriteFailure

Alert Description: Error detected during Backup Manager Service state DB write operation.

Reason: This alert is triggered when the Backup Manager Service fails to write state to the database.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.477

Platform: All

CE00610053 KubernetesSourceDeploymentFailure

Alert Description: Service deployment failed after Kubernetes source registration

Reason: This alert is triggered when a service cannot be deployed following the registration of a Kubernetes source.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.469

Platform: All

CE00610054 KubernetesSourceCleanupFailure

Alert Description: Cleanup failed after Kubernetes source deregistration.

Reason: This alert is triggered when the cleanup process does not complete successfully following the deregistration of a Kubernetes source.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.470

Platform: All

CE00610056 GlobalProtectionPolicyUpdateRejected

Alert Description: Failed to apply the protection policy update

Reason: This alert is triggered when the application of the protection policy update is unsuccessful.

Action: The reason for the update rejection is provided in the alert body. For more information, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.472

Platform: All

CE00610055 CertRenewalFailure

Alert Description: Error encountered while renewing certificates for apps linked to the source domain.

Reason: This alert is triggered when an error occurs while renewing certificate for an app in the source domain.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.471

Platform: All

CE00610050 YodaNotifyError

Alert Description: Failed to notify the Indexing Service.

Reason: This alert is triggered when a task fails to notify the indexing service and all retry attempts have been exhausted.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.458

Platform: All

CE00610059 MagnetoMasterHighMemoryUsage

Alert Description: High memory usage detected on the backup service master node.

Reason: This alert is triggered when the backup service master node encounters high memory usage during operation.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 1800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.482

Platform: All

CE00610060 MagnetoHighMemoryUsageDuringRecovery

Alert Description: High memory usage detected during the backup service recovery process.

Reason: This alert is triggered when the backup service experiences unusually high memory consumption while performing a recovery operation.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.484

Platform: All

CE00610058 CertificateExpired

Alert Description: Expired certificate(s) detected.

Reason: This alert is triggered when one or more certificates have passed their expiration date and are no longer valid. Retaining expired certificates may pose security and compliance risks.

Action: Delete the expired certificate manually or contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.481

Platform: All

CE00610061 CassandraSourceSnapshotDeletionFailed

Alert Description: A Cassandra snapshot created for instant recovery could not be deleted.

Reason: The snapshot deletion process failed on the Cassandra source.

Action: Manually delete the snapshot from the Cassandra environment using the snapshot tag/name provided in the alert message. If the problem continues, contact Cohesity Support.

Related Error Code: Not Applicable.

Severity: Warning

Dedup Interval: N/A

SNMP OID: 1.3.6.1.4.1.47421.0.485

Platform: All

CE00610063 SourceRegistrationFailure

Alert Description: A data source could not be registered with the Cohesity cluster.

Reason: The system encountered an error during the source registration process.

Action: Review the error details shown in the UI and refer to the Knowledge Base for troubleshooting guidance.

Related Error Code: Not Applicable.

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.488

Platform: All

CE00610064 MagnetoAPIUnhealthy

Alert Description: Secure Connector response time has exceeded acceptable limits.

Reason: Secure Connector is under high load and unable to respond to API requests quickly.

Action: Refer to the Knowledge Base to investigate Secure Connector performance degradation and apply recommended resolutions.

Related Error Code: Not Applicable.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.491

Platform: All

CE00610062 BackupJobPrePostScriptFailure

Alert Description: Alert is raised when a pre-script or post-script fails during a Physical Server Protection job run.

Reason: This alert is triggered when a pre-backup or post-backup script associated with a Physical Server Protection job (Physical or PhysicalFiles type) fails during execution. The failure can occur due to:

- Script execution error (non-zero exit code)
- Script did not exit properly
- Script encountered a runtime error

Action: Check the alert cause and Protection run's pulse log for more information. Ensure that the script runs successfully when executed externally on the protected host.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.486

Platform: All

Archival and Restore Alerts

CE00810008 ProtectionGroupArchivalBacklog

Alert Description: The archival tasks are queued and might take longer to complete.

Reason: This alert is triggered when a backup job has archival tasks with a high queue wait time for the specified archival targets.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.125

Platform: All

CE00820001 MediaErrorDuringArchival

Alert Description: The archival job is waiting to correct the error.

Reason: This alert is triggered when there are no more tapes available for archiving data.

Action: Add new tapes.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: Not Applicable

Platform: All

CE00820002 MediaErrorDuringRestore

Alert Description: The restore task is waiting to correct the error.

Reason: This alert is triggered when the tape required to restore data is not available in the tape drive. One or more tapes that are required to restore data are not available.

Action: Insert the required tapes to continue with the restore. The required tapes are listed in the alert.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: Not Applicable

Platform: All

CE00820003 kIntegralVolumeReachingMaxCapacity

Alert Description: The current integral volume on the QStar host is reaching its max capacity limit.

Reason: A QStar Integral Volume can handle about 1PB (Petabyte) of data. This alert is triggered when an integral volume is close to this capacity and there are no integral volumes available which have not already reached this threshold.

Action: Add a new QStar Integral Volume.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 21600 seconds

SNMP OID: Not Applicable

Platform: All

CE00820004 ArchiveJobFailed

Alert Description: The crchive task for the job failed.

Reason: This alert is triggered when the archive task of a job fails. This alert is triggered only once per archive task. This alert can be caused by one of the following conditions:

- When an archive task fails due to some external target connectivity and/or credentials issues.
- When an archive task fails due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues. If the archive task is still failing after fixing these issues, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: Not Applicable

Platform: All

CE00820005 ArchiveJobStuck

Alert Description: The archive task for the job is stuck.

Reason: This alert is triggered when the archive task of a job is stuck and does not make any progress for more than 3 hours. If the archive task continues to be stuck/queued, this alert is triggered once a week per archive task. This alert can be caused by one of the following conditions:

- When an archive task doesn't make progress due to some external target connectivity and/or credential issues.
- When an archive task doesn't make progress due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: Not Applicable

Platform: All

CE00820005 ArchiveJobStuck

Reason: This alert is triggered when the archive task of a job is stuck and does not make any progress for more than 3 hours. If the archive task continues to be stuck/queued, this alert is triggered once a week per archive task.

Description: This alert can be caused by one of the following conditions:

- When an archive task doesn't make progress due to some external target connectivity and/or credential issues.
- When an archive task doesn't make progress due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues.

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: Not Applicable

Platform: All

CE00820006 RestoreJobStuck

Alert Description: The restore task for the job is stuck.

Reason: This alert is triggered when the restore task of a job is stuck and does not make any progress for more than one day. If the restore task continues to be stuck/queued, this alert is triggered once a week per Restore task. This alert can be caused by one of the following conditions:

- When a restore task doesn't make progress due to some external target connectivity and/or credentials issues.
- When a restore task doesn't make progress due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues. If the restore task still doesn't make progress after fixing these issues, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: Not Applicable

Platform: All

CE00820007 RestoreJobFailed

Alert Description: Restore task has failed.

Reason: This alert is triggered when the restore task of a job fails. This alert is triggered only once per archive task. This alert can be caused by one of the following conditions:

- When a restore task fails due to some external target connectivity and/or credentials issues.
- When a restore task fails due to an internal Cohesity issue.

Action: Check and fix external target connectivity or credential issues. If the restore task is still failing after fixing these issues, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: Not Applicable

Platform: All

CE00820008 IceboxDedupCacheFull

Alert Description: Icebox dedup cache is full.

Reason: The alert is generated when the number of archived chunks maintained in the Distributed Key Value Store exceeds the default threshold value. When this happens, the effectiveness of deduplication for the archived data is impacted and could potentially result in transferring more data to the external target.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.179

Platform: All

CE00806027 ArchiveValidationFailedAlert

Alert Description: Cohesity has determined an inconsistency in an archive created on the external target.

Reason: This alert is triggered if OrionPipeline detects any inconsistencies in the Cloud Archives (CAv2/CADv2). This alert is raised only once for the given inconsistency.

Action: Contact Cohesity Support and file a field issue with the Cloud-Dataplane-Core team.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

Remote Replication Alerts

CE00910010 ProtectionGroupReplicationBacklog

Alert Description: Protection group has replication tasks with high queue wait time.

Reason: This alert is triggered when a backup job has replication tasks with a high wait time for the specified remote clusters.

Action: Contact Cohesity Support to help resolve this issue.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.131

Platform: All

CE00910012 ProtectionGroupReplicationFailed

Alert Description: Replication to remote cluster failed.

Reason: This alert is triggered when a replication task fails. The alert is caused due to network connectivity issues between the clusters involved in the replication

Action: If you notice error messages related to connectivity, verify the connectivity between the clusters involved in the replication. If not, Cohesity Support.

Related Error Code:

Job Type	Error Codes	Description
All	`kInvalid`	The alert with the error code is raised when the snapshot on the primary Cohesity cluster is deleted before the Cohesity cluster can replicate the snapshot to the remote cluster.
	`kUnsupported`	The alert with the error code is raised when: • An older version of the Cohesity software is running on the remote cluster. • The remote cluster does not support the workload that is being replicated. • The remote cluster does not support cascaded replication, and you attempt to perform cascaded replication.
	`kDisabled`	The alert with the error code is raised for Azure or AWS native workloads if their disaster recovery failback support is disabled
	`kIneligible`	The alert with the error code is raised when no objects are copied to the remote cluster.
	`kCancelled`	The alert with the error code is raised when replication is canceled on the source Cohesity cluster.
	`kPending`	The alert with the error code is raised when an attempt to perform replication fails. After a backoff, the replication is retried.
	`kTransportError`	The alert with the error code is raised when the RPC to the remote cluster to initiate replication fails.
	`kTimeout`	The alert with the error code is raised when the RPC to the remote cluster to initiate replication timed out.
	`kRejected`	The alert with the error code is raised when the remote cluster rejects the replication request. This scenario can happen if the remote cluster has been rebuilt recently.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.231

Platform: All

CE00909011 RemoteReplicationTaskStuck

Alert Description: Replication to the remote cluster is stuck.

Reason: This alert is triggered when a Remote Replication task does not make progress for more than 1 day. It is only triggered once for a single Remote Replication task.

Action: Check the connectivity between the local and remote Cohesity cluster and fix any issues. Check that ports 22, 11111 and 20000 are open. If connectivity is not an issue and Replication Tasks are still stuck, Cohesity Support

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.136

Platform: All

CE00909012 RemoteReplicationTaskFailed

Alert Description: Replication to the remote cluster failed.

Reason: This alert is raised when a Remote Replication task fails. It is only triggered once for a single Remote Replication task.

Action: Check connectivity between local and remote Cohesity cluster and fix any issues. Check that ports 22, 11111 and 20000 are open. If connectivity is not an issue and Replication Tasks are still failing, Cohesity Support

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.137

Platform: All

CE00909013 RemoteReplicationFileSkipped

Alert Description: Replication to the remote cluster skipped an unsupported file.

Reason: This alert is triggered when the local Cohesity cluster replicates a view Snapshot containing special files (pipes, block devices, sockets or character devices) to a remote cluster. The new Snapshot on the remote cluster does not contain these files.

Action: No action required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.94

Platform: All

CE00909016 RemoteReplicationBackupUserQuotaSkipped

Alert Description: Replication to the remote cluster skipped the user quota data backup.

Reason: This alert is triggered on the cluster transmitting replication data. It is triggered during the replication of a view for which a user quota is set, but the user quota settings cannot be extracted and transmitted to the receiving cluster. Replication continues to completion. The view is replicated but the user quota is not set.

Action: Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.92

Platform: All

CE00909017 RemoteReplicationRestoreUserQuotaSkipped

Alert Description: Replication to the remote cluster skipped the user quota data backup.

Reason: This alert is triggered on the cluster receiving replication data. It is triggered during the replication of a view for which a user quota is set, but the user quota cannot be set on the receiving cluster. Replication continues to completion. The view is replicated but the user quota is not set.

Action: Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.93

Platform: All

CE00909043 RemoteReplicationTrackingViewUpdateFailed

Alert Description: Update of tracking view for replication failed.

Reason: This alert is triggered when the tracking view is not updated at the end of replication.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.285

Platform: All

CE00909064 RemoteReplicationEncryptionKeyMismatch

Alert Description: Mismatch in the encryption key for replication to the remote cluster.

Reason: This alert is triggered if the encryption key for replication to a remote cluster differs in both clusters' remote cluster pairing section.

Action: Ensure that the encryption key is the same in both clusters' remote cluster pairing sections. For more details, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.384

Platform: All

CE00910046 ProtectionGroupStuckReplicationTask

Alert Description: The protection group has replication tasks that are not progressing.

Reason: This alert is triggered when one or more replication tasks in the Protection Group fail to make progress for a specified duration.

Action: Perform the following:
- Verify that the clusters are still paired.
- Check if the communication between the clusters is functioning correctly.
- Review for any replication backlog.

If the issue persists, contact Cohesity Support for assistance.

Related Error Code: Not applicable.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.445

Platform: All

CE00909079 VaultWindowIncompleteReplications

Alert Description: Incomplete replication tasks for remote cluster and view box.

Reason: This alert is triggered when there are incomplete replication tasks for the remote cluster and view box at the end of the vault window.

Action: Review the replication tasks for the specified remote cluster and view box. Consider increasing the vault window if required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.487

Platform: FortKnox

Quota Alerts

CE01006005 StorageDomainQuota

Alert Description: Storage domain usage exceeded the quota alert limit.

Reason: This alert is triggered when the Storage Domain physical quota alert threshold is reached. This alert may be triggered daily.

Action: Delete backup Objects such as Snapshots that are stored on the Storage Domain or increase the Storage Domain quota.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE01006006 ViewQuota

Alert Description: View usage exceeded quota alert limit.

Reason: This alert is triggered when the view logical quota alert threshold is reached. This alert may be triggered daily.

Action: Reduce the logical data stored in the view or increase the view logical data quota.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE01006007 UsersExceededQuota

Alert Description: Some users exceeded their logical alert limit. Further writes will be rejected for these users.

Reason: This alert is triggered when one or more users have exceeded their usage quota for the view named in the alert. The users will be unable to write additional data to the view. This alert may be triggered daily.

Action: Have users reduce their view usage or consider increasing their quota. For more information, see User and Directory Quotas for Views. For current quota usage information, run the User Quotas.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE01006008 UserExceededQuotaAlertLimit

Alert Description: User exceeded the logical alert limit.

Reason: This alert is triggered when a user has reached the alert threshold for the view named in the alert. The user can continue to write data to the view until they exceed their actual quota. This alert may be triggered daily.

Action: Have the user reduce their view usage or consider increasing their quota. For more information, see User and Directory Quotas for Views. For current quota usage information, run the User Quotas.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE01006010 DirectoriesExceededQuota

Alert Description: Some directories exceeded their logical alert limit. Further writes will be rejected for these directories as soon as their usage exceeds the configured quota threshold.

Reason: This alert is triggered when one or more directories in the specified view have reached the alert threshold (90% by default) of their directory quota. Data can still be written to the directory.

Action: Bring the directory usage below the quota by deleting files, or increasing the directory quota limit.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE01006011 DirectoryExceededQuotaAlertLimit

Alert Description: Directory exceeded the logical alert limit quota.

Reason: This alert is triggered when a directory in the specified view has reached 100% of its directory quota. Further writes to the directory will be rejected.

Action: Bring the directory usage below the quota by deleting files, or increasing the directory quota limit.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

View Failover Alerts

CE04109039 ViewFailoverStateUpdateFailed

Alert Description: Update of state maintained for failover of the view failed.

Reason: This alert is triggered when any kind of update to the internal state fails during a view failover.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.283

Platform: All

CE04109040 ViewFailoverStuck

Alert Description: Planned failover of the view is stuck.

Reason: This alert is triggered to inform the user to finalize the view failover.

Action: Perform the final failover.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.281

Platform: All

CE04109041 ViewFailoverReplicationFailed

Alert Description: Replication for planned failover of view failed.

Reason: This alert is triggered when a replication is triggered as part of the view failover fails.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.280

Platform: All

CE04109042 ViewFailoverDuplicateViewUID

Alert Description: Multiple views exist on the cluster with the same universal view id.

Reason: This alert is triggered when there are multiple views on the cluster with the same universal id.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.279

Platform: All

CE04109044 ViewFailoverFailed

Alert Description: Unsuccessful failover of the view.

Reason: This alert is triggered when the view failover operation fails.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.282

Platform: All

Data Source Connector Alerts

CE04870004 DataSourceConnectorCertReachingExpiry

Alert Description: The data source connector certificate is nearing expiry.

Reason: This alert is triggered when the certificate for the data source connector is approaching its expiration date.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 1800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.447

Platform: All

CE04870006 DataSourceConnectorPatchFailed

Alert Description: The patch installation for the connector has failed.

Reason: This alert is triggered when the data source connector patch fails to download due to lack of access to the patch URL.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.455

Platform: All

Maintenance Alerts

Download a spreadsheet containing **Maintenance** alerts.

License Alerts
Security Alerts
Upgrade Alerts
Cluster Management Alerts
Networking Alerts
Configuration Alerts
Fault Tolerance Alerts
Storage Usage Alerts

License Alerts

CE01101046 LicenseAboutToExpire

Alert Description: License is about to expire.

Action: Generate a new license from Helios and apply it to the cluster. See Cohesity License.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 207

Platform: All

CE01101047 LicenseExpired

Alert Description: License has expired.

Action: Upload a new license file generated from Helios. See Cohesity License.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 206

Platform: All

CE01101048 LicenseIdentityInvalid

Alert Description: License identity is invalid.

Action: Generate a new license from Helios and apply it to the respective cluster. See Cohesity License.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 212

Platform: All

CE01101050 LicenseServerUnreachable

Alert Description: License server is unreachable.

Action: Check the connection between the cluster and Helios. Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 189

Platform: All

CE01101051 InvalidLicense

Alert Description: License is invalid.

Action: Generate a new license from Helios and apply it to the respective cluster. See Cohesity License.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 190

Platform: All

CE01101052 UnlicensedCluster

Alert Description: Cluster is unlicensed.

Action: License your cluster. See Cohesity License.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 191

Platform: All

CE01101086 LicenseStatsHistoryCollection

Alert Description: License stats history collection.

Action: This is to notify that license usage stats history is being collected.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 604800 seconds

SNMP OID: 228

Platform: All

CE01101112 ClusterPendingAuditReportUpload

Alert Description: Audit report is due.

Reason: This alert is triggered when the Cohesity cluster is not registered with Helios for licensing, even after the grace period of 30 days. The alert is caused if the Cohesity clusters not registered with Helios do not upload the audit report within 90 days.

Action: From the **License** page on the cluster UI, click **Download Audit Report** and follow the steps.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 271

Platform: All

CE01101111 UnregisteredCluster

Alert Description: The cluster is unregistered.

Reason: This alert is triggered when the Cohesity cluster is not registered with Helios for licensing, even after the grace period of 30 days. You must register every Cohesity cluster with Helios for licensing. If a cluster is not registered even after the grace period of 30 days, the alert is raised.

Action: Register the Cohesity cluster with Helios for licensing. For more information, see Cohesity License.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 270

Platform: All

Security Alerts

CE01509025 AVServerConnectionFailed

Alert Description: Antivirus server with ICAP service went down.

Reason: This alert is triggered when the Cohesity Cluster is unable to reach or connect to the registered Antivirus ICAP server.

Action: Check the ICAP server state and make sure the correct ports are open on the AV Server for file scans.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.174

Platform: All

CE01509026 AVServerLicenseError

Alert Description: The license of the Antivirus server with ICAP service is expired.

Reason: This alert is raised when Antivirus server licensing fails.

Action: Update the license of the service.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.56

Platform: All

CE01509027 AVThreatFound

Alert Description: A threat is found in the view.

Reason: This alert is triggered when a virus file is detected by the Antivirus ICAP server.

Action: Review the threats in the Antivirus Threats page in the Cohesity Dashboard by selecting **SmartFiles > Antivirus**.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 300 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.50

Platform: All

CE01512001 EncryptionKeyCreated

Alert Description: A new encryption key has been created.

Reason: This alert is triggered when a new encryption key is generated because a Storage Domain is created and the Storage Domain has either Encryption or Cloud Tier enabled. This alert is triggered for auditing.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE01512002 EncryptionKeyRotated

Alert Description: Encryption key has been rotated.

Reason: This alert is triggered when an encryption key is rotated.

Action: No action. This alert is logged for auditing.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE01512101 KeyRotationPolicyChanged

Alert Description: Key rotation policy has been changed.

Reason: This alert is triggered when the key rotation period is changed. This alert is triggered for auditing purposes.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE01512201 KeyObjectCreated

Alert Description: A key object has been created.

Reason: This alert is triggered when a new key manager is attached to the Cohesity cluster, which can be used to get new encryption keys. This alert is triggered for auditing.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE01512202 KeyObjectDestroyed

Alert Description: A key object has been destroyed.

Reason: This alert is triggered when a new key manager is removed from the Cohesity cluster. This alert is triggered for auditing.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE01512206 KMSError

Alert Description: Error occurred while communicating to KMS server.

Reason: A communication failure or request was rejected by the KMS Server. This alert may be triggered hourly. This alert can be caused by errors in communication, authentication, insufficient privileges, KMS setup, etc.

Action: Review the KMS Server settings.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: Not Applicable

Platform: All

CE01501102 UserDefaultPassword

Alert Description: The cluster is using the default password on the node.

Reason: This alert is triggered when the Cohesity UI or Cohesity CLI '**admin**' user or the shell '**cohesity**' user uses the default password to log in to the node(s) in the cluster.

Action: For added security, ensure that you change the default password of the user accounts. For detailed instructions, see the *Update Default User Passwords* topic in the Cohesity Platform: Security Hardening Guide.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.254

Platform: All

CE01516011 DataIngestAnomalyAlert

Reason: This alert is triggered when an anomalous change in the data ingest rate for a protected Source is detected, and is only generated if the cluster is registered with Helios. The change might be a symptom of a ransomware attack on your primary environment.

Action: Consider restoring the Source from a Snapshot. This alert provides a link to begin an Instant Recovery using the latest clean Snapshot. For more information on how to detect anomalies and ransomware attacks, see Detect Anomalies and Detect Ransomware Attacks.

Severity: Critical

SNMP OID: Not Applicable

Platform: All

CE01501090 AuditLogFailure

Alert Description: The cluster is at risk of failing to process audit logs as required.

Reason: This alert is triggered when the Cohesity cluster fails to generate audit logs for an operation. Issues in writing of the audit log record might cause the alert.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.237

Platform: All

CE01512203 KMSUnreachable

Alert Description: The KMS server is unreachable.

Reason: This alert is triggered if the Cohesity cluster is unable to reach the KMS server. The alert can be caused due to a networking issue, or the KMS server is down.

Action: Fix any networking issue or verify the KMS server status.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.218

Platform: All

CE01512204 KMSCertificateWillExpire

Alert Description: A KMS server certificate will expire soon in the cluster.

Reason: This alert is triggered when the certificates related to the KMS server are about to expire.

Action: Update the certificates related to the KMS server.

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.213

Platform: All

CE01512205 KMSCertificateExpiration

Alert Description: A KMS server certificate will expire soon in the cluster.

Reason: This alert is triggered when the certificates related to the KMS server have expired.

Action: Replace the certificates related to the KMS server immediately.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.214

Platform: All

CE01501104 WormChangeForPolicyAlert

Alert Description: The cluster is at potential security risk due to changes in the policy settings.

Reason: This alert is triggered when a WORM-enabled policy is modified to reduce WORM retention. The alert signifies that the Cohesity cluster is at potential security risk due to the reduction in WORM retention. If WORM retention is reduced and expires before the data lock (WORM) is removed, the snapshots are immediately deleted after WORM gets disabled.

Action: Ensure that you have configured WORM retention appropriately in your policy settings.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.257

Platform: All

CE01501127 PrincipalWithNonViewerRole

Alert Description: The cluster has one or more IRIS principals with roles other than viewer.

Reason: This alert is triggered whenever an AD group, SSO user/group with a non-viewer role is created or updated.

Action: For security reasons, it is recommended that the third-party users do not have any non-viewer role.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.337

Platform: All

CE01501103 WormChangeForGroupAlert

Alert Description: The cluster is at potential security risk due to changes in the protection group settings.

Reason: This alert is triggered when a Protection Group configured with a WORM-enabled policy is updated to associate a Protection Policy without WORM settings. The alert signifies that the Cohesity Cluster is at potential security risk due to the removal of WORM (data lock) settings from the policy associated with an existing Protection Group. It might lead to data loss if the snapshots are deleted after the Protection Group is updated with a Protection Policy without WORM settings.

Action: Ensure that you do not remove WORM (data lock) settings from the policy associated with an existing Protection Group.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.256

Platform: All

CE01501104 WormChangeForPolicyAlert

Alert Description: The cluster is at potential security risk due to changes in the policy settings.

Reason: This alert is triggered when a WORM-enabled policy is modified to reduce WORM retention. The alert signifies that the Cohesity Cluster is at potential security risk due to the reduction in WORM retention. If WORM retention is reduced and expires before the data lock (WORM) is removed, the snapshots are immediately deleted after WORM gets disabled.

Action: Ensure that you have configured WORM retention appropriately in your policy settings.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.257

Platform: All

CE01512207 KMSFailover

Alert Description: KMS in the cluster has failed over to the backup server.

Reason: This alert is triggered when a KMS server fails over to a different IP. It is possible to configure multiple KMS server IPs for KMIP KMS. If one of them is unreachable, the service tries to connect with the next available KMS server IP.

Action: Check whether the KMS servers configured are up, certificates are valid and license is up to date.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.367

Platform: All

CE01509070 InsecureClientDetected

Alert Description: The external client is not using mTLS to connect to the Cohesity cluster.

Reason: Cohesity supports secure communication using mTLS for external clients such as Cohesite plugins such as Oracle SBT, SAP HANA, etc. This alert would report any clients that are not using mTLS-based secure connection.

Action: Identify the type of clients based on the reported source IP addresses of the insecure clients and follow the documentation to move those clients to secure mTLS based connection.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.406

Platform: All

CE01511004 CertificateIsAboutToExpire

Alert Description: The certificate is about to expire alert.

Reason: This alert can be raised across various components, indicating that a specific type of certificate deployed on a specific component is nearing expiration. Certificate type, component, and expiration threshold can be individually configured while defining the alert on the component side.

Action: Contact Cohesity Support for details.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.404

Platform: All

CE01511005 UserAccountIsAboutToBeLocked

Description: The user account is going to be locked soon.

Reason: The user's account has been inactive for a certain period of time.

Action: Please log in to prevent your account from being locked.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.419

Platform: All

CE01512208 KMSImportExportError

Alert Description: Failed to import/export KMS configuration.

Reason: This alert is triggered when the import or export of the KMS configuration fails for the tenant.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.451

Platform: IBM BaaS

CE01511006 AdminAccountWillBeForcedToChangePassword

Alert Description: An admin account will soon be required to change its password.

Reason: The account has met one or more password policy enforcement conditions.

Action: Log in with the admin account and update the password to avoid forced expiration.

Related Error Code: Not Applicable.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.490

Platform: All

CE01516051 DataIngestAnomalyAlert

Alert Description: Anomalous file system changes were detected, indicating possible ransomware activity.

Reason: The system detected unusual data ingestion behavior on the protected source.

Action: Perform Instant Recovery using the most recent clean snapshot and investigate the source environment.

Related Error Code: Not Applicable.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.480

Platform: All

CE01511167 CertInventoryCertAboutToExpire

Alert Description: Certificate approaching expiry on the Cohesity cluster.

Reason: This alert is triggered when the remaining certificate validity reaches 15% based on the default threshold configuration. Customers can override these thresholds using the alert-threshold override configuration.

Action: Renew or replace the certificate before it expires to avoid service disruption.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.507

Platform: All

CE01511167 CertInventoryCertAboutToExpire

Alert Description: Certificate approaching expiry on the Cohesity cluster.

Reason: This alert is triggered when the remaining certificate validity reaches 10% based on the default threshold configuration. Customers can override these thresholds using the alert-threshold override configuration.

Action: Renew or replace the certificate before it expires to avoid service disruption.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.507

Platform: All

CE01511167 CertInventoryCertAboutToExpire

Alert Description: Certificate approaching expiry on the Cohesity cluster.

Reason: This alert is triggered when the remaining certificate validity reaches 5% based on the default threshold configuration. Customers can override these thresholds using the alert-threshold override configuration.

Action: Renew or replace the certificate before it expires to avoid service disruption.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.507

Platform: All

CE01511168 CertInventoryCertExpired

Alert Description: Certificate expired on the Cohesity cluster.

Reason: This alert is triggered when a certificate in the Cohesity cluster inventory was not renewed before its expiry date. Expired certificates can cause service communication failures and security vulnerabilities.

Action: Renew or replace the expired certificate immediately to restore service security.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.508

Platform: All

Upgrade Alerts

CE03301003 UpgradeFailed

Alert Description: Upgrade failure.

Reason: This alert is triggered when an upgrade fails and is aborted. This is triggered once per failed upgrade. After the Cohesity cluster upgrade is initiated, if the first node is unable to download the upgrade package and install, then the upgrade is marked as failed and aborted. Downloading the upgrade package may fail for a number of reasons such as network failures, download checksum mismatch, or mismatched versions.

Action: Investigate why the upgrade failed. After correcting any invalid input, initiate the upgrade again using the Cohesity UI or Cohesity CLI.

Related Error Code:

Error Codes	Description
`KUpgradeSoftwareDownloadError`	The alert with the error code is raised when the Cohesity cluster cannot download the upgrade package.
`KUpgradeHealthCheckError`	The alert with the error code is raised when the Cohesity cluster state is not healthy.
`KUpgradeDiskMarkedForRemoval`	The alert with the error code is raised when a disk in the Cohesity cluster is marked for removal and the upgrade fails.
`KUpgradeFeatureCompatabilityError`	The alert with this error code is raised when the existing version's features matrix is incompatible with the upgrade version.

Error Codes	Description
`KUpgradeManagedBaseOSMissingPackages`	The alert with the error code is raised when the customer-managed OS is missing the required packages and the upgrade fails.

Severity: Warning

Dedup Interval: 600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.4

Platform: All

CE03301006 NewUpgradePackagesAvailable

Alert Description: New upgrade packages are available on the node.

Reason: This alert is triggered when new compatible upgrade packages are available on the Cohesity cloud upgrade package server. If your Cohesity cluster has automatic polling enabled, the cluster raises this alert whenever Cohesity releases a new upgrade package and publishes it to the Cohesity cloud upgrade server.

Action: If desired, upgrade the Cohesity cluster. For more information, see Upgrade Cluster.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.154

Platform: All

CE03301011 UpgradeSucceeded

Alert Description: Upgrade succeeded on the node.

Reason: This alert is triggered when an upgrade finishes successfully. This alert is triggered once per successful upgrade.

Action: Check out the new features or fixed customer issues documented in the Release Notes.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.143

Platform: All

CE03301012 UpgradeStarted

Alert Description: Upgrade started on the node.

Reason: This alert is triggered when an upgrade is started. This alert is triggered once per initiated upgrade.

Action: You can continue to use the Cohesity cluster during the upgrade.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.144

Platform: All

CE03301018 UpgradeStuck

Alert Description: The upgrade is taking longer than usual.

Reason: This alert is triggered when the Cohesity cluster detects that the upgrade of a node is stuck and not progressing.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.68

Platform: All

CE03301020 ServiceVersionsMismatched

Alert Description: Cohesity nodes are running different versions of software in the same cluster.

Reason: This alert is triggered when the software versions are not consistent across all the nodes in the Cohesity cluster. This alert is suppressed during the upgrade and is alerted until it is resolved.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.184

Platform: All

CE03301125 RpmPackageVersionMismatched

Alert Description: The rpm package version is mismatched on the node.

Reason: This alert is triggered when there is a mismatch in rpms or rpm versions installed across the cluster nodes. This alert is triggered every 1 hour. This alert can be caused when the customer/support team installs rpms, but fails to install the same version rpm consistently across the cluster nodes.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.330

Platform: All

CE03301126 BaseOSVersionMismatched

Alert Description: BaseOS version is mismatched on the node.

Reason: This alert is triggered when there is a mismatch in the baseOS version across the cluster nodes. This alert is triggered every hour. This alert can be caused when the customer sets up nodes with different baseOS versions, before installing the Cohesity software.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.331

Platform: All

CE03301133 UpgradeCancelled

Alert Description: Upgrade cancelled on the node.

Reason: This alert is triggered when a user cancels the upgrade.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.361

Platform: All

CE03360000 KnownIssueDetectedPatch

Alert Description: A known issue is detected on the cluster that can be resolved by applying a patch.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.364

Platform: All

CE03360000 KnownIssueDetectedPatch

Alert Description: A known issue is detected on the cluster that can be resolved by applying a patch.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.364

Platform: All

CE03360000 KnownIssueDetectedPatch

Alert Description: A known issue is detected on the cluster that can be resolved by applying a patch.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.364

Platform: All

CE03360004 KnownIssueDetectedHotfix

Alert Description: A known issue is detected on the cluster that can be resolved by applying a hotfix.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.368

Platform: All

CE03360004 KnownIssueDetectedHotfix

Alert Description: A known issue is detected on the cluster that can be resolved by applying a hotfix.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.368

Platform: All

CE03360004 KnownIssueDetectedHotfix

Alert Description: A known issue is detected on the cluster that can be resolved by applying a hotfix.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.368

Platform: All

CE03360007 ImportantAwarenessCommunication

Alert Description: There is a general awareness message available for your platform.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.365

Platform: All

CE03319009 AppRollingUpgradeFailure

Alert Description: The app failed the rolling upgrade.

Reason: This alert is triggered when a rolling upgrade fails for an app after a new version becomes available. Rolling upgrades ensure seamless updates without downtime, but if the process encounters an issue, this alert is generated to indicate the failure.

Action: The issue typically resolves after a few retries. If the problem persists, contact the Cohesity Support team for further assistance.

Related Error Code: Not Applicable.

Severity: Warning

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.435

Platform: All

CE03319010 AppRollingUpgradeSuccess

Alert Description: The app's rolling upgrade succeeded.

Reason: This alert is triggered when a rolling upgrade for an app successfully completes after a new version becomes available. If an AppRollingUpgradeFailure alert was previously raised for the same app and version, this alert resolves it.

Action: No action is required.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.436

Platform: All

Cluster Management Alerts

CE03401016 AddNodeFailed

Alert Description: The node failed to join the cluster.

Reason: This alert is triggered when an unexpected error occurs when adding a node to a Cohesity cluster.

Action: Contact Cohesity Support to investigate.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.73

Platform: All

CE03401039 RemoveNodeSuccess

Alert Description: The removal of a node is successful.

Reason: The node was removed successfully. This alert is triggered daily.

Action: If a node is intentionally being removed from the cluster, no action is required. Otherwise, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.161

Platform: All

CE03401040 ForcefulRemoveNodeWarning

Alert Description: The node is being removed from the cluster forcefully.

Reason: This alert is triggered daily.

Description: The node is forcefully being removed from the Custer.

Action: If a node is intentionally being removed from the cluster, no action is required. Otherwise, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.173

Platform: All

CE03401041 AddNodeSuccess

Alert Description: The node has joined the cluster.

Reason: This alert is triggered every hour.

Description: The node has successfully joined the cluster.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.172

Platform: All

CE03401084 AddNodesStarted

Alert Description: The node has joined the cluster.

Reason: This alert is triggered when node addition has started.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.221

Platform: All

CE03401085 NodeMarkedForRemoval

Alert Description: The node is marked for removal from the cluster.

Reason: This alert is triggered when the node is marked for removal.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.221

Platform: All

CE03401119 CreateSnapshotStarted

Alert Description: Create snapshot process started on the cluster.

Reason: The alert is triggered when a snapshot creation operation is started.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.336

Platform: PXG1 only

CE03401120 CreateSnapshotSucceeded

Alert Description: Create snapshot process succeeded on the cluster.

Reason: The alert is triggered when a snapshot is successfully created.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.322

Platform: PXG1 only

CE03401121 CreateSnapshotFailed

Alert Description: Create snapshot process failed on the cluster.

Reason: The alert is triggered when snapshot creation fails.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.332

Platform: PXG1 only

CE03401122 RestoreSnapshotStarted

Alert Description: Restore snapshot started on the cluster.

Reason: The alert is triggered when you start to restore a snapshot.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.333

Platform: PXG1 only

CE03401123 RestoreSnapshotSucceeded

Alert Description: Restore snapshot succeeded on the cluster.

Reason: The alert is triggered when the snapshot is successfully restored.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.334

Platform: PXG1 only

CE03401124 RestoreSnapshotFailed

Alert Description: Restore snapshot failed on the cluster.

Reason: The alert is triggered when the snapshot fails to restore.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.335

Platform: PXG1 only

CE03401145 SecureShellEnabled

Alert Description: Secure shell is enabled on the cluster.

Reason: The alert is triggered when the Secure Shell is enabled on the Cohesity cluster.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.393

Platform: All

CE03401146 SecureShellDisabled

Alert Description: Secure shell disabled on the cluster for an hour.

Reason: The alert is triggered when the Secure Shell is disabled on the Cohesity cluster for an hour.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.394

Platform: All

CE03415014 MetadataMovedToFlashblade

Alert Description: All the metadata has moved to the pure flash array.

Reason: This alert is triggered when metadata moves to Flashblade.

Action: No action required.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.316

Platform: Pure FlashBlade

CE03415013 DiskIdMismatch

Alert Description: Metadata disk re-addition error due to improper disk removal, expected disk ID is different from the saved disk ID.

Reason: This alert is triggered due to an improper disk removal and re-addition.

Description: This alert can be caused when the scribe folder is not deleted before the re-addition.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.232

Platform: All

CE03401136 UILoginSuccess

Alert Description: UI login is successful.

Reason: This alert is triggered when a user is able to log in successfully with `allow_ui_login_alerts` GFLAG enabled.

Action: Check if the login is done by an authorized user.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.373

Platform: All

CE03401137 UILoginFailure

Alert Description: UI login failed.

Reason: This alert is triggered when a user login fails with `allow_ui_login_failure_alerts` GFLAG enabled.

Action: Check if the UI login is done by an authorized user or any unauthorized user is trying to log in with an invalid password.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.372

Platform: All

CE03401138 UIUserLockout

Alert Description: UI user locked out.

Reason: This alert is triggered when a user lockout happens with `allow_user_lockout_alerts` GFLAG enabled.

Action: Check if the UI login is done by an authorized user, or any unauthorized user is trying to log in with an invalid password.

Related Error Code: Not Applicable

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.371

Platform: All

CE03401141 SSHLoginSuccess

Alert Description: SSH login is successful.

Reason: This alert is triggered when the `cohesity` user is able to log in successfully through SSH with `allow_ui_login_alerts` GFLAG enabled. A successful SSH login by the `cohesity` user triggers this alert.

Action: Check if the login is done by an authorized user.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.376

Platform: All

CE03401142 SSHLoginFailure

Alert Description: SSH login failed.

Reason: This alert is triggered when a user's SSH login attempt to a cluster node fails. This includes failures caused by an incorrect username or password, an account locked out after repeated failed attempts, or a failed second-factor verification when multi-factor authentication is configured.

Action: Check if the login is done by an authorized user.

Related Error Code: Not Applicable

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.377

Platform: All

CE03460001 KnownIssueDetectedGflag

Description: A known issue is detected on the cluster that can be resolved by applying a gflag.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.363

Platform: All

CE03460001 KnownIssueDetectedGflag

Description: A known issue is detected on the cluster that can be resolved by applying a gflag.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.363

Platform: All

CE03460001 KnownIssueDetectedGflag

Description: A known issue is detected on the cluster that can be resolved by applying a gflag.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.363

Platform: All

CE03460002 KnownIssueDetectedCommand

Description: A known issue is detected on the cluster that can be resolved by running cli commands.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.362

Platform: All

CE03460002 KnownIssueDetectedCommand

Description: A known issue is detected on the cluster that can be resolved by running cli commands.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.362

Platform: All

CE03460002 KnownIssueDetectedCommand

Description: A known issue is detected on the cluster that can be resolved by running cli commands.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.362

Platform: All

CE03460003 KnownIssueDetectedAlert

Description: A known issue is detected on the cluster.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.361

Platform: All

CE03460003 KnownIssueDetectedAlert

Description: A known issue is detected on the cluster.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.361

Platform: All

CE03460003 KnownIssueDetectedAlert

Description: A known issue is detected on the cluster.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.361

Platform: All

CE03460005 KnownIssueDetectedDoc

Description: A known documented issue is detected on the cluster.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.367

Platform: All

CE03460005 KnownIssueDetectedDoc

Description: A known documented issue is detected on the cluster.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Warning

SNMP OID: 1.3.6.1.4.1.47421.0.367

Platform: All

CE03460005 KnownIssueDetectedDoc

Description: A known documented issue is detected on the cluster.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.367

Platform: All

CE03460006 NewPatchAvailable

Description: This alert is triggered when a new patch is available for your cluster.

Action: Click *http://cluster_ip/siren/v2*, navigate to the Siren Support Toolbox page, and select Known Issues for more information.

Severity: Informational

SNMP OID: 1.3.6.1.4.1.47421.0.366

Platform: All

CE03401155 ConsoleLoginSuccess

Alert Description: The console login is successful.

Reason: This alert is triggered when a console login is successful.

Action: No action needed

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0. 460

Platform: All

CE03401156 ConsoleLoggedOut

Alert Description: The console has been logged out.

Reason: This alert is triggered when the console is logged out.

Action: No action needed

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.461

Platform: All

CE03401157 ConsoleLoginFailure

Alert Description: The console login failed.

Reason: This alert is triggered when the console fails to log in.

Action: Check if the login is done by an authorized user.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0. 459

Platform: All

Networking Alerts

CE03613002 NodePingFailed

Alert Description: The node is not pingable.

Reason: This alert is triggered when a node is unreachable. (In previous releases, this alert was NodeIsDown.)

Action: See the ALERT: CE03613002 NodePingFailed KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.2

Platform: All

CE03613029 LinkStatsIsNotNormal

Alert Description: Errors are detected in the network's physical layer.

Reason: This alert is triggered when a node's network interface statistics are abnormal.

Action: Check the network connections.

Related Error Code:

Error Codes	Description
`kLinkRxCRCError`	Ethernet Checksum errors exceed the threshold. The alert with the error code is raised when the error percentage exceeds the threshold percentage of the incoming packets within a five-minute interval on an interface.
`kLinkLscIntError`	The state of the link state is frequently changing. The alert with the error code is raised when the state of the link changes the number of times <actual_count> that exceeds the set threshold <threshold>.

Error Codes	Description
kLinkRxError	The number of ethernet errors exceeds the threshold. The alert with the error code is raised when the percentage of error percentage <error_percentage> exceeds the threshold percentage <threshold_percentage> of the incoming packets within a five-minute interval on an interface.
kLinkRXCSumOffloadError	The number of NIC checksum offload errors exceeds the threshold. The alert with the error code is raised when the number of NIC checksum offload errors <error_count> exceeds the threshold <threshold>.
kLinkRXFlowControlError	NIC Rx frame buffer backlog exceeding the threshold. actual count: ${cur_count}; threshold: ${threshold}

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.142

Platform: All

CE03613030 SupportChannelAboutToClose

Alert Description: Cohesity support channel about to close on the cluster.

Reason: This alert is triggered every five minutes until the Support Channel is closed.

Action: Reopen the Support Channel if necessary. For more information, see Manage the Support Channel.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 1800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.210

Platform: All

CE03601013 BondingModeUnexpected

Alert Description: Bonding mode does not match the expected configuration on the node.

Reason: This alert is triggered by interface bond0 and not by any other bonding interfaces on the cluster and the alert may be triggered daily. The alert can be caused by one of the following conditions:

- The bonding mode is changed directly on the bond interface of node(s) but the bonding mode configured on the Cohesity cluster is not updated.

- The bonding mode configured on the Cohesity cluster is changed but the bonding mode on the bond interface of a node is not updated.

Action: Using the Cohesity CLI, correct the bonding mode configured on the Cohesity cluster. See the ALERT: CE03601013 BONDINGMODEUNEXPECTED KB article.

Related Error Code:

Error Codes	Description
`kBondingModeError`	The alert with the error code is raised when the bonding mode is incorrect.

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.145

Platform: All

CE03601014 MtuUnexpected

Alert Description: MTU does not match the expected configuration on the node.

Reason: This alert is triggered when the MTU configured on the bond interface of node (s) is different than the MTU specified in the cluster configuration. This alert can be caused by one of the following conditions:

- The MTU is changed directly on the bond interface of node(s) and the MTU configured on the Cohesity cluster is not updated.
- The MTU configured on the Cohesity cluster is changed but the MTU on the bond interface of a node is not updated.

This alert may be triggered daily.

Action: Using the Cohesity CLI, correct the MTU configured on the Cohesity cluster.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.140

Platform: All

CE03601015 GatewayUnexpected

Alert Description: The gateway configuration does not match the expected configuration on the node.

Reason: This alert is triggered when the gateway configured on the bond interface of node(s) is different than the gateway specified in the cluster configuration. This alert can be caused by one of the following conditions:

- The gateway is changed directly on the bond interface of node(s) and the gateway configured on the Cohesity cluster is not updated.
- The gateway configured on the Cohesity cluster is changed but the gateway on the bond interface of a node is not updated.

This alert may be triggered daily.

Action: Using the Cohesity CLI, correct the gateway configured on the Cohesity cluster.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.141

Platform: All

CE03601021 ProxyServerUnreachable

Alert Description: Proxy server is unreachable on the node.

Reason: This alert is triggered when the Cohesity Cluster cannot connect to the specified Proxy Server.

Action: Check the network and proxy server configurations. If you are unable to resolve the issue, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.183

Platform: All

CE03601031 IpMigrationOnBondFailed

Alert Description: The node migration to bond failed.

Reason: This alert is triggered when the Cohesity cluster is unable to create bond interfaces because the physical interfaces have more than one IP address.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.165

Platform: All

CE03601032 RoutesMigrationOnBondFailed

Alert Description: The node routes migration to bond failed.

Reason: This alert is triggered when the Cohesity cluster is unable to create bond interfaces because the physical interfaces have more than one static route file.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.162

Platform: All

CE03601035 EtcHostsFileError

Alert Description: The etc hosts file on node has errors.

Reason: The etc host file on the specified node contains errors. This alert is triggered weekly until resolved.

Action: Review and correct the hosts file content.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.169

Platform: All

CE03601036 FirewallError

Alert Description: The cluster firewall on the node has errors.

Reason: The cluster firewall on the specified node has errors and the alert is triggered weekly until resolved.

Action: See the Alert: CE03601036 Cluster firewall on node has errors KB article. If you are unable to resolve the issue, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.166

Platform: All

CE03601053 VerifyRemoteClusterConnection

Alert Description: The cluster requires a connection confirmation from another remote cluster.

Reason: The cluster needs to confirm a connection to a remote cluster. This alert is triggered yearly until resolved.

Action: Do one or more of the following:

- Verify the remote cluster connection on both Clusters: Log in to the Cohesity Dashboard, select **Infrastructure > Remote Clusters** and verify the remote cluster connection and details.
- Verify the remote cluster connection by using the CLI on both clusters. Run the command: `iris_cli remote-cluster ls`.
- Verify the static routes on both clusters by using the CLI. Run the command: `iris_cli route ls`.
- Verify the static route file, for example:

```
[cohesity@node-1 ~]$ cat /etc/sysconfig/network-scripts/route-
bond1198.16.0.0/20 via 10.2.32.1 dev bond1198.16.10.179/32 via
10.2.32.1 dev bond1198.16.10.180/32 via 10.2.32.1 dev
bond1198.16.10.181/32 via 10.2.32.1 dev bond1
```

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 31536000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.192

Platform: All

CE03601059 NtpServerUnReachable

Alert Description: NTP server is not reachable for the cluster.

Reason: Incorrect input was provided for creating the cluster or there is a networking issue in the current node. This alert is triggered when the NTP server is not reachable.

Action: Check the state of the NTP server and provide an alternate NTP server if necessary.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.198

Platform: All

CE03601060 DnsServerUnReachable

Alert Description: The DNS server is not reachable for the cluster.

Reason: Incorrect input was provided for creating the cluster or there is a networking issue in the current node.

Action: Check the state of the DNS server and provide an alternate DNS server if necessary.

Related Error Code:

Error Codes	Description
`kDNSserverUnreachableError`	The alert with the error code is raised when the DNS server is not reachable for the Cohesity cluster.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.81

Platform: All

CE03601061 ClusterGatewayUnReachable

Alert Description: The cluster gateway is not reachable from the node.

Reason: Incorrect input was provided for creating the cluster or there is a networking issue in the current node. This alert is triggered when a cluster gateway is not reachable.

Action: Check the state of the cluster gateway and provide an alternate gateway if necessary.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.80

Platform: All

CE03601062 IpmiNetworkConfigFailed

Alert Description: IPMI network config failed on the node.

Reason: Incorrect input was provided for creating the cluster or there is a networking issue in the current node. This alert is triggered when an IPMI network configuration failed on a node.

Action: Check and reconfigure the IPMI on the node.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.38

Platform: All

CE03601334 VLANUnexpected

Alert Description: Unexpected VLAN found on the node.

> **Note:** This alert is triggered with different severities based on the scenario.

Reason: This alert is triggered when an unexpected VLAN is found on a node. The unexpected VLAN was not configured by the cluster. The VLAN's VIPs are duplicated, so the alert severity is Critical.

Action: Remove the unexpected VLAN or reconfigure the VIPs so they are not duplicated.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.146

Platform: All

CE03601065 VerifyVlanHostName

Alert Description: Error while verifying hostname entry in DNS server for the interface group.

Reason: This alert is triggered during verification of the hostname entry on the DNS server for the specified interface group. Incorrect input was provided for creating the cluster or there is a networking issue in the current node.

Action: Change the hostname in the DNS server or enter the hostname that matches the hostname in the DNS server.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.82

Platform: All

CE03601066 BondModeAutoMigrated

Alert Description: Node bond mode has been migrated.

Reason: This alert is triggered when the bond mode of the specified node and IP address has been automatically migrated from bond mode=6 (balance-alb) to bond mode=1 (active-backup).

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE03601072 UnexpectedInterfaceIpSetting

Alert Description: Static IP on some interfaces is not configured in the interface group.

Reason: This alert is triggered when an interface group has interfaces configured with and without static IP addresses. All the interfaces in the group must be configured the same (all with static IP addresses or no static IP addresses).

Action: Use the `iris_cli cluster update-interfaces` command on the node to configure these interfaces.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 21600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.64

Platform: All

CE03601075 BondSlavesSpeedInconsistent

Alert Description: Bond slaves' speed is inconsistent on the interface.

Reason: This alert is triggered when there is a periodic monitoring of the system. The alert can be caused if the speed of different network interfaces are inconsistent.

Action: Adjust the physical port/interface to make the speed consistent.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.63

Platform: All

CE03601077 BondUnhealthy

Alert Description: The bond interface is unhealthy on the node.

Description: Either the bond interface is unhealthy due to the interface down or there is an aggregatorId mismatch across different interfaces under this bond device.

Action: Check and correct network interfaces and interface configuration.

Related Error Code:

Error Codes	Description
`kBondUnhealthyError`	The alert with the error code is raised when the bond interface is unhealthy.

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.62

Platform: All

CE03601078 SwitchCongestion

Alert Description: Ethernet traffic is experiencing congestion on the interface on a node.

Reason: This alert is triggered if the interface stats `rx_flow_control_xon` exceeds the threshold value, 30000. The alerts are triggered in intervals of 5 minutes .

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.59

Platform: All

CE03601080 InconsistentStaticRoutes

Alert Description: One or more inconsistent static routes in the cluster configuration on the node.

Reason: The static routes are compared and an alert is triggered if a difference is found between the cluster configuration and nodes. The alerts are triggered every 5 minutes until the issue is resolved.

Action: Contact Cohesity Support.

Related Error Code:

Error Codes	Description
kExtraStaticRoutes	The alert with the error code is raised when there are additional routes in the routing table than what is configured in the Cohesity cluster configuration.
kMissingStaticRoutes	The alert with the error code is raised when there are missing routes in the routing table than what is configured in the Cohesity cluster configuration.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.89

Platform: All

CE03609036 NisDomainNotReachable

Alert Description: NIS domain is not reachable.

Reason: This alert is triggered when multiple NIS requests fail due to server reachability issues. This alert can be caused by one of the following conditions:

- The NIS servers are down.
- The NIS servers are not properly configured in DNS.
- Due to network issues.

Action: Verify the NIS servers' state and reachability. Also, verify the DNS configuration.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 86

Platform: All

CE03609037 DnsServerConnectionFailed

Alert Description: DNS server is not reachable.

Reason: This alert is triggered when the cluster fails to retrieve the SRV records from the configured DNS servers. If the issue is not resolved, the alert is triggered again after 10 minutes. The DNS server settings are not configured correctly.

Action: Verify the DNS settings in your cluster.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.229

Platform: All

CE03609035 KeystoneNotReachable

Alert Description: Keystone is not reachable.

Reason: This alert is triggered when an invalid Keystone server is configured on the Cohesity cluster. The alert can be caused due to the following reasons:

- Incorrect information is provided in the Keystone configuration on the Cohesity cluster. The information includes URL, domain name, user name, password, and project name.
- Admin credentials are not provided in the Keystone configuration on the Cohesity cluster.
- The Keystone server is not reachable due to a network issue.

Action:

- Verify the information provided in the Keystone configuration on the Cohesity cluster.
- Suppose the Keystone server is not reachable due to a network issue, then fix the network issue.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.157

Platform: All

CE03601105 InconsistentInterfaceGroupSpeed

Alert Description: The Interface speed is inconsistent on the node.

Reason: This alert is triggered when the interface group has a mismatching interface across nodes in the cluster. This alert can be caused:

- If one of the interfaces is down in a bond.
- If the interface is negotiating different speeds with the peer switch or router.

Action: Check the interface operations configured in the interface group.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.258

Platform: All

CE03601105 InconsistentInterfaceGroupSpeed

Alert Description: The Interface speed is inconsistent on the node.

Reason: This alert is triggered when the interface group has a mismatching interface across nodes in the cluster. This alert can be caused:

- If one of the interfaces is down in a bond.
- If the interface is negotiating different speeds with the peer switch or router.

Action: Check the interface operations configured in the interface group.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.258

Platform: All

CE03601095 RoutingDaemonDown

Alert Description: The routing daemon is down.

Reason: This alert is triggered when the routing daemon is down in the system. The BGP configuration can cause this alert in the node.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.244

Platform: All

CE03601056 BridgeMigrationRevertFailed

Alert Description: The bridge migration revert failed on the node.

Reason: This alert is triggered when reverting the bond interface migration to the ovs bridge interface failed on the node.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.195

Platform: All

CE03601055 BridgeMigrationFailed

Alert Description: The bridge migration failed on the node.

Reason: This alert is triggered when the bond interface migration to the ovs bridge interface fails on the node.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.194

Platform: All

CE03601101 DnsServerUnassigned

Alert Description: DNS server is not assigned on the node.

Reason: This alert is triggered when a DNS server is not assigned to one or more nodes in the Cohesity cluster.

Action: Verify the DNS configuration.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.252

Platform: All

CE03601087 AppNetworkMasterUnhealthy

Alert Description: App Network Master status is unhealthy.

Reason: This alert is triggered when the master axon error checks fail. The network virtualization controller might have issues.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.234

Platform: All

CE03601088 AppNetworkResourceUnhealthy

Alert Description: App Network Resources status is unhealthy.

Reason: This alert is triggered when axon resources' error checks fail. The network virtualization resources are in a bad state.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 604800 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.236

Platform: All

CE03601030 ConfigureVipFailed

Alert Description: Failed to configure VIP on the node.

Reason: This alert is triggered when the Cohesity cluster is unable to configure a virtual IP (VIP) address on a node.

Action: Contact Cohesity Support.

Related Error Code:

Error Codes	Description
`kVipConfigurationError`	The alert with the error code is raised when the VIP configuration has failed.

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.164

Platform: All

CE03613037 NodeConnectivityFailed

Alert Description: Node connection failure.

Reason: The alert is triggered when the TCP connections on any of the following ports fail for the node-to-node communication:

- 2049
- 11111
- 12222
- 23456
- 23457
- 20000
- 22222

The alert is triggered after three retry attempts. This alert can be caused by one of the following conditions:

- There is a network issue that prevents connection on any of the ports.
- The applications listening on either one of the ports are down, so the TCP connection failed.

Action: Perform the following:

- Verify the networking configuration to check the ping connectivity between the nodes is working.
- Verify if the Gateway ping is working.
- Verify if any of the Cohesity services are failing.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.323

Platform: All

CE03613038 NetworkStatsIsNotNormal

Alert Description: Node netstats is not normal.

Reason: The alert is triggered when TCP/UDP/IP checksum error counters are more than 10% of the actual TCP/UDP/IP packets received in a node. This alert can be caused by one of the following conditions.

- There is an issue with the network cable.
- There is an issue with the SFP Hardware.
- The link speed is not negotiated accurately with the peer switch.
- The packet gets corrupted in the network if there is an error in any of the intermediate switches or routes.

Action: Perform the following:

- Check the network cable.
- Verify the SFP re-seat
- Verify if the peer switch has auto-neg turned on.
- Check if there are any network issues.
- Check the peer node or switch to find any particular hardware issues.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.320

Platform: All

CE03619004 KubernetesObjectsCleanupAlert

Alert Description: Failed to clean up Kubernetes objects on stale incarnation id.

Reason: When the Marketplace Apps Services subnet is updated, all the Kubernetes objects are deleted automatically. The Kubernetes objects are stuck in a terminating state in certain scenarios, triggering the alert.

Action: Manually delete all the Kubernetes namespaces.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.327

Platform: All

CE03601129 AppsModeDisabledDueToSubnetClash

Alert Description: Apps mode and Marketplace apps mode are disabled due to a subnet clash.

Reason: This alert is triggered if the configured Apps Subnet clashes with any other subnets configured in the cluster.

Action: In the Cohesity Dashboard:

1. Go to **Settings > Summary** and click **Configure**.
2. In the **Configure Apps management network**field, set a non clashing subnet for apps management.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.344

Platform: All

CE03609045 HighRPCLatencies

Alert Description: RPC latencies for metadata service between nodes are high.

Reason: This alert is triggered due to poor network conditions or if there is an error during the creation of the cluster or the network configuration.

Action: Contact Cohesity Support. and verify the network setup.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 2592000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.288

Platform: All

CE03613039 VipPingFailed

Alert Description: The VIPs are not pingable from the node.

Reason: Ping to the VIPs failed.

Action: Check if the VIP is programmed in the nodes or not. Make sure that the connectivity exists for that node.

Related Error Code:

Error Codes	Description
`kNonExistent`	If the parent entity cannot be found in the registered sources or there are no eligible sources for the protection run, an alert with the error code is raised.

Severity: Critical

SNMP OID: 1.3.6.1.4.1.47421.0.396

Platform: All

CE03619013 InvalidAppsSubnet

Alert Description: Invalid apps subnet provided.

Reason: This alert is triggered when the provided subnet prefix length is greater than 30 for IPv4 and greater than 126 for IPv6.

Action: Ensure the app's subnet is within the valid range, which should be less than 30 for IPv4 and less than 126 for IPv6.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.442

Platform: All

CE03619014 ValidAppsSubnet

Alert Description: Valid apps subnet provided.

Reason: This alert is triggered when the provided subnet prefix length is less than or equal to 30 for IPv4 and less than or equal to 126 for IPv6.

Action: No action is required

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 2592000 seconds

Severity: 1.3.6.1.4.1.47421.0.443

Platform: All

CE03601158 NetworkManagerMigrationFailure

Alert Description: Error encountered during network manager migration.

Reason: This alert is triggered when a node reports an error while migrating the network manager.

Action: Contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 21600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.462

Platform: All

CE03613040 SupportChannelExtension

Alert Description: Support channel extended until a specified date.

Reason: This alert is triggered when the support channel duration for a cluster is automatically extended beyond its original end date due to continued access by Cohesity Support Personnel.

Action: Review whether the support channel should remain open and is in active use.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 43200

SNMP OID: 1.3.6.1.4.1.47421.0.478

Platform: All

CE03601163 AppsModeEnabledDueToSubnetClashResolution

Alert Description: Apps mode has been enabled after a subnet conflict was resolved.

Reason: The configured applications subnet no longer conflicts with internal cluster networking.

Action: No action required. The subnet issue is resolved.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.495

Platform: All

CE03601166 AthenaSubnetFirstIPMatch

Alert Description: The subnet configuration of the Marketplace Apps Services is now valid.

Reason: The subnet CIDR configuration now matches the correct network boundary.

Action: No action required. This is an informational alert.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.498

Platform: All

CE03601169 ConfigureVipSucceeded

Alert Description: VIP configured successfully on the Cohesity cluster node.

Reason: This alert is triggered when a VIP configuration succeeds after a previous failure. The Cohesity cluster detected that the VIP is now correctly assigned to the node and raises this informational alert to clear the earlier failure state.

Action: No action needed. The VIP has been successfully configured.

Related Error Code: Not Applicable.

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.515

Platform: All

Configuration Alerts

CE03701025 ClusterServiceGflagsFound

Alert Description: The upgrade process found that cluster service Gflags are set on the cluster.

Reason: This alert is triggered when during an upgrade, the Cohesity cluster determines that cluster service gflags are set.

Action: No action is required.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 31536000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.180

Platform: All

CE03701044 UnusedGflags

Alert Description: The node has Gflags that are no longer in use.

Reason: The node is configured with deprecated Gflags that are no longer in use. This alert is triggered weekly until resolved.

Action: Remove the unused Gflags.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 31536000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.209

Platform: All

CE03701063 IpmiUserUpdateFailed

Alert Description: IPMI user update failed on the node.

Reason: Incorrect input was provided for creating the cluster or there is a networking issue in the current node. This alert is triggered when an IPMI user update fails on a node.

Action: Check and reconfigure the IPMI user.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 31536000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.39

Platform: All

CE03701074 IpmiConfigAbsent

Alert Description: IPMI config is absent on the cluster.

Reason: This alert is triggered when there is periodic monitoring of the system. The alert can be caused if IPMI is not configured on the physical cluster.

Action: Configure IPMI.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 7776000 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.40

Platform: All

CE03701076 SshTimeoutChanged

Alert Description: The SSH timeout configuration has changed on the node.

Reason: This alert is triggered when there is periodic monitoring of the system. The alert can be caused if the SSH timeout configuration has changed on the node.

Action: No action is needed.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.61

Platform: All

CE03701082 SshKeyMismatch

Alert Description: Check Failed, All nodes in the cluster must have the same ssh keys.

Reason: This alert is triggered when all the nodes on the Cohesity cluster do not have the same SSH keys.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.223

Platform: All

CE03701139 InactivityUITimeoutChanged

Alert Description: UI inactivity timeout changed.

Reason: This alert is triggered when the UI inactivity timeout is changed.

Action: Check if the configuration change is done by an authorized user.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.374

Platform: All

CE03701140 AbsoluteUITimeoutChanged

Alert Description: UI absolute timeout changed.

Reason: This alert is triggered when the UI absolute timeout is changed.

Action: Check if the configuration change is done by an authorized user.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.375

Platform: All

CE03701141 SSHLoginSuccess

Alert Description: SSH login is successful.

Reason: This alert is triggered when the `cohesity` user is able to log in successfully through SSH with `allow_ui_login_alerts` GFLAG enabled.

Action: Check if the login is done by an authorized user.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.376

Platform: All

Fault Tolerance Alerts

CE03915005 MetadataCannotTolerateAnymoreFailures

Alert Description: Metadata fault tolerance level has been reached. Autohealing from an additional failure will not be possible.

Reason: This alert is triggered when automatic healing of failed Nodes is stalled and any more failures would result in unavailability and possible data loss. This alert is triggered daily. The metadata server cannot tolerate any more failures because the automatic healing of failed nodes has stalled. Possible causes include an insufficient number of nodes online or a failed metadata server on available nodes.

Action: Bring any failed nodes back online. Look for any metadata server failed Alerts and resolve them. If the problem persists, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.115

Platform: All

CE03915006 MetadataCannotAchieveConfiguredFaultToleranceFactor

Alert Description: Metadata cannot achieve the configured fault tolerance factor.

Reason: This alert is triggered if a metadata fault tolerance factor change is stalled. The cluster does not have enough nodes for the metadata server to achieve the configured fault tolerance factor. This alert is triggered daily.

Action: Make sure there are enough nodes in the cluster to support the configured metadata fault tolerance factor. Look for any metadata server failed Alerts and resolve them.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.116

Platform: All

CE03902123 NodeFailureIsNotTolerated

Alert Description: The effective usage space in the cluster is high. The cluster may not have enough space to heal if a node failure happens.

Reason: This alert is triggered when the Cohesity cluster detects that a potential node failure could result in an effective Cohesity cluster disk usage of 99% or greater. This alert Cause field indicates the current actual disk usage, which is typically less than 99%. If a node in the Cohesity cluster fails, its contents need to be copied to the other nodes. If the Cohesity cluster detects that the contents of a potentially failed node cannot be copied to the remaining nodes due a lack of disk space, this alert is triggered to indicate that node failures cannot be tolerated.

Action: Do one or more of the following:

- Add more nodes to your Cohesity cluster.
- Clean up Snapshots that are no longer required.
- Archive Snapshots to an External Target.
- Configure Cloud Tier to hold rarely used data.
- Consider changing Storage Domains to use Erasure Coding. Capacity utilization will be reduced over time in the background.
- See the Alert: CE03902123 Cluster Space Usage High KB article.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.106

Platform: All

CE03901064 LowFaultTolerantStorageDomain

Alert Description: Low fault tolerant storage domain found for the cluster.

Reason: This alert is triggered when a low fault tolerant Storage Domain is found. The metadata failure tolerance factor for the cluster was increased to a value higher than the failure tolerance factor of the Storage Domain mentioned in the alert.

Action: Consider increasing the fault tolerance/EC factor of the Storage Domain to match the metadata failure tolerance for the cluster.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.83

Platform: All

CE03906004 NotEnoughNodes

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: More nodes need to be added.

Reason: This alert is triggered when the Cohesity cluster detects that more nodes must be added to the Cohesity cluster to support the current Erasure Code settings. This alert is triggered once per day.

Action: Add more nodes to the Cohesity cluster.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

CE03906009 DataUnavailable

Reason: This alert is triggered when the Cohesity cluster detects that some chunk files have been lost due to disk or node failures. This alert may be triggered daily.

Action: Contact Cohesity Support.

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: Not Applicable

Platform: All

Storage Usage Alerts

CE03815001 MetadataSizeExceedsThreshold

Alert Description: Metadata size on disk exceeding threshold.

Reason: This alert is triggered when the Cohesity cluster detects that the metadata usage on an SSD disk exceeds 95% of the total capacity of the disk. After the alert is triggered, if the metadata size continues to exceed 95% of disk capacity, this alert is triggered every 24 hours.

Action: See the ALERT: CE03815001 METADATASIZEEXCEEDSTHRESHOLD KB article.

- Review the Protection Policy and reduce the frequency of any full backups or reduce the Snapshot retention period.
- Delete old Snapshots or add more nodes to expand the cluster.
- If you are performing data dumps on Views (Oracle RMAN, Robocopy, SQL Dumps, Remote Adapter) and view backups, review how long such dumps are being retained.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 86400 seconds

SNMP OID: 156

Platform: All

CE03815007 MetadataCanRunOutOfSpaceUponFailure

Alert Description: Automatically healing from a single domain failure is no longer guaranteed due to space constraints.

Reason: This alert is triggered when the cluster metadata usage is high and there is not enough free metadata space to heal a potential failure.

Action: Increase the metadata capacity by adding more nodes or bringing any offline nodes or SSDs back online. Resolve any other metadata alerts.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.118

Platform: All

CE03802121 ClusterSpaceUsageHigh

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: Cluster is full.

Reason: This alert with a Critical severity is triggered when the Cohesity cluster is over 90% full. After the alert is initially raised, if the Cohesity cluster continues to be over 80% full, this alert is raised every 12 hours.

Action: Do the following:

- Add more nodes to your Cohesity cluster.
- Clean up Snapshots that are no longer required.
- Archive Snapshots to an External Target.
- Configure Cloud Tier to hold rarely used data.
- Consider changing Storage Domains to use Erasure Coding. Capacity utilization will be reduced over time in the background.
- See the Alert: CE03802121/CE03802122 Cluster Space Usage High KB article.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 43200 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.104

Platform: All

CE03802122 StorageDomainSpaceUsageHigh

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: Storage Domain is full.

Reason: This alert with a Critical severity is triggered when a Storage Domain is over 90% full. After the alert is initially raised, if the Storage Domain continues to be over 90% full, this alert is raised every hour.

Action: Do one or more of the following:

- Add more nodes to your Cohesity cluster.
- Clean up Snapshots that are no longer required.
- Archive Snapshots to an External Target.
- Configure Cloud Tier to hold rarely used data.
- Consider changing Storage Domains to use Erasure Coding. Capacity utilization will be reduced over time in the background.
- See the Alert: CE03802121/CE03802122 Cluster Space Usage High KB article.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.105

Platform: All

CE03801007 DiskSpaceLow

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The available space in the disk is low on the node.

Reason: This alert is triggered when the partitions (logs, root, data, boot) usage is higher than the informational level threshold value.

Action: If a resolution is not found for log partition usage, see Alert: CE03801007 LogsDiskSpaceLow KB article. For other partitions usage, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Informational

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.6

Platform: All

CE03801007 DiskSpaceLow

> **Note:** This alert is triggered with different severities based on the scenario.

Alert Description: The available space in the disk is low on the node.

Reason: This alert is triggered when the partitions (logs, root, data, boot) usage is higher than the critical level threshold value.

Action: If a resolution is not found for log partition usage, see Alert: CE03801007 LogsDiskSpaceLow KB article. For other partitions usage, contact Cohesity Support.

Related Error Code: Not Applicable

Severity: Critical

Dedup Interval: 3600 seconds

SNMP OID: 1.3.6.1.4.1.47421.0.6

Platform: All

CE03815008 MetadataDiskRanOutOfSpace

Alert Description: Scribe reservation failed.

Reason: This alert is triggered when the Metadata disk has insufficient disk space on the node.

Action: Contact Cohesity Support for assistance.

Related Error Code: Not Applicable

Severity: Warning

Dedup Interval: 86400 seconds

SNMP OID: N/A

Platform: All

Alert Not Documented or Link is Broken

If an alert is not documented or the link is broken, click **No** for the **Was this page helpful?** question in the **Feedback** section located at the top right of this topic. Provide the alert name, code, and severity so that we can add it to the documentation.

Reports

> **Important:** Helios Reporting supports global reporting and has an enhanced library of built-in reports. You can find equivalent reports in the new and improved Helios Reporting. Please create new schedules with the equivalent reports in Helios Reporting to start using these new features.

In today's data-driven business environment, it is more important than ever to analyze your data and identify patterns and trends that help organizations make better decisions. With reporting, Cohesity extends its built-in reporting features to give you the flexibility to produce several reports to understand cluster capacity, usage metrics and help you address your organization's specific needs most effectively.

For information about how long the cluster keeps audit data, see Data Retention for Audits.

Run Reports

1. Navigate to **Reporting** and click **Proceed to cluster Reporting**.
2. Select a report from the **Select Report** drop-down.
3. If the **Filters** panel is provided for the report, you can optionally filter your results by the options listed. Select the options and click **Apply Filters**.

Email Reports

You can email most reports as needed or on a recurring basis.

> **Note:** To email reports, an email address must be specified for the local System Admin account (admin) of the Cohesity cluster. Select **Settings > Access Management** and if necessary, edit the local admin user to provide an email address.

1. Navigate to **Reporting** and click **Proceed to cluster Reporting**.
2. Select a report from the **Select Report** drop-down.
3. If the **Filters** panel is provided for the report, you can optionally filter your results by the options listed. Select the options and click **Apply Filters**.
4. Click the **Email** icon ( ) and type one or more email addresses in the **Email Recipients** field. (Not all reports can be emailed.)
5. Select a format for the report in the **Format** drop-down. A CSV format report is sent as an email attachment. An HTML format report is sent in the body of the email.
6. To email the report one time, toggle **Schedule Recurring Email** off.

   To send recurring emails, select one or more days in **Email Every** and set the time the emails should be sent.
7. Optionally, select the **Time Zone** that matches your mailing schedule.
8. The current report filters are displayed. You can use them as is or modify them.
9. Click **Save Schedule**.

If you scheduled emails, click the **Scheduled Emails** tab to display all scheduled emails.

If you did not provide schedule information, the report is emailed immediately.

Delete a Scheduled Email

1. Navigate to **Reporting** and click **Proceed to cluster Reporting**.
2. Click the **Scheduled Emails** tab.

   This tab lists the scheduled emails for all reports.
3. To edit a scheduled email, click the edit icon ( ) and edit the recipients, format, filters and scheduling information and then click **Save Schedule**.
4. To permanently delete a scheduled email, click its delete icon ( ).

Download a Report as a CSV File

For most reports, you can download the data generated for the currently displayed report as described below. (You can also schedule an email that includes the report as a CSV

attachment as described previously in this topic.)

1. Navigate to **Reporting** and click **Proceed to cluster Reporting**.
2. Select a report from the **Select Report** drop-down.
3. If the **Filters** panel is provided for the report, you can optionally filter your results by the options listed. Select the options and click **Apply Filters**.
4. Click the **CSV** icon (↓). The file is downloaded immediately to your device.

Cohesity Reports

Cohesity provides the following reports:

Report Name	Description
Available Local Snapshots Per Job	This report provides Snapshot summary statistics for a leaf Protection Source Object (such as a VM) for all Protection Groups or a specific Protection Group. For detailed information about this report, see Available Local Snapshots.
Agent Status Summary	This report provides information about the Cohesity Agents installed on servers registered as Sources on the Cluster. For detailed information about this report, see Agent Status.
Backup Summary	This report shows an overview of the Protection Groups run on the Cohesity cluster for the specified filter criteria. For each Protection Group, this report shows how many times the Job has run, total bytes read and statistics about the last time the Job ran. For detailed information about this report, see Backup Summary.
Clone and Recover VM Tasks Summary	This report shows an overview of Clone and Recover tasks for the specified date range. For detailed information about this report, see Clone and Recover VM Tasks Summary.
Cluster Wide Storage	This report provides Cluster storage statistics that are useful for administrators to plan out disk space utilization. For detailed information about this report, see Cluster Wide Storage.
Data Transferred to External Targets	This report provides statistics about the data transferred to External Targets for archiving Snapshots created by Protection Groups, and moving cold data from the Cohesity cluster to Cloud Tier. For detailed information about this report, see Data Transferred to External Targets.

Report Name	Description
Failed Objects	This report lists all the Objects for which an error occurred during the last protection runs and prevented the Objects from being backed up. For detailed information about this report, see Failed Objects.
GDPR Object Summary	This report provides information about whether an Object is protected, who has access to it, where Object Snapshots are stored and how long they are retained, and any actions performed on the Object. For detailed information about this report, see GDPR Objects Summary.
Objects Protected by Multiple Protection Groups	This report lists all the Objects that have been backed up (protected) by more than one Protection Group. For detailed information about this report, see Objects Protected by Multiple Protection Groups.
Protected Objects Heatmap	This report provides an interactive grid of cells that indicate whether the Backup Tasks associated with Protection Group runs were successful (green cells) or encountered errors (red cells). For detailed information about this report, see Protected Objects Heatmap.
Protection Details Per Object	This report provides Protection Run (Snapshot) details for a leaf Protection Source Object (such as a VM). For detailed information about this report, see Protection Details Per Object.
Protection Groups Inventory and Schedule	This report lists all the Protection Groups on the Cohesity cluster matching the specified filter criteria. For each Protection Group, this report lists all the Objects that will be backed up by this group and the schedule for when the Objects will be backed up. For detailed information about this report, see Protection Groups Inventory and Schedule.
Protection Runs Summary	This report provides information about Protection Group runs and related replication and archival tasks. For detailed information about this report, see Protection Runs Summary.
Protection Summary by Object Type	This report shows the protection status of Objects as of the last protection runs that occurred in the specified Date Range. For detailed information about this report, see Protection Summary By Object Type.
Source Growth and Variance	This report shows how the size of protected Servers per registered Source are changing over time. For detailed information about this report, see Source Growth & Variance.

Report Name	Description
Sources Protected	This report provides overview statistics about the number of Servers protected or unprotected for registered Sources on the Cohesity cluster. For detailed information about this report, see Sources Protected.
Storage Consumed by Backups	This report provides statistics about the Protection Groups that are consuming the most disk space on the Cohesity cluster. For detailed information about this report, see Storage Consumed by Protection Groups.
Storage Consumed by File Categories	This report provides data statistics about File Categories in the latest Snapshots for all Servers stored on the Cohesity cluster. Statistics are only reported for files indexed by the Cohesity cluster. For detailed information about this report, see Storage Consumed by File Categories.
Storage Consumed by Servers	This report provides statistics about the Servers that are using the most disk space on the Cohesity cluster and Primary Storage. For detailed information about this report, see Storage Consumed by Servers.
Storage Consumed by Storage Domains	This report provides information about the Storage Domains that are storing the most logical data and raw data data on the Cohesity cluster. For detailed information about this report, see Storage Consumed by Storage Domains.
Storage Consumed by Organizations	This report provides information about the organizations consuming the most storage. For detailed information about this report, see Storage Consumed by Organizations.
Top Jobs	This report shows the top Protection Groups by the average amount of data transferred, the length of Protection Group run time, the number of Objects protected and the number of SLA violations. For detailed information about this report, see Top Protection Group.
Unprotected VMs	This report provides lists of VMs that are protected and not protected by Cohesity for the registered Sources on the Cohesity cluster. In addition, overview statistics about protected and unprotected VMs are provided. For detailed information about this report, see Unprotected/Protected VMs.
User Quota	This report provides information about the default user quota defined for a View and any overrides defined for individual users. For detailed information about this report, see User Quotas.

Understanding Storage Metrics

The Cohesity UI and reports use the following terminology to describe storage metrics:

Metrics	Description
Logical Data	The combined total of data in the objects that are protected by Cohesity.
Data In	Total unique logical data residing in Cohesity.
Data Written	Data written on the Cohesity after the unique logical data has been reduced by data deduplication and data compression.
Storage Consumed	The total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency). This data can be stored either all locally or as a combination of local storage and in the Cloud Tier (if Cloud Tier is enabled).
Resiliency Impact	The space required to honor resiliency settings (Erasure Coding / Replication Factor) in addition to the data written.
Storage Consumed Local and Cloud	Data stored locally and in the Cloud Tier (if Cloud Tier is configured) including the space required for honoring the resiliency settings (EC/RF).

Custom Reporting

The Custom Reporting feature addresses the limitations of monitoring options for historical data. This feature helps in creating custom reports to capture the information outside of the built-in reports available on the Cohesity cluster.

The Custom Reporting in Cohesity is achieved by reading the data from a Custom Database (Postgres) that stores relevant performance metrics, stats, and information of a (standalone) Cohesity cluster.

Note:

- In earlier releases, there was no firewall rule for custom reporting. After upgrading to the latest version, you may encounter connection failures to the Reporting Database because, starting from version 6.8, the new firewall rule is set to `deny` by default.
- To use custom reporting in a Cohesity cluster, you should edit the firewall rule for Reporting Database and set it to Allow. For more information, refer to the *Configure Firewall Rule* section in Manage Application-Based Firewall Rules.
- The IP address of the Reporting Database may change, as it can fail over to another node in the cluster during service restarts, patching, upgrades, or similar events.
- Contact your Cohesity account team to use the Custom Reporting feature.

The users can use the third-party SQL reporting tools, such as JasperSoft, TablePlus, Orange, DevArt plugin for Excel, Microsoft SSRS / Visual Studio, and Tableau for Desktop to access the Custom Database and build reports tailored to specific business requirements.

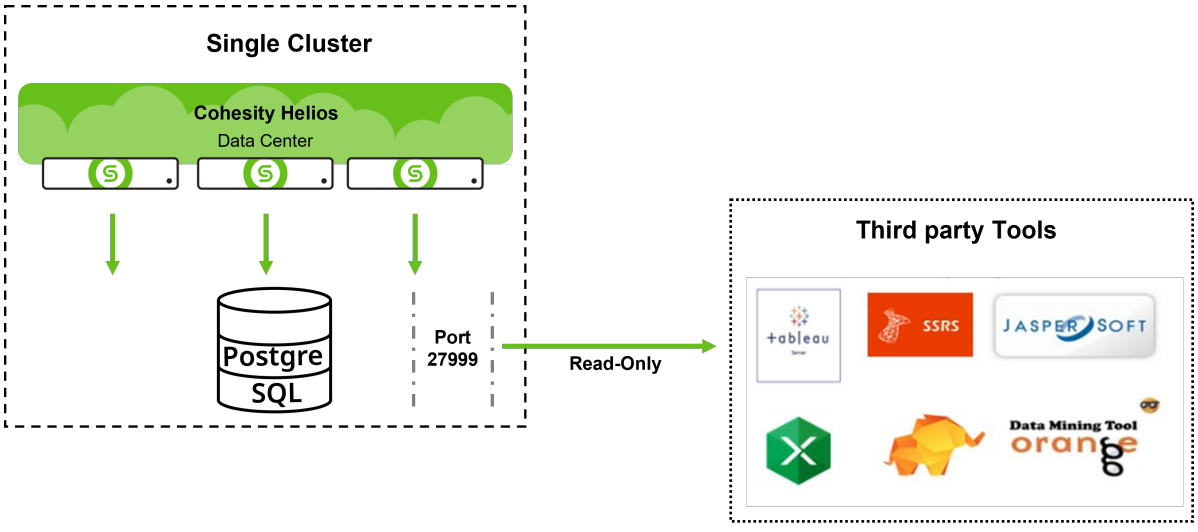

Custom Reporting Architecture

For more information, see the Cohesity Custom Reporting Solution Guide.

Limitations

- If you change the fault tolerance values of the Postgres database in the Cohesity Dashboard (**Settings > Summary > Configure**), the changes are not applied to the

Postgres database.

- Data archival is not supported for the Postgres database.
- Custom Reporting does not track expired snapshot data for Protection Group runs.

Connecting to the Postgres Database

To avoid a single point of failure, the data is replicated across the nodes in the Cohesity cluster. A service manager manages redundant Postgres databases inside the cluster. This redundancy is set during cluster configuration. The number of replicas depends on the fault tolerance of the cluster.

1. Connect to the Cohesity CLI. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

   > **Note:** The Cohesity Admin user should perform the steps listed in this section.

2. Execute the following command to get the default credentials to connect to the Postgres database.

   ```
   iris_cli custom-reporting database
   ```

   > **Note:** The default credentials returned by the command are those generated on the cluster, not the global default.

3. Connect to the Postgres database using any third-party SQL reporting tools with the following parameters.
   - IP: Node IP running Postgres
   - Port: 27999
   - Database: postgres
   - User: reporting_read
   - Password: default password from the previous step

     > **Note:** The default password of the reporting database is changed after you upgrade a Cohesity cluster to the 6.7 release. You should update the password of the reporting database in your database clients. The connection to the reporting database will be broken after the upgrade if the connection is established using the default password provided in the previous releases.

- Schema: reporting

  Cohesity recommends changing the password after connecting to the Postgres database.

4. To change the password, connect to the Postgres database as described below and execute the following query:

```
alter role reporting_read with encrypted password 'new_strong_pass_
here';
```

> **Note:** This changes the password for the Postgres instance on the current node.

Repeat the above 4 steps on all Postgres instances.

This changes the password for the Postgres instance on the current node. Repeat these steps on all Postgres instances.

Understanding the Database

The users can build reports by understanding the underlying database schema. The database schema is intuitive and user-friendly. The table names and the field names in the database are straightforward and self-explanatory. Click the link below to refer to the Entity Relationship diagram and the Data Dictionary for the Custom Reporting database.

- Entity Relationship Diagram and Data Dictionary

If you update your Cohesity version to 6.4 and later, the following changes are made to the Custom Reporting tables:

- `Snapshot_expiry_time_usecs` attribute added in the *protection_job_runs* table.
- `Logical_size_in_bytes` attribute added in the *leaf_entities* table.
- The status' of the job run has been updated and from Cohesity version 6.4 the *job_run_status* table will report more updated status as listed in the table below.

Existing job run status	Updated job run status
1, 'success' 2, 'in_progress' 3, 'error' 4, 'cancelled'	0, 'Accepted' 1, 'Running' 2, 'Canceling' 3, 'Canceled' 4, 'Success' 5, 'Warning 6, 'Failure' 7, 'OnHold' 8, 'Missed' (Not currently used)

- The status' of the *restore_status_types* table from Cohesity version 6.4 will report more updated status types as listed in the table below.

Existing restore status types	Updated restore status types
0, 'kReadyToSchedule' 1, 'kAdmitted' 2, 'kInProgress' 3, 'kFinished' 4, 'kProgressMonitorCreated' 5, 'kFinishingProgressMonitor' 6, 'kRetrievedFromArchive' 7, 'kCancelled' 8, 'kInternalViewCreated' 9, 'kZipFileRequested	0, 'Accepted' 1, 'Running' 2, 'Canceling' 3, 'Canceled' 4, 'Success' 5, 'Warning 6, 'Failure' 7, 'OnHold'

- `is_latest_attempt` attribute added to *protection_job_run_entities* table to track the latest attempt for entities (in case of failures).
- `timezone` attribute added to the *cluster* table.
- `restore_source_type_id` attribute added to *restored_objects* table to identify the source from which the restore task was performed.
- Added support for AD leaf entities and Protection Groups.

Agent Status

This report provides information about the Cohesity Agents installed on servers registered as Sources on the Cluster. An Agent is required to protect physical servers and some hypervisor environments.

You can filter the report results by changing the Agent health status, operating system and software version and then clicking **Apply Filters**. If you do not have the Agent Version, run the report without this filter, copy the desired version from the report results and paste it into this field.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The report contains the following columns.

Column Name	Description
Host IP	The IP address of the host on which the Agent is installed.
Host OS Type	The operating system running on the host.
Agent Health Status	The status of the Agent: • **Healthy**–The Agent is healthy and is able to contact the Cluster. • **Degraded**–The Agent is running but in a degraded state because of internal issues on the server. • **Unknown**–The Agent is unknown and the Cluster does not attempt to contact the Agent. • **Unreachable**–The Cluster cannot contact the Agent. The server might be powered off. • **Unhealthy**–An attempt to refresh the status of the Agent failed.
Software Version	The software version of the Agent.

Column Name	Description
Upgradability	Indicates the whether the Agent can or should be upgraded to the same software version as the Cluster: • **Upgradable**–The Agent is running a back-level version and can be upgraded. For information about upgrading, see Install and Manage Cohesity Agents. • **Current**–The Agent is running the same software version as of Cluster. No upgrade is needed. • **Unknown**–The Agent is running a non-production version. This is unlikely in a production environment, but possible if the Agent is running a pre-release version for testing purposes. • **NonUpgradableAgentisNewer**–When Agent is running a later version than the Cluster version. • **NonUpgradableInvalidVersion**–The Agent is running an invalid version. This is unlikely in a production environment, but possible if the Agent is running a pre-release version for testing purposes. • **NonUpgradableAgentisOld**–The Agent is running a version that is too old to upgrade to the Cluster version. If possible, install a newer version of the Agent on the server.
Last Upgrade Status	Indicates whether the last upgrade was successful or not. This is the message text from the last upgrade attempt.

Available Local Snapshots

This report provides Snapshot summary statistics for a leaf Protection Source Object (such as a VM) for all Protection Groups or a specific Protection Group.

In the Actions panel, you can also optionally filter your results by Date Range or Job Name. By default, Snapshot summary statistics are reported for all Jobs. In the **Job Name** field, you can optionally specify a Job name to limit the Snapshots summary statistics to a single Job. To apply the filters, click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

If data is found matching the specified filters, a Job table is displayed.

Protection Group Table

The columns for the Protection Group table are described below. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again.

Column Name	Description
Object	The leaf Protection Source Object (such as a VM). Drill down on this Object link to get details about the protection runs backing up this Object. (See the Object table below.)
Total Snapshots	The total number of Snapshots that backed up the leaf Protection Source Object based on the specified filter criteria
Last Snapshot	List the last Snapshot that backed up this Protection Source Object based on the specified filter criteria.
First Snapshot	List the first Snapshot that backed up this Protection Source Object based on the specified filter criteria
Error Snapshots	List the number of protection runs that failed when attempting to back up this Protection Source Object based on the specified filter criteria
Last Error Snapshot	List the last protection run that reported an error while attempting to back up this Protection Source Object based on the specified filter criteria

> **Tip:** To view the Snapshots summary statistics for all Reports again after viewing a single Job, clear the Job Name field and click **Apply Filters**.

Object Table

When you drill down on an Object in the Job table, an Object table displays providing details about the Snapshots backing up the Object that match the specified filter criteria. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again.

Column Name	Description
Start Time	The time this protection run started.
Protection Group Name	The name of the Protection Group.
End Time	The time this protection run ended.

Column Name	Description
Status	The current status of the protection run: • **Success**–Indicates that the protection run was successful. • **Running**–Indicates that the protection run is currently running. • **Warning**–Indicates that the protection run was successful but warnings were detected. • **Canceled**–Indicates that the protection run was canceled before completing. • **Error**–Indicates the protection run encountered an error and did not run to competition.
Schedule Type	The type of run specified in the Policy for the Job: • **Incremental**–When possible, utilizes Change Block Tracking (CBT) when capturing a Snapshot of this Object. • **Full**–Always captures a Full Snapshot of this Object and does not utilize Change Block Tracking (CBT). • **Log**–Captures database transaction logs to allow rolling back to a specific point in time.
Data Read	The total number of bytes read for this Object during this protection run.
Logical	The total number of logical bytes protected for this Object during this protection run. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication.
Copy Task Status	The status of copy tasks for this protection run such as replication and archival tasks. These tasks copy Snapshots to a Remote Cohesity Cluster or External Targets.

Backup Summary

This report shows an overview of the Protection Groups run on the Cohesity cluster for the specified filter criteria. For each Protection Group, this report shows how many times the Job has run, total bytes read and statistics about the last time the Job ran.

For information about how long the cluster keeps audit data, see Data Retention for Audits.

This report contains the following charts:

- **Overview of Last Run Statuses Doughnut Chart**—This chart shows only the last run status for all the Protection Groups in the cluster that meet the specified filter criteria. For example if the first run of a Protection Group fails with an error but then is successful, the chart shows the status of success for this Protection Group.
- **Protection Runs Hourly Bar Chart**—This chart provides a status snapshot of all the Protection Group runs in the last twenty-four hours that meet the specified filter criteria. A single Protection Group can be represented multiple times in this chart. For

example, if one runs four times in the last twenty-four hours, each run is represented by a one-run bar in this chart.

In the Actions panel, you can also optionally filter your results by Data Range, Job Name, Registered Sources or Last Run Status of the Protection Group. You can specify a Date Range, enter part of a Job name, select one or more Sources from the Registered Source drop-down list and only report on Jobs that the last run status is successful. To apply the filters, click **Apply Filters**. You cannot filter on Storage Domain Sources.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The columns for the table at the bottom of the Protection Groups Summary page are described below.

Column Name	Description / Action
Job	Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.
Source	Drill down on this link to get details and protection statistics about this Source on the Cohesity cluster. If you drill down on a VMware Source, the reported VM statistics reflect all VMs in the Source and is not specific to the current Protection Group.
Runs	For the specified filter criteria, the number of times this Protection Group has run.
Last Run Success/Error	This column lists how many Objects were successfully backed up and how many failed during the last Run of this Protection Group based on the specified filter criteria. Drill down on this link to see Backup Run Details for this Job. A list of protection runs for this Job are listed. If this is a Server Protection Group, you can also drill down on the protection run to get a listing of the Servers that were backed up.
Data Read (total)	This column reports the cumulative number of bytes read from the Source for all the Runs of this Protection Group that meets the specified filter criteria.

Column Name	Description / Action
SLA Violation	For the last Protection Group Run based on the specified filter criteria, this column reports the SLA status. • **Pass**–The last run of this Protection Group completed within the SLA time period specified for this Protection Group. • **Fail**–The last run of this Protection Group did not complete within the SLA time period specified for this Protection Group. • **n/a**–Pass/Fail statistics are not available. The last Run of this Protection Group is still running or did not complete. The service-level agreement (SLA) defines how long the administrator expects protection run to take.
Bandwidth	The average amount of data transferred over the network per second per backup job.
Last Run Status	For the last Protection Group Run based on the specified filter criteria, this column reports the current status: • **Succeeded**–The protection group run completed without error. • **Canceled**–A protection group run started and then was canceled in the Cohesity Dashboard The protection group run did not complete. • **Canceling**–In the Cohesity Dashboard, the cancel option was selected for the protection group run and the cluster is in the process of canceling the run but the run is not yet canceled. • **Failed**–The protection group run did not complete due to an error. • 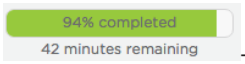 –The protection group run is currently running but is not completed yet.
Last Run Replication/Archival Status	For the last Protection Group Run based on the specified filter criteria, this column indicates the status of any replication ( ) and archival ( ) tasks. Mouse over the icon to see the replication or archival target, or click the icon to display additional details about the task.

Clone and Recover VM Tasks Summary

This report shows an overview of Clone and Recover tasks for the specified date range.

This report contains the following charts:

- **Recover Operations**—This doughnut chart shows how many VM Recover and Clone tasks have run in the specified date range.

- **Latest Operations**—In the specified date range, this bar chart represents the run time for the five last VM Recover and VM Clone tasks.

In the Actions panel, you can filter the results by Date Range. Select the range and click **Apply Filters**.

In the Search field above the table, you can filter the results by entering part of the Recover Task Name and Return.

The table lists all the cloning VM and recovering VM tasks in the specified date range. In the table, drill-down on the Task Name to find out details about the Task.

Cluster Wide Storage

This report provides Cluster storage statistics that are useful for administrators to plan out disk space utilization.

This page provides the following data use statistics about the entire local Cohesity cluster. The disk space available from External Targets is not included in the reported statistics. For example, the value reported in the Cluster Storage Capacity field does not include the disk space available in the Cloud Tier.

Field	Field Description
Cluster Storage Capacity	The amount of total disk space on the Cohesity cluster available for storage of your data such as Snapshots of Protection Groups, files and directories. **Note:** • Storage Domain settings such as number of failures tolerated and erasure coding affect how much data can be stored on the Cluster. • On All Flash clusters, Cohesity reserves extra space for metadata. This space is automatically made available for data (cluster capacity) if the cluster usage exceeds 90%. Due to this behavior, cluster capacity would vary when the cluster usage is 90% or exceeds 90%.
Cluster Storage Used	The amount of disk space on the Cohesity cluster currently used for storing your data such as Snapshots of Protection Groups, files and directories.
Daily Growth Rate	Daily growth rate for the **Storage Consumed** value. Where ***Storage Consumed*** is the total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for.

Field	Field Description
Data Reduction Ratio	The data reduction ratio is calculated by dividing the Logical Data Managed by the Physical Data Used. **Logical Data Managed**–Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication. **Physical Data Used**–The size of the data physically stored on the Cohesity cluster, after the data has been reduced by change block tracking, compression and deduplication, including the disk space used for data resiliency (fault tolerance and erasure coding).
Utilization Ratio	Ratio of **Storage Consumed** to **Raw Cluster Capacity**. Where ***Storage Consumed*** is the total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for.
Date to Exceed 80% Utilization	Based on the current daily growth rate, what is the projected date that the Cohesity cluster is expected to be 80% full.
Raw Cluster Capacity	The total physical capacity of all storage devices. This is pre-reservation of capacity for Cohesity software operations.
Minimum Usable	The total physical capacity of all storage devices after subtracting the space reserved for Cohesity software operations.

Below the list of statistics is the following chart:

Cluster Storage Used Over 7 Days—This chart graphs the physical growth of the disk space used on the Cohesity cluster over the last seven days.

You can optionally Filter your results by **Date Range**. Specify a date range to check the physical growth of the disk space for the specified dates.

Data Transferred to External Targets

This report provides statistics about the data transferred to External Targets for archiving Snapshots created by Protection Groups, and moving cold data from the Cohesity cluster to Cloud Tier.

You can filter the report by date range, target type and target name.

The report provides the By Target and By Protection Group tabs as described below.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

By Target

This tab provides a table of statistics for each External Target.

Column Name	Description
Target Name	The name of the External Target.
Protection Groups	The number of Protection Groups that transferred data to the External Target.
Logical Data Transferred	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type. • **VMs**: The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. • **All Other Workloads**: The data size reported is the actual front end data residing on the server. If a server with 1TB capacity contains 100GB of data, the server reports 100GB. **Note:** - These metrics do not include unprotected objects. - Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000GB (1TB).
Physical Data Transferred	Data written to the External Target after the unique logical data has been reduced by data deduplication and data compression.
Total Storage Consumed	The total size of data stored on the External Target after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for. **Note:** The Total Storage Consumed on an external target may vary or be unavailable for collection when using an S3 bucket.

By Protection Group

This tab provides a table of the Projection Job statistics for each External Target.

> **Note:** This tab does not reflect data that has been moved from the Cohesity cluster to theCloud Tier.

Column Name	Description
Protection Group	The name of the Protection Group.
Logical Data Transferred	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type. • **VMs**: The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. • **All Other Workloads**: The data size reported is the actual front end data residing on the server. If a server with 1TB capacity contains 100GB of data, the server reports 100GB. **Note:** - These metrics do not include unprotected objects. - Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000GB (1TB).
Physical Data Transferred	Data written to the External Target after the unique logical data has been reduced by data deduplication and data compression.
Storage Consumed	The total size of data stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for.

Failed Objects

This report lists all the Objects for which an error occurred during the last protection runs and prevented the Objects from being backed up. Objects are entities such as VMs and Views that are protected by Protection Groups.

For information about how long the cluster keeps audit data, see Data Retention for Audits.

In the Actions panel, you can optionally filter the report results. Specify a Date Range, part of a Job Name, Object Type and Registered Source, and click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The summary table at the top of the report contains the following columns.

Column Name	Description
Type	The Object type.
Sources	The number of registered Sources for the Object type.
Total Objects	The total number of Objects that were selected to be backed up by the last protection runs.
Failed Objects	The number of Objects for which an error occurred during the last protection runs and prevented the Objects from being backed up.

Detailed tables are provided for each Object type and contain the following columns.

Column Name	Description
Object	The name of the Object, such as VM, View or Volume.
Failed Backup	The date and time of the protection run that encountered an error.
Run Type	The type of run specified in the Policy for the Job: • **Incremental**–When possible, utilizes Change Block Tracking (CBT) when capturing a Snapshot of this Object. • **Full**–Always captures a Full Snapshot of this Object and does not utilize Change Block Tracking (CBT). • **Log**–Captures database transaction logs to allow rolling back to a specific point in time.
Message	The message that indicates why the last protection run failed to back up this Object.

GDPR Objects Summary

This report provides information about whether an Object is protected, who has access to it, where Object Snapshots are stored and how long they are retained, and any actions

performed on the Object. Objects are entities such as VMs, physical servers, Views and storage volumes. You would typically run this report for Objects that contain Personally Identifiable Information (PII) as defined in the General Data Protection Regulation (GDPR).

Begin typing an Object name in the **Object Name** field and then select one or more Objects from the list.

You can filter the report results by changing the date range for the audit logs and then click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The report is organized by Objects and provides the following information. Mouse over the Object's icon to display its description.

Column Name	Description
Source	The name of the Source that hosts the Object.
Accessible By	The local users and Active Directory users or groups who have access to the Object based on their Cohesity role.
Protection Status	The name of the Protection Groups, if any, that back up the Source.
Data Locations	**Location:** Where Snapshots (backups) of the Source are located, including the local Cohesity cluster and any replication, archive and cloud tier targets. **Snapshots:** The number of Snapshots stored at each location. **Retention:** The retention period for the Snapshots as specified by the Policy associated with the Protection Group. **Expiration:** How long the Snapshot is available before it is deleted (based on the retention period). The **Data Locations** section of the report is available only if the Source is protected by a Protection Group.
Audit Logs	Shows the actions performed on the Source during the time period specified in the **Date Range** filter field. **Date:** When the action was performed. **User:** Who performed the action and whether the user is a local or Active Directory user. **Action:** Which action was performed, such as `ran a Protection Group`, `created a view` or `started recover task`.

Objects Protected by Multiple Protection Groups

This report lists all the Objects that have been backed up (protected) by more than one Protection Group. Objects are entities such as VMs and Views that are protected by Protection Groups.

A table is provided for each Object that has been backed up by multiple Protection Groups that match the specified search criteria. In the sub-heading of the table, a link to the protected Source is provided. Drill down on this link to get details and protection statistics about this Source on the Cohesity cluster. If you drill down on a VMware Source, the reported VM statistics reflect all VMs in the Source and is not specific to the current Protection Group.

In the table, a link to each of the Protection Groups is provided. Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.

Using the top right **Sort results by...** drop-down, you can sort by the Object tables by the following criteria:

- **Sort by Object Name**
- **Sort by Projection Job Count**
- **Sort by Source**

By default, this page reports on all the Protection Groups for all the Sources in the Cohesity cluster. In the Actions panel, you can limit the results to report on the Protection Groups for specific Sources. In the **Registered Source** drop-down list, select one or more Sources and click **Apply Filters**. You cannot filter on Storage Domain Sources.

Protected Objects Heatmap

The Protected Objects Heatmap Report displays the status of the backup tasks of Objects that are associated with a Protection Group. This report lists the protected Objects.

The report provides an interactive grid of cells. By default, the report shows the backup tasks for the last seven days. Each cell indicates whether the backup task runs were successful (green cells) or encountered errors (red cells). The Red cells indicate both the failed and missed job runs. Gray cells indicate a period for which there were no Job runs.

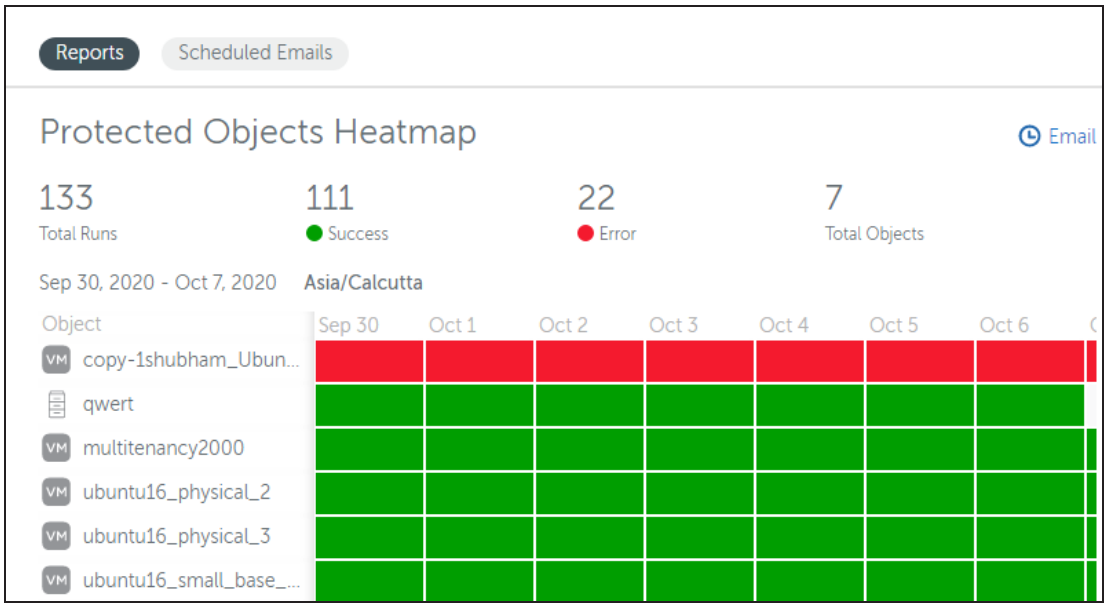

To find the root cause of an error, click a red cell to drill down to a list of Protection Group runs, then click the **Start Date** to drill down to job run details.

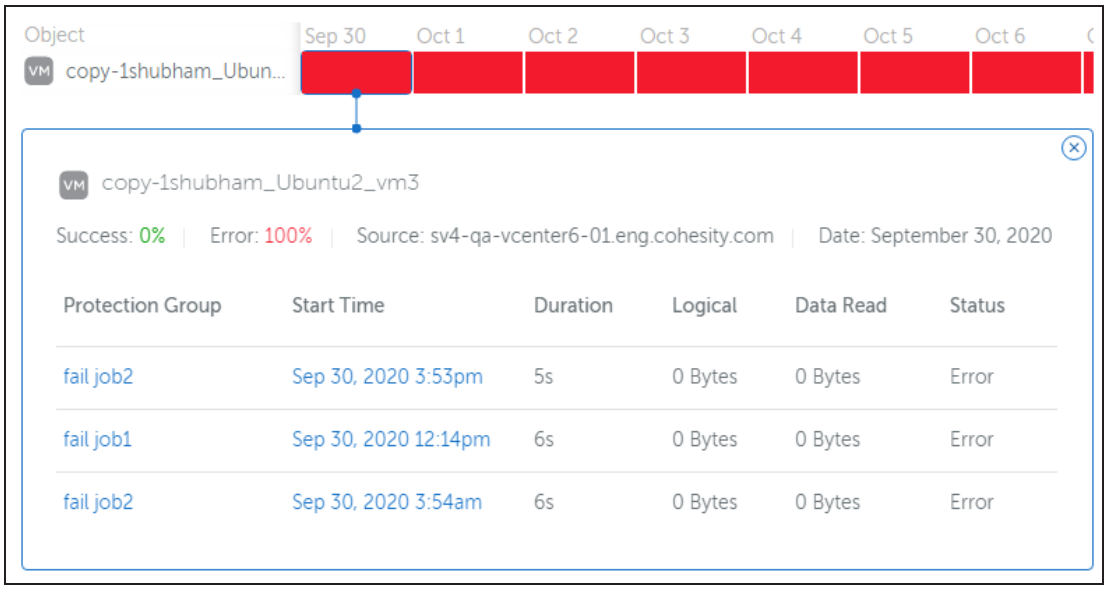

The **Message** column displays the error.

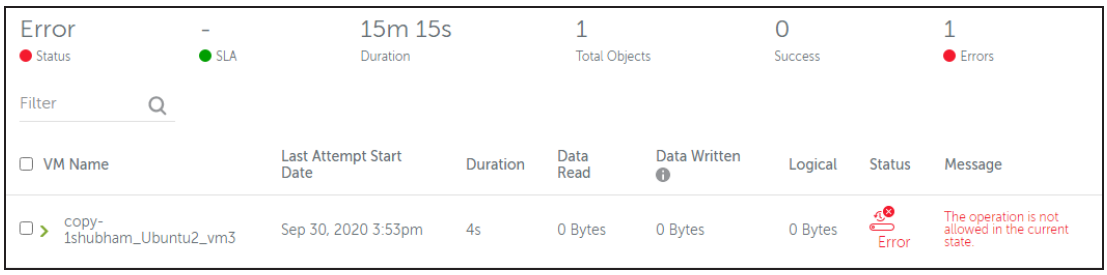

The right pane allows you to filter the report results. Use Date Range, Aggregate By, Time Zone, Object Type, Registered Source, Protection Group Name, Object Name, and Organizations to filter the report based on specific criteria. Click **Apply Filters** to view the report based on the specified criteria.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

> **Note:** The Protected Heatmap Report only allows to send an email in HTML format.

The report does not reflect the status of other Protection Group run tasks, such as replication or archive tasks. The Protection Runs Summary provides some of that information for Protection Groups.

To view another report, in the right-pane, from the **Select Report** drop-down, select the required report.

Protection Details Per Object

This report provides Protection Run (Snapshot) details for a leaf Protection Source Object (such as a VM).

For information about how long the cluster keeps audit data, see Data Retention for Audits.

To filter on Objects (such as VMs), in the **Object Name** field enter one or more leaf Protection Source Object names.

In the Actions panel, you can also optionally filter your results by Date Range or Run Status of the protection run. The **All Jobs** option clears the other run status options.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The columns for the table of this report are described below. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again.

Column Name	Description
Start Time	The time this protection run started.
Protection Group Name	The name of the Protection Group.

Column Name	Description
End Time	The time this protection run ended.
Status	The current status of the protection run: • **Success**–Indicates that the protection run was successful. • **Running**–Indicates that the protection run is currently running. • **Warning**–Indicates that the protection run was successful but warnings were detected. • **Canceled**–Indicates that the protection run was canceled before completing. • **Error**–Indicates the protection run encountered an error and did not run to competition.
Schedule Type	The type of run specified in the Policy for the Job: • **Incremental**–When possible, utilizes Change Block Tracking (CBT) when capturing a Snapshot of this Object. • **Full**–Always captures a Full Snapshot of this Object and does not utilize Change Block Tracking (CBT). • **Log**–Captures database transaction logs to allow rolling back to a specific point in time.
Data Read	The total number of bytes read for this Object during this protection run.
Logical	The total number of logical bytes protected for this Object during this protection run. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication.
Copy Task Status	The status of copy tasks for this protection run such as replication and archival tasks. These tasks copy Snapshots to a Remote Cohesity Cluster or External Targets.

Protection Groups Inventory and Schedule

This report lists all the Protection Groups on the Cohesity cluster matching the specified filter criteria. For each Protection Group, this report lists all the Objects that will be backed up by this group and the schedule for when the Objects will be backed up. Objects are entities such as VMs and Views that are protected by Protection Groups.

A table is provided for each Protection Group and all the Objects that will be protected by this Protection Group are listed that match the specified filter criteria. The title for the table is the name of the Protection Group and a link. Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.

In the sub-heading of the table, a link to the Source is provided. Drill-down on the link to get more information about the Source:

- **Source is a hypervisor**Drill-down on this link to get details and protection statistics about this registered Source on the Cohesity cluster. The VM statistics reported in the drill-down page reflect all VMs that use the Source and is not specific to the current Protection Group. In the table, a link to each of the VMs is provided. Drill down on this link to get statistics about the files stored in this Server.
- **Source is a Storage Domain**—Drill down on this link to get storage statistics about the Storage Domain.

Using the top right **Sort results by...** drop-down, you can sort by the Job tables by the following criteria:

- **Sort by Job Name**
- **Sort by Source**

In the Actions panel, you can also optionally filter your results by Job or VM name. You can enter part of a Job name and/or VM name and click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

Protection Runs Summary

This report provides information about Protection Group runs and related replication and archival tasks. If the protection run has multiple replication or archival targets configured, only the first of each is displayed in the report.

You can filter the report results by changing the date range and then clicking **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The report contains the following columns.

Column Name	Description
Protection Group Name	The name of the Protection Group.
Start Time	The date and time the Protection Group started running.

Column Name	Description
Duration	The amount of time the Protection Group run required.
Bytes Read	The size of the set of protected objects as read by Cohesity for a single backup run. This number is a per protection run statistic and is not additive across backup runs.
Status	Whether the protection run was successful or not.
Replication	If snapshots are replicated for this protection run, this column indicates the name of the remote replication cluster. Otherwise, the column is empty. You can drill down on the cluster name for details.
Replication Status	Indicates whether the replication was successful or not.
Replication Duration	The amount of time the replication required.
Bytes Transferred	Total unique logical data received by Cohesity. For Data Protection, this is the sum of your first full backup and subsequent change data captured by Cohesity.
Archival	If snapshots are archived for this Protection Group, this column indicates the name of the archival destination. Otherwise, the column is empty. You can drill down on the archival destination name for details.
Archival Status	Indicates whether the archival was successful or not.
Archival Duration	The amount of time the archival required.
Bytes Transferred	The number of bytes transferred to the archival destination.

Protection Summary By Object Type

This report shows the protection status of Objects as of the last protection runs that occurred in the specified Date Range. Objects are entities such as VMs and Views that are protected by Protection Groups.

For information about how long the cluster keeps audit data, see Data Retention for Audits.

In the Actions panel, you can optionally filter the report results. Specify a Date Range, part of a Job Name, Object Type and Registered Source, and click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The summary table at the top of the report contains the following columns.

Column Name	Description
Type	The Object types displayed in the report.
Sources	The number of registered Sources for the Object type.
Objects	The number of protected Objects in the source.
Success	The number of Objects that were successfully backed up by the last protection runs that occurred during the specified Date Range.
Error	The number of Objects for which an error occurred (and the Object was not backed up) during the last protection runs that occurred during the specified Date Range.
Running	The number of Protection Groups currently running.
Data Read	The total number of bytes read for this Object by the last protection runs that occurred during the specified Date Range.
Logical	The total number of logical bytes protected for this Object by the last protection runs that occurred during the specified Date Range. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication.

The Object type tables contain the following columns.

Column Name	Description
Object	The name of the Object, such as VM, View or Volume.

Column Name	Description
Last Run Status	The status of the last protection run: • **Success**–Indicates that the protection run was successful. • **Running**–Indicates that the protection run is currently running. • **Warning**–Indicates that the protection run was successful but warnings were detected. • **Canceled**–Indicates that the protection run was canceled before completing. • **Error**–Indicates the protection run encountered an error and did not run to completion.
Start Time	The time this protection run started.
End Time	The time this protection run ended.
Run Type	The type of run specified in the Policy for the Job: • **Incremental**–When possible, utilizes Change Block Tracking (CBT) when capturing a Snapshot of this Object. • **Full**–Always captures a Full Snapshot of this Object and does not utilize Change Block Tracking (CBT). • **Log**–Captures database transaction logs to allow rolling back to a specific point in time.
Data Read	The total number of bytes read for this Object during this protection run.
Logical	The total number of logical bytes protected for this Object by protection runs that occurred during the specified Date Range. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication.

Source Growth & Variance

This report shows how the size of protected Servers per registered Source are changing over time.

This report provides the following chart:

- **Sources Growth**—This chart shows how the size of the protected Servers per Source is changing over the specified number of days. The size of a Server is reported if the Server has been selected to be backed up (protected).

In the **Calculate Variance Over Specified Days** field of the Actions panel, you can specify the number of days used to calculate the following:

- The number of days to show the growth in the **Source Growth** chart.
- Calculate the **# Objects Backed Up (<Num> days ago)** column in the table
- Calculate the **Unique on Primary (<Num> Days Ago)** column in the table.
- Calculate the **Variance of Unique on Primary** column in the table.

Enter a number of days (from 1 to 180) and click **Apply Filters**.

In the Actions panel, you can also optionally filter your results by Source. In the **Registered Source** drop-down list, select one or more Sources and click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The columns for the table at the bottom of this report are described below. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Primary Physical** column.

Column Name	Description / Action
Source	Drill-down on this link to get details and protection statistics about this registered Source on the Cohesity cluster. **Note:** The VM statistics reported in the drill-down page reflect all VMs that use the Source and is not specific to the current Protection Group.
# Objects Backed Up	As of today, the number of Objects that are successfully backed up by the Cohesity cluster for the current Source.
# Objects Backed Up (<Num> Days ago)	As of a past date, the number of VMs that were successfully backed up by the Cohesity cluster for the current Source. The past date is calculated from the current date and time minus the number of days specified in the **Calculate Variance Over Specified Days** field in the Actions panel.
Logical	As of today, the total amount of logical data on the protected Servers per Source. Indicates the size of the data if the data is fully hydrated or expanded.

Column Name	Description / Action
Unique on Primary	As of today, the total amount of physical storage required to store all the protected VMs on the primary storage device. (This is not the secondary storage consumed on the Cohesity cluster.) For VMware VMs, this metric is provided by the VMware API and corresponds to the `uniqueSize` property as defined in the VMware documentation.
Unique on Primary (<Num> Days Ago)	The same as **Unique on Primary**, but as of a past date. The past date is calculated from the current date and time minus the number of days specified in the **Calculate Variance Over Specified Days** field in the Actions panel.
Variance of Unique on Primary	This field shows the percentage difference between the two **Unique on Primary** fields defined above.The past date is calculated from the current date and time minus the number of days specified in the **Calculate Variance Over Specified Days** field in the Actions panel. The Variance is calculated using the following formula: `( (Unique on Primary Data Used Today - Unique on Primary Data Used on a Past Date) / Unique on Primary Data Used on a Past Date) * 100`

Sources Protected

This report provides overview statistics about the number of Servers protected or unprotected for registered Sources on the Cohesity cluster.

This report contains the following chart:

- **Total Protection Across Sources**—This doughnut chart shows the protection level of all Sources registered in the Cohesity cluster.

This report contains the following Table:

- **Protection by Source**—This table provides statistics about all Sources in the Cohesity cluster as described below.

Column Name	Description / Action
Source	Drill-down on this link to get details and protection statistics about this registered Source on the Cohesity cluster. > **Note:** The VM statistics reported in the drill-down page reflect all VMs that use the Source and is not specific to the current Protection Group.

Column Name	Description / Action
Protected Servers	Lists the number of protected Servers in this Source that are backed up by Jobs.
Unprotected Servers	Lists the number of unprotected Servers in this Source that are not backed by Jobs.
Protection Groups	Lists the number of Protection Groups that back up this Source.
Registered	Date and time that the Source was registered on the Cohesity cluster.
Refreshed	Date and time that the Object hierarchy for the Source has been refreshed on the Cohesity cluster. The Cohesity cluster tracks all the objects (for example VMs, resource pools, datacenters, folders, and ESXi hosts) found within a source in an object hierarchy. This object hierarchy is automatically refreshed every four hours. If a new object such as a VM is added to the source, the Cohesity Dashboard does not list the new VM until the object hierarchy is refreshed, either manually or automatically after four hours. To manually force a refresh of the Object hierarchy, select **Data Protection > Sources** , find the Source in the list and click .

Storage Consumed by Protection Groups

This report provides statistics about the Protection Groups that are consuming the most disk space on the Cohesity cluster. This report provides statistics for all types of Objects. Objects are entities such as VMs and Views that are protected by Protection Groups.

This report provides the following charts:

- **Top <Num> Protection Groups by Logical Data Managed**—This chart lists up to ten Protection Groups that match the specified filter criteria and have stored the largest amount of logical data from a Source. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication.

In the **Calculate Variance Over Specified Days** field of the Actions panel, you can specify the number of days used to calculate the **Logical Managed (At <Num> Days Ago)** and **Variance of Logical Managed** columns in the table. Enter a value from 1 to 180 and click **Apply Filters**. In the Actions panel, you can also optionally filter your results by Registered Sources or Storage Domains. Select a Storage Domain and Sources to filter by and click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The columns for the table at the bottom of this report are described below. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Total Physical Written** column from smallest to largest.

Column Name	Description / Action
Protection Group	Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.
Total Runs	The total number of times this Protection Group has run based on the specified search criteria.
Objects	The total number of Objects backed by all the Runs of this Protection Group based on the specified search criteria. For example if a Protection Group backs up 3 VMs and runs twice, a total of 6 Objects is reported in this column.
Storage Domain	The Storage Domain that is used to store the data backed up by this Job.
Logical (total)	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. These metrics are different depending on workload type. • **VMs**: The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. • **All Other Workloads**: The data size reported is the actual front end data residing on the server. If a server with 1TB capacity contains 100GB of data, the server reports 100GB. **Note:** - These metrics do not include unprotected objects. - Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the Logical Data to be 1000GB (1TB).
Logical (total at <Num> days ago)	As of a past date, the total amount of logical data stored on the Cluster by all the Runs of this Protection Group that matches the specified filter criteria. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication. The past date is calculated from the current date and time minus the number of days specified in the **Calculate Variance Over Specified Days** field in the Actions panel.

Column Name	Description / Action
Variance of Logical Managed	The percentage difference between the values of Logical (total) today, and Logical (total) x days ago, where x is a user input value.
Data In	Total unique logical data residing on the Cohesity cluster. This is the sum of your first full backup and subsequent change data captured by Cohesity as measured in units of source-side bytes.
Data Written	Data written and retained on the Cohesity cluster, as measured in post-dedupe and post-compression bytes.
Storage Consumed	The total size of data stored on the Cohesity file system, as measured in bytes after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for. This data can be stored either all locally or as a combination of local storage and in the Cloud Tier (if Cloud Tier is enabled). For example, if the Data Written is 100GB and resiliency is set to RF3, leading to the need for an additional 200GB of space, the resulting Storage Consumed is 300GB. If the Data Written is 100GB and resiliency is set to EC4:2, leading to the need for an additional 50GB (as for every 4 data fragments, 2 parity fragments are added), the resulting Storage Consumed is 150GB. > **Note:** If you have enabled Cloud Tiering, then Storage Consumed is the sum of space consumed on local storage and space consumed on the cloud tier. In such cases, the sum of Storage Consumed value and Capacity Available value for the cluster can exceed the cluster capacity.
Storage Consumed (<Num> days ago)	The total size of data stored on the Cohesity file system, as measured in bytes after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for over the number of days.
Growth Rate	Shown as a percentage, the increase in “storage consumed” over the time period specified in the report’s filter.
Daily Growth Rate	Daily growth rate for the Storage Consumed value over the time window specified in the report’s filter. Where Storage Consumed is the total size of data as measured in bytes stored on the Cohesity file system after all operations on the data (data deduplication, data compression, and data resiliency) have been accounted for.
Storage Reduction	The ratio of Logical Data to Storage Consumed.

Storage Consumed by File Categories

This report provides data statistics about File Categories in the latest Snapshots for all Servers stored on the Cohesity cluster. Statistics are only reported for files indexed by the Cohesity cluster.

> **Note:** The report only provides data statistics about File Categories in the latest Snapshots for all Servers and not other Objects like Views.

A file in a Snapshot is not indexed due to one of the following reasons:

- In the Protection Group that creates the Snapshot, the directory where the file is located was not selected to be indexed.

To get File Category statistics for a specific Server, go to the Storage Consumed by Servers and drill-down on the Servers in the Protected Object column.

This report provides the following chart:

- **Top Categories by Size in Cluster**—This chart lists the collections of files of similar types that store the largest amount of logical space for the last Snapshots of all the Servers on the Cohesity cluster. For example, the **image** file category reports all the files that end with image file extensions such as `*.png`, `*.tif`, `*.jpg`, `*.gif` etc.

In the **Calculate Variance Over Specified Days** field of the Actions panel, you can specify the number of days used to calculate the **Data Size (total, <Num> Days Ago)** and the **Variance** columns in the table. Enter a value from 1 to 180 and click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The columns for the table at the bottom of this report are described below. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Total Logical Data Written** column.

Column Name	Description
File Category	The name of the file category. For example, the **image** file category reports all the files that end with image file extensions such as `*.png`, `*.tif`, `*.jpg`, `*.gif` etc.
Number of Files	The number of files of the current category found in the latest Snapshots for all the Servers in the Cohesity cluster. Statistics are only reported for files indexed by the Cohesity cluster. See notes above.

Column Name	Description
Data Size (total)	As of today, the total size of the data, prior to deduplication and compression, for all the files in this Category found in the latest Snapshots for all the Servers in the Cohesity cluster. Statistics are only reported for files indexed by the Cohesity cluster.
Data Size (total, <Num> Days Ago)	The same as **Data Size (total)**, but as of a past date. The past date is calculated from the current date and time minus the number of days specified in the **Calculate Variance Over Specified Days** field in the Actions panel.
Variance	This field shows the difference between the two **Data Size** fields defined above. Statistics are only reported for files indexed by the Cohesity cluster. The Variance is calculated using the following formula: `( (Data Size as of Today - Data Size on a Past Date)/ Data Size on a Past Date) * 100`

Storage Consumed by Organizations

This report provides information about the organizations consuming the most storage. For more information about organizations, see Manage Organizations (Tenants).

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

This report consists of the following two tabs:

- Physical Usage
- Workload

Physical Usage Tab

The **Physical Usage** tab is displayed only in the Service Provider login. You must use the **Storage Consumed by Organizations** report available in the Service Provider login for billing purposes.

> **Note:** The data displayed on the **Physical Usage** tab is refreshed between 2 to 7 days.

By default, the report displays all storage domains and organizations in the cluster, but you can filter the data by:

- Storage domain(s)
- Organization(s)

The **Physical Usage** tab consists of the following two charts:

- Top organizations by Storage Consumed
- Top organizations by Storage Reduction Ratio

The following table describes the fields displayed on the **Physical Usage** tab:

Column Name	Description
Organization Name	The name specified for the organization when added to the cluster.
Storage Domain	The name of the storage domain assigned to the organization.
Logical Data	Data residing on the primary object.
Data In	Data sent from primary objects to Cohesity.
Data Written	Data to be written post-deduplication.
Storage Consumed	Data written with all copies required for honoring the resiliency setting.
Dedupe Ratio	Part of the data-reduction process that eliminates redundant copies of data to reduce storage required on the cluster. The ratio reflects the original size of data versus the size of the data after eliminating redundancy.
Compression Ratio	Part of the data-reduction process that performs a compression operation within a file or a set of files and further reduces the storage requirements. The ratio reflects the size of uncompressed data versus the size of compressed data.
Storage Reduction Ratio	The ratio of Logical Data to Storage Consumed.

Workload Tab

The **Workload** tab displays the following two charts:

- Top Organizations by Backup Data
- Top Organizations by View Data

> **Note:** The data displayed on the **Workload** tab is refreshed in real time.

> **Note:** The **Physical Usage** tab is not displayed in the tenant login. Do not use the **Storage Consumed by Organizations** report in the tenant login for billing purposes.

By default, the report displays all organizations in the cluster, but you can filter the data by organization(s). In the **Calculate Variance Over Specified Days** field, you can specify the number of days used to calculate the **Daily Logical Storage Growth Rate** column in the table. Enter a value and click **Apply Filters**.

The following table describes the fields displayed on the **Workload** tab:

Column Name	Description
Organization Name	The name specified for the organization when added to the cluster.
Storage Domain	The name of the storage domain assigned to the organization.
Data In	Data sent from primary objects to Cohesity.
Logical Used by Backups	Logical storage consumed by backups.
Protected Objects Count	The number of objects protected.
Physical Used by Backups	Physical storage consumed by backups.
Daily Logical Storage Growth Rate	This field shows the percentage daily growth rate of the total logical size of the views for the organization in this Storage Domain. The percentage daily growth rate is calculated using the following formula: `( (Sum of logical size of views Today - Sum of logical size of views of a Past Date) / ( (Number of Days Specified + 1) * Sum of logical size of views of a Past Date ) ) * 100` The past date is calculated from the current date and time minus the number of days specified in the **Calculate Growth Over Specified Days** field in the Actions panel.
Logical Used by Views	Logical storage consumed by views.

Storage Consumed by Servers

This report provides statistics about the Servers that are using the most disk space on the Cohesity cluster and Primary Storage. This report provides statistics about Servers only and not other Objects like Views.

This report provides the following charts

- **Top <Num> Servers by Primary Physical**—This chart lists up to twenty Servers whose Snapshots are using the most disk space on the Cohesity cluster that match the specified filter criteria.

In the **Calculate Variance Over Specified Days** field of the Actions panel, you can specify the number of days (from 1 to 180) used to calculate the **Primary Physical (At <Num> Days Ago)** and **Variance of Primary Physical** columns in the table. In the Actions panel, you can also optionally filter your results by Job Name, Registered Sources or a Storage Domain. You can enter part of a Job name, select one or more Registered Sources or a Storage Domain from the drop-down lists and click **Apply Filters**.

> **Note:** The data displayed in the **Storage Consumed by Servers** report is refreshed every six minutes.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The columns for the table at the bottom of this report are described below. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Primary Physical** column from smallest to largest.

Column Name	Description / Action
Protected Object	Drill down on this link to get statistics about the files stored in this Server.
Job Name	Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.
Storage Domain	The Storage Domain that is used to store the Snapshots made by the listed Protection Group of this Protected Object (Server).

Column Name	Description / Action
Snapshots	The total number of Snapshots of this Server created by all the Runs of the listed Protection Group that are currently stored on the Cohesity cluster and match the specified filter criteria. Snapshots that have been deleted because they exceeded the retention policy limit defined by the Protection Group are not counted.
Last Snapshot	The creation time and date of the last Snapshot of this Object (Server) created by the listed Protection Group that matches the specified filter criteria and currently stored on the Cohesity cluster.
Logical	The total capacity of the protected Source object, including free and used space. For more information, see Understanding Storage Metrics
Unique on Primary	Current value of actual data residing on source. Protected data is either equal to or less than this value for this source. The use-case where certain drives are excluded from protection is where this value is greater than the value of data protected from the object / source in the latest run.
Unique on Primary (<Num> Days Ago)	The same as **Unique on Primary**, but as of a past date. The past date is calculated from the current date and time minus the number of days specified in the **Calculate Variance Over Specified Days** field in the Actions panel.
Variance of Unique on Primary	Percentage difference between the values of Unique on Primary today, and Unique on Primary x days ago, where x is a user input value.

Storage Consumed by File Categories for Server: <myServer>

This report provides logical data statistics about File Categories and File Sizes for the last Snapshot of this Server created by the listed Protection Group. Statistics are only reported for files indexed by the Cohesity cluster.

A file in a Snapshot is not indexed due to one of the following reasons:

- In the Protection Group that creates the Snapshot, the directory where the file is located was not selected to be indexed.

This page provides the File Categories and File Size tabs that each contain a chart and table as described below.

File Categories

This tab provides the following chart:

- **Top Categories by Size**—This chart lists the collections of files of similar types that have the largest amount of logical space stored the Cohesity cluster for the last

Snapshot of this Server created by the listed Protection Group. Statistics are only reported for files indexed by the Cohesity cluster. See above for details. For example, the **image** file category reports all the files that end with image file extensions such as `*.png`, `*.tif`, `*.jpg`, `*.gif` etc.

In the **Calculate Daily Variance Over Specified Days** field of the Actions panel, you can specify the number of days used to calculate the **Total Logical Data Written (<Num> Days Ago)** and the **Variance** columns in the table. Enter a value from 1 to 180 and click **Apply Filters**.

This tab also provides the following table that lists statistics about the file categories found in the last Snapshot of this Server created by the listed Protection Group. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Total Logical Data Written** column.

Column Name	Description
File Category	The name of the file category. For example, the **image** file category reports all the files that end with image file extensions such as `*.png`, `*.tif`, `*.jpg`, `*.gif` etc.
Number of Files	The number of files of this category found in the last Snapshot of this Server created by the listed Protection Group. Statistics are only reported for files indexed by the Cohesity cluster.
Total Logical Data Written	As of today, the total size of the logical data for all the files in this Category is found in the last Snapshot of this Server created by the listed Protection Group. Statistics are only reported for files indexed by the Cohesity cluster. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication.
Total Logical Data Written (<Num> Days Ago)	As of a past date, the total size of the logical data for all the files in this Category found in the last Snapshot of this Server created by the listed Protection Group. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication. The past date is calculated from the current date and time minus the number of days specified in the **Calculate Variance Over Specified Days** field in the Actions panel.
Variance	This field shows the difference between the total logical data size for all the files in this Category found in the last Snapshot of this Server created by the listed Protection Group as of today compared to the total logical data size for all the files in this Category found in the last Snapshot of this Server created by the listed Protection Group as of a past date. The Variance is calculated using the following formula: `( (Total Logical Data Written as of Today - Total Logical Data Written on a Past Date)/ Total Logical Data Written on a Past Date) * 100`

File Size

This tab provides the following chart:

- **Top 20 Files by Size**—This chart lists the top twenty files with the largest amount of logical space found in the last Snapshot of this Server created by the listed Protection Group.

This tab also provides the following table that lists the top twenty files with the largest amount of logical space found in the last Snapshot of this Server created by the listed Protection Group. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Latest Logical Data Size** column.

Column Name	Description
File Name	The name of the file containing the logical data.
Latest Logical Data Size	The total size of the logical data for this file is found in the last Snapshot of this Server created by the listed Protection Group. Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication.

Storage Consumed by Storage Domains

This report provides information about the Storage Domains that are storing the most logical data and raw data data on the Cohesity cluster.

This report provides the following charts:

- **Top <Num> Storage Domains by Storage Used**—This chart lists up to ten Storage Domains that have stored the largest amount of data.
- **Top <Num> Jobs by Logical Data Managed**—This chart lists up to ten Storage Domains that have stored the largest amount of logical data from a Source.

In the **Calculate Variance Over Specified Days** field of the Actions panel, you can specify the number of days used to calculate the **Daily Storage Growth Rate** column in the table. Enter a value and click **Apply Filters**.

> **Note:** The data displayed in the **Storage Consumed by Storage Domains** report is refreshed every six minutes.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The columns for the table at the bottom of this report are described below. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Total Physical Data Written** column from smallest to largest.

Column Name	Description
Storage Domain	The name of the Storage Domain.
Logical	The sum of the logical sizes of all the protected Source objects, such as a VM, physical server or NAS share, stored in this Storage Domain. The logical size is the total capacity of the object, including free space and used space. For more information, see Understanding Storage Metrics.
Raw Used	Total space consumption which is the sum of storage consumption and space reserved for Cohesity software operations
Raw Used (<Num> Days Ago)	Same as **Raw Used**, but as of a past date. The past date is calculated from the current date and time minus the number of days specified in the **Calculate Growth Over Specified Days** field in the Actions panel.
Daily Storage Growth Rate	Daily growth rate for the **Storage Consumed** value.

Top Protection Group

This report shows the top Protection Groups by the average amount of data transferred, the length of Protection Group run time, the number of Objects protected and the number of SLA violations. Objects are entities such as VMs and Views that are protected by Protection Groups.

This report provides the following charts:

- **Data Read (average per run)**—This chart shows the top Protection Groups that have the highest average of data transferred. The average is calculated over all the runs of a Protection Group.

- **Average Run Time**—This chart shows the top Protection Groups that have the longest average run time. The average is calculated over all the runs of a Protection Group.
- **Total Objects Backed Up**—This chart shows the top Protection Groups that have the largest number of backed-up Objects. Objects are entities such as VMs and Views that are protected by Protection Groups.
- **Total SLA Violations**—This chart shows the top Protection Groups that have the largest number of SLA violations. An SLA violation occurs when the run time of the Protection Group exceeds the SLA time period specified in the Protection Group.

In the Actions panel, you can also optionally filter your results by Data Range, Job Name or Registered Source. You can specify a Date Range, enter part of a Job name, or select one or more Sources from the Registered Source drop-down list and click **Apply Filters**. You cannot filter on Storage Domain Sources.

In addition, the Top Jobs Report provides the Data Transferred, Run Time and SLA Violations tabs that each contain a table as described below.

Data Transferred Tab

This table lists the top Protection Groups that have the highest average of data transferred. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Average Transferred** column from largest to smallest.

Column Name	Description / Action
Job Name	Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.
Source	The Source that contains the Objects the Job is protecting. Objects are entities such as VMs and Views that are protected by Protection Groups.
Objects	The number of Objects protected by this Protection Group matches the filter criteria. Objects are entities such as VMs and Views that are protected by Protection Groups.
Data Read (average per run)	The average number of bytes read from Sources over all runs of this Protection Group matching the filter criteria
Data Read (total)	The cumulative number of bytes read from Sources for all the runs of this Protection Group matches the filter criteria. The amount of data actually stored on the Cluster is reduced by compression and deduplication.

Column Name	Description / Action
Data Reduction	The data reduction ratio is calculated by dividing the Logical Data Managed by the Physical Data Used. **Logical Data Written**–Indicates the size of the data if the data is fully hydrated or expanded. The actual physical data stored on the Cohesity cluster is reduced by change-block tracking, compression and deduplication. **Physical Data Written**–The size of the data physically stored on the Cohesity cluster, after the data has been reduced by change block tracking, compression and deduplication, including the disk space used for data resiliency (fault tolerance and erasure coding).

Run Time Tab

This table lists the top Protection Groups that have the longest average run time. The average is calculated over all the runs of a Protection Group. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **Avg Run Time** column from longest to shortest.

Column Name	Description / Action
Job Name	Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.
Source	The Source that contains the Objects the Job is protecting. Objects are entities such as VMs and Views that are protected by Protection Groups.
Protection Runs	The number of times this Protection Group has run within the specified filter criteria.
Avg Run Time	The average amount of run time for all the Runs of this Protection Group that match the specified filter criteria.
Min Run Time	The minimum amount of run time for all the Runs of this Protection Group that match the specified filter criteria.
Max Run Time	The maximum amount of run time from all the Runs of this Protection Group that match the specified filter criteria.

SLA Violations Tab

This table lists the top Protection Groups that have the largest number of SLA Violations. To sort the list in the table by Column type, click the Column header. To reverse the sorting, click the Column header again. By default, this table is sorted by the **SLA Violations** column from smallest to largest.

Column Name	Description / Action
Job Name	Drill down on this link to get statistics about this Protection Group such as throughput, megabytes read, a list of job runs, job details and a job audit trail.
Source	The Source that contains the Objects the Job is protecting. Objects are entities such as VMs and Views that are protected by Protection Groups.
Protection Runs	The number of times this Protection Group has run within the specified filter criteria.
Avg Run Time	This column reports the average amount of time the Protection Group Run takes to complete over all Runs of this Protection Group that match the specified filter criteria.
SLA	The SLA time period specified for this Protection Group when the Job was created or last edited. The service-level agreement (SLA) defines how long the administrator expects protection run to take.
SLA Violations	The number of Runs of this Protection Group that match the specified filter criteria that have taken longer than the SLA time period specified for this Protection Group.

Unprotected/Protected VMs

This report provides lists of VMs that are protected and not protected by Cohesity for the registered Sources on the Cohesity cluster. In addition, overview statistics about protected and unprotected VMs are provided.

This report provides the following tabs:

- **Unprotected VMs**—Lists the VMs that are not protected by Jobs on the Cohesity cluster grouped by registered Source.
- **Protected VMs**—Lists the VMs that are protected by Jobs on the Cohesity cluster grouped by registered Source.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

Each tab provides tables with the columns described below. You can sort by VM Size by clicking the triangle next to the VM Size heading.

Column Name	Description / Action
VM Name	The name of the VM.
VM Size	**Data Protection**: The combined total of data in the objects that are protected by Cohesity. The data size reported by VMware is the provisioned amount, not the actual data residing in the VM. For example, if a VM is provisioned for 1TB but contains only 100GB of data, VMware reports it as 1TB. **Note:** - These metrics do not include unprotected objects. - Currently, the logical data value shown on the UI is a sum of the logical data values captured across all the protection runs. For instance, if the source has 100GB of logical data, and assuming it remains at 100GB for the first 10 protection runs, Cohesity would report, after 10 runs, the **Logical Data** to be 1000GB (1TB).
Indexed	This applies to Protected VMs only. The current state of indexing of the VM: • **Fully-Indexed**–All the files in the VM were successfully indexed. • **Not-Indexed**–No files in the VM were indexed. Maybe indexing was not enabled for the Protection Group protecting this VM. • **Partially-Indexed**–Some files in the VM were indexed and some files were not. Indexing was enabled but not all files were indexed because the Cohesity cluster encountered a problem during indexing such as an unsupported file system. Indexing enables you to search and recover files or VMs.

User Quotas

This report provides information about the default user quota defined for a View and any overrides defined for individual users. For information about adding user quotas, see User and Directory Quotas for Views.

Begin typing a View name in the **View Name** field and then select a View from the drop-down.

In the Actions panel, you can optionally filter the report results by a Unix user ID or an Active Directory domain and user name, and click **Apply Filters**.

Click the **Email** icon to provide information for scheduling report emails. For more information, see Email Reports. The **Scheduled Emails** tab on the Reports page displays those emails.

Click the **CSV** icon to download the current report as a CSV file to your device.

The summary table at the top of the report contains the following columns.

Column Name	Description
Users	The number of users who have a quota that overrides the default quota for the View.
Default Quota per User	The default quota defined for all users of the View.
Users with Quota Exceptions	The number of users who have used 100% of their quota for View capacity.
Users Over Alert Thresholds	The number of users who have used 85% (the default threshold) of their quota.

The table at the bottom of the report contains the following columns.

Column Name	Description
User	The user for which a quota override is defined. Click the user name to drill down to more details.
Alert Threshold	The amount of View capacity the user can use before an Alert is triggered.
Quota	The quota defined for the user.
Usage	The percentage and graphical representation of how much of the quota is used.

Audit Logs

Cohesity records all events and operations occurring on your Cohesity cluster in the form of Audit Logs. An audit log contains details about an event, such as date and time, category, type, and user.

Audit logs help you as a cluster administrator to monitor and analyze the events that occurred if you need to troubleshoot issues and investigate unauthorized or suspicious activities.

> **Note:** Cohesity captures the audit logs at the cluster level and file services-level (NFS and SMB Views).

You can enable or disable the logging of the events in the audit logs and manage the audit logs retention period and access privileges.

For information about Cohesity Audit Logs, see the following topics:

- Cluster Audit Logs
- File Services Audit Logs
- Add a Syslog Server

Cluster Audit Logs

The Cluster audit logs feature records all events that occur at the cluster level. You can view these logs in the Cohesity UI and download the log files for analysis purposes.

You can enable or disable cluster audit logs and set retention periods. For more information, see Settings Cluster Audit Logs

The **Audit Logs** page lists the following actions that are logged by the Cohesity cluster:

- Details
- Time
- User
- Action
- Category
- Entity Name

You can use the filters to narrow the visible logs.

> **Note:** It can take up to 30 minutes for logged actions to appear in the Cohesity UI.

You can also download the audit logs as a CSV file. The CSV file may contain more records than the Cohesity UI can display. For more information, see Download Cluster Audit Logs.

Categories in Cluster Audit Logs

All cluster audit logs are logged under different predefined categories for you to find the relevant audit logs fast and analyze easily.

- API Key
- Access Token
- Active Directory
- Alert
- Alert Notification Rule

- Antivirus Service Group
- App
- Csr
- Chassis
- Clone Refresh Task
- CloudSpin
- Cluster
- Cluster Partition
- Cluster Services
- Disk
- Encryption Key
- Firewall Setting
- Group
- Helios Event
- Hybrid Extender
- IOTier
- IDP Configuration
- Infected File
- Interface
- IP
- KMS Configuration
- Keystone
- LDAP
- Network
- NIS
- Network Interface Group
- Node
- Tenant
- Physical Agent
- Preferred Domain Controller
- Object
- Protection Group

- Protection Run
- Protection Policy
- Source
- Proxy Server
- QoS
- Remote Cluster
- Resolution
- Recovery Task
- Role
- SMTP Server
- SNMP Config
- SSL Certificate
- Scheduler
- Search Job
- Service Flag
- Snapshot
- Static Route
- Storage Domain
- Support Server
- Swift Roles
- Tags
- Vault
- View
- Share
- VLAN

Logged Actions

The following table describes the actions that are logged.

Action	Description
Assign	A user assigns a source to a tenant.
Unassign	A user removes a source from a tenant.

Action	Description
Accept	A user accepted the license agreement.
Activate	A user activated a Protection Group.
Cancel	A user canceled an entity such as a running Protection Group or Recover task.
Clone	A user cloned an entity such as a Snapshot, VM, View or SQL Server.
Close	A user closed an SMB file open.
Create	A user created an entity such as a Protection Group.
Deactivate	A user deactivated a Protection Group.
Delete	A user deleted an entity such as a Protection Group, Protection Policy or View.
Disjoin	A user disjoined the Cluster from an AD domain.
Download	A user downloaded a VMX file or a file from a VM Snapshot.
Join	A user joined the Cluster to an AD domain.
Login	A user logged in to the Cohesity cluster.
Logout	A user logged out of the Cohesity cluster.
Mark	A user marked an entity for removal such as a disk.
Modify	A user modified an entity such as a User, Protection Group or Remote Cluster.
Pause	A user paused an entity such as a running Protection Group.
Recover	A user recovered an entity such as a VM, file, or SQL Database.
Refresh	A user refreshed the Object hierarchy in a Source.
Register	A user registered an entity such as an External Target (Vault).
Rename	A user renamed an entity such as a Storage Domain.
Resume	A user performed a resume action on a Protection Group.

Action	Description
Run Now	A user performed a Run Now action on a Protection Group.
Unregister	A user unregistered an entity such as a Source.
Upgrade	A user upgraded the Cohesity cluster.

Settings for Cluster Audit Logs

Enable or Disable Cluster Audit Logs

By default, the Cluster Auditing setting is enabled. You can disable or enable it back from the **Log Settings** tab of the **Audit Logs** page. When you disable the option, the old logs continue to retain until the retention period ends.

To enable or disable cluster auditing, complete the following in the Cohesity UI:

1. Select **System> Audit Logs>Log Settings**.
2. In the **Log Settings** tab, toggle the **ClusterAudit** button to enable or disable.

After enabling the **Cluster Audit**, option the **Log Retention Period** option appears. Using this option you can set the audit logs retention period.

Set Log Retention Period for Cluster Audit Logs

You can set the retention period for cluster audit logs. When you set a retention period, the logs are retained on the cluster until the retention period ends.

> **Note:** The default retention period is 90 days.

To set a retention period for cluster audit logs, complete the following in the Cohesity UI:

1. Select **System> Audit Logs>Log Settings**.
2. In the **Log Settings** tab, click the pencil icon ( ) for the **Log Retention Period** audit information option, type the desired number, and select retention period (Days, Weeks, Months, or Years).
3. The **Log Retention Period** is displayed in days as Cohesity converts weeks, months, or years into days.
4. Click **Save**.

Monitor Cluster Audit Logs

You can view cluster audit logs to know the activities on your cluster.

To monitor cluster audit logs, complete the following in the Cohesity UI:

1. Select **System>Audit Logs>Cluster.**
2. In the **Cluster** tab of the **Audit Logs** page, apply filters on the required fields (**Date Range**, **Users**, **Category**, or **Action**) and monitor logs.

Download Cluster Audit Logs

You can download the cluster audit logs for analyzing the logs outside your cluster. Cohesity cluster allows you to download the logs in a .CSV format.

> **Note:** The download .CSV file has more details than what the Cohesity displays. The file contains details about the IP addresses of the systems from which the cluster is accessed, tenants, impersonation, and so on.

To download a cluster audit log, complete the following in the Cohesity UI:

1. Select **System>Audit Logs>Cluster**.
2. In the **Cluster** tab of the **Audit Logs** page, click the ⬇ icon.

The CSV file automatically downloads to your local system.

Sample Audit Log

```
{Timestamp,User Name,Domain,Action,Entity Name,Entity Type,Entity Id,De-
tails,Previous Record,New Record,IP,ClusterInfo,Tenant,Original Ten-
ant,Impersonation "Sunday, 14-Nov-21 19:00:31 UTC",scheduler-
bot,local,ScheduleReport,Protection Summary by Object Report,Sched-
uler,kProtectionSummaryByObjectTypeReport,scheduler-bot@local has scheduled
email reports.,,,,"ClusterName: system-test-03, ClusterId:
633139267370196",,, "Sunday, 14-Nov-21 19:00:30 UTC",scheduler-bot,-
local,ScheduleReport,Protection Summary by Object Report,Sched-
uler,kProtectionSummaryByObjectTypeReport,scheduler-bot@local has scheduled
email reports.,,,,"ClusterName: system-test-03, ClusterId:
633139267370196",,, "Sunday, 14-Nov-21 16:30:25 UTC",ad-
min,local,Logout,admin,User,S-1-100-21-18873556-4717255-1,"admin@local
logged out from ""127.0.0.1"".",,,127.0.0.1,"ClusterName: system-test-03,
ClusterId: 633139267370196",,, "Sunday, 14-Nov-21 16:30:24 UTC",ad-
min,local,Login,admin,User,S-1-100-21-18873556-4717255-1,"admin@local logged
in from ""127.0.0.1"" using Iris CLI.",,,127.0.0.1,"ClusterName: system-
test-03, ClusterId: 633139267370196",,, "Sunday, 14-Nov-21 10:00:33
UTC",scheduler-bot,local,ScheduleReport,Available Local Snapshot
Report,Scheduler,kAvailableLocalSnapshotsReport,scheduler-bot@local has
scheduled email reports.,,,,"ClusterName: system-test-03, ClusterId:
633139267370196",,, "Sunday, 14-Nov-21 08:00:33 UTC",scheduler-bot,-
local,ScheduleReport,Agent deployment status report,Sched-
uler,kAgentDeploymentReport,scheduler-bot@local has scheduled email
```

```
reports.,,,,"ClusterName: system-test-03, ClusterId: 633139267370196",,,
"Sunday, 14-Nov-21 08:00:33 UTC",scheduler-bot,local,ScheduleReport,Agent
deployment status report,Scheduler,kAgentDeploymentReport,scheduler-bot@-
local has scheduled email reports.,,,,"ClusterName: system-test-03,
ClusterId: 633139267370196",,, "Sunday, 14-Nov-21 08:00:33 UTC",scheduler-
bot,local,ScheduleReport,Agent deployment status report,Sched-
uler,kAgentDeploymentReport,scheduler-bot@local has scheduled email
reports.,,,,"ClusterName: system-test-03, ClusterId: 633139267370196",,,
"Sunday, 14-Nov-21 06:58:25 UTC",admin,local,Create,admin,Access Token,ad-
min,admin@local has generated new access token for user admin on domain
local from 10.2.77.242.,,,10.2.77.242,"ClusterName: system-test-03,
ClusterId: 633139267370196",,, "Sunday, 14-Nov-21 06:58:25 UTC",ad-
min,local,Create,admin,Access Token,admin,admin@local has generated new
access token for user admin on domain local from
10.2.77.242.,,,10.2.77.242,"ClusterName: system-test-03, ClusterId:
633139267370196",,, "Sunday, 14-Nov-21 06:58:19 UTC",ad-
min,local,Create,admin,Access Token,admin,admin@local has generated new
access token for user admin on domain local from
10.2.77.242.,,,10.2.77.242,"ClusterName: system-test-03, ClusterId:
633139267370196",,, "Sunday, 14-Nov-21 06:58:02 UTC",ad-
min,local,Modify,system-test-01,Remote Cluster,711211522818758,"admin@local
modified remote cluster ""system-test-01"" with id
""711211522818758"".","autoRegisterTarget: false}
```


File Services Audit Logs

The File Services Audit Logs feature enables you to record and monitor events occurring on your SMB and NFS views.

Unlike the cluster-level audit logs, the logs for views are not displayed in the Cohesity GUI. The file services audit logs are saved in an internal view called ***filesystem_audit***. You can access these views by mounting the path ***COHESITY_AUDIT_VIEW\filesystem_audit***.

The **File Services** tab in the **Audit Log** page provides two mount paths for NFS and SMB. You can mount these paths to access and view the audit logs.

You can set permissions for the view and restrict access only to required subnets or users. Only the subnet's IP addresses in the allowlist can access and view the audit logs.

> **Note:**
>
> By default, when you create a view, the **Audit Log** option is disabled. You must manually enable it when creating the view or after creation.

Events Logged for NFS Views

Event	Description
Mount	To mount a view using NFS.
File Create	To create a file.
Link	To create a hard link to a file.
Folder Create	To create a folder.
Mknod	To create a special file used for resource like a device.
File Remove	To delete a file.
Folder Remove	To delete a folder.
Rename	To rename a file or folder.
Change Mode	To change the Unix mode bits of a file or folder for permission.
Change Ownership	To change the Unix UID or GID of a file or folder.
Modify Content	To modify the file's content.
Symlink	To create a soft link to a file or folder.

The following information is logged for NFS views. The information logged depends on the event type.

- Timestamp
- RecordID
- Protocol
- Client IP
- UserID
- ViewName
- Fsid
- RequestType
- Result
- EntityType

- EntityPath
- EntityInode
- EntityMode
- PrevEntityMode
- EntityUserID
- PrevEntityUserID
- EntityGroupID
- PrevEntityGroupID
- EntityModifyTimestamp
- DestinationEntityPath
- DestinationEntityInode
- InodeSizeBytes
- SourceParentEntityInode

Events Logged for SMB Views

Event	Description
Mount	To mount a view using SMB.
Logon	To establish an SMB session.
Logoff	To disconnect the current SMB session.
Close	To close a file or folder.
Create	To create a file or folder.
File Open	To open a file.
Delete	To delete a file or folder.
Rename	To rename a file or folder.
Modify Content	To modify the content of a file.
Change Ownership	To change the owner of a file or folder.
Change Permission	To add or remove the ACL permission of a file or folder.
File Change Protection	To enable or disable the Windows file protection of a file or folder.

The following information is logged for SMB views. The information logged depends on the event type.

- Timestamp
- RecordID
- Protocol
- Client IP
- Username
- UserID
- DomainName
- UserSID
- ViewName
- ViewAliasName
- Fsid
- RequestType
- Result
- EntityType
- EntityPath
- EntityInode
- EntityMode
- EntityUserID
- EntityGroupID
- SecurityDescriptor
- PrevSecurityDescriptor
- EntityModifyTimestamp
- EntityAccessMask
- DestinationEntityPath
- DestinationEntityInode
- DeleteOnClose
- InodeSizeBytes
- SourceParentEntityInode

Settings for File Services Audit Logs

Enable or Disable File Services Audit Logs

By default, the File Services Audit setting is disabled. You should enable the File Services Audit option for Cohesity to start logging operations on the cluster's views for auditing purposes.

You can disable or enable it back from the **Settings** tab of the **Audit Logs** page. When you disable the option, the old logs are retained until the retention period ends.

To enable or disable File Services Audit, complete the following:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **Log Settings** tab, turn on/off the **File Services Audit** toggle to enable or disable.

After enabling the **File Services Audit**, the **Log Retention Period** audit information option appears. Use this option to set the audit logs retention period.

Set Log Retention Period for File Services Audit Logs

You can set the retention period for file services audit logs on the Cohesity cluster.

> **Note:** The default retention period is 90 days. You can increase or decrease the renewal period. Once the retention period expires, the audit logs are deleted automatically.

Prerequisite:

- You must enable the File Services Audit option. For more information, see Settings for File Services Audit Logs.

To set a retention period for File Services Audit Logs, complete the following:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **File Services Audit** section of the **Log Settings** tab, click the icon for the **Log Retention Period** option, type the desired number, and select the retention period (Days, Weeks, Months, or Years).

   The Log Retention Period is displayed in days as Cohesity converts weeks, months, or years into days.
3. Click **Save**

Add IP Allowlist for File Services Audit Logs

By default, every client that is part of the Global Allowlist set at the cluster level can access the File Services audit logs. You can change access permission by the **Override Allowlist**

option to allow or deny access to a subnet. Additionally, access can also be restricted using SMB Permissions on the view for AD or LOCAL principals.

As the file services audit logs are saved in an internal view, you can provide access only to required networks by adding them to the allowlist along with Allow Deny setting for NFS/SMB Clients in the allowlist.

To override the global IP allowlist set at the cluster level and add a subnet to the allowlist for file service audit:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **IP Allowlist** section of the **Log Settings** page, turn on the **Override Global IP Allowlist** toggle.
3. Click **Add**.
4. In the **Add Allowlist** dialog, complete the following:
    1. Type the Subnet IP address in CIDR format.

       Example: Subnet in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).
    2. Select **NFS Permissions**, **SMB Permissions**, or **S3 Permissions**. (**Read/Write**, **Read Only**, or **Disabled**).
    3. Optionally, type a description.
    4. Click **Add**.

    The added allowlist IP subnet details appear in the File Services Audit section of the page.

Edit IP Allowlist for File Services Audit Logs

You can edit the IP allowlist for File Services audit logs if you need to change the subnet or change the permission for access.

To provide access to a network by adding to the IP allowlist for file service audit:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **IP Allowlist** section of the **Log Settings** page, click the for the required subnet.
3. In the **Edit Allowlist** dialog, complete the following:
    1. Modify the Subnet IP address.

       Example: Subnet in CIDR format (IPv4 - 10.0.0.0/24 or IPv6 - FE80:CD00::211E:729C/60).
    2. Modify **NFS Permissions**, **SMB Permissions**, or **S3 Permissions**. (**Read/Write**, **Read Only**, or **Disabled**).

3. Optionally, type a description.
4. Click **Add**.

The added allowlist IP subnet details appear in the File Services Audit section of the page.

Delete Allowlist for File Services Audit Logs

You can delete the allowlist IP subnet to remove access to a network. After deleting, clients in the deleted subnet cannot access the audit log files.

To delete an allowlist for File Services Audit logs:

1. Navigate to **System> Audit Logs>Log Settings**.
2. In the **IP Allowlist** section of the **Log Settings** page, click the delete icon for the required subnet.
3. In the **Delete Allowlist** dialog, enter **YES** to confirm, and click **Delete**.

The subnet disappears from the list and the audit logs view is restricted for the deleted subnet.

Add SMB Permissions for File Services Audit Logs

You can provide permission for SMB domain users and groups. The access permission is Read-Only and access type can be Allow or Deny.

To add SMB permissions for the File Services audit log:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **SMB Permissions** section of the **Log Settings** page, click the **Add**.
    1. In the **Add SMB Permission** dialog, complete the following:
    2. Select an **AD Domain** from the list.
    3. Select a **Principal** from the list.
    4. Select Type: **Allow** or **Deny**.
    5. Click **Add**.

    The new permission appears under the SMB Permissions section.

Delete SMB Permissions for File Services Audit Logs

You can delete an SMB permission to remove the set permission (Allow or Deny) for a principal (user or group).

To delete SMB permissions for file services audit log, complete the following:

1. Navigate to **System > Audit Logs** > **Log Settings**.
2. In the **SMB Permissions** section of the **Log Settings** page, click the delete icon for the required SMB permission.

3. In the **Delete SMB Permission** screen, type **Yes** for confirmation.
4. Click **Delete**.

Access NFS Audit Logs using a Mount Path

Prerequisite:

- The subnet you are using should have read-only permission to the File services audit log view.

To access NFS audit logs, complete the following steps:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Log Location** section of the **File Services** tab, click the copy icon for NFS and mount using the following steps:
    1. Mount the path using the syntax given below:

        ```
        mount cluster_name:/COHESITY_AUDIT_VIEW/filesystem_audit /mnt/path
        ```

    2. Connect to a hidden internal view 'COHESITY_AUDIT_VIEW'.
    3. Connect to `mount cluster_name:/COHESITY_AUDIT_VIEW/filesystem_audit /mnt/path`.

Sample NFS Audit Log


```
{"Timestamp" : "2022-03-10T22:12:17.838Z", "RecordID" : "bd65:2afc4", "Pro-
tocol" : "NFS3", "ClientIP" : "10.2.223.127", "UserID" : 1000, "ViewName" :
"test", "Fsid" : 4377, "RequestType" : "Delete", "Result" : "NFS3_OK",
"EntityType" : "File", "EntityPath" : "bbb", "EntityInode" : 9422800,
"SourceParentEntityInode" : 9422329}
```


Access SMB Audit Logs using a Mount Path

Prerequisites:

- The subnet you are using should have read-only permission to the File services audit log View.
- As an SMB user, you must have permission to access the File services audit log View.

To access SMB audit logs, complete the following steps:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Log Location** section of the **File Services** tab, click the copy icon for SMB and mount using the following steps:
    1. Connect to a hidden internal view 'COHESITY_AUDIT_VIEW'.
    2. Access `\\cluster_name\COHESITY_AUDIT_VIEW\filesystem_audit`.

Sample SMB Audit Log


```
{"Timestamp" : "2022-03-04T00:01:40.503Z", "RecordID" : "30fe:216", "Pro-
tocol" : "SMB", "ClientIP" : "10.14.0.41", "UserName" : "Administrator",
"UserID" : 500, "DomainName" : "w2k16-0.qa.local", "UserSID" : "S-1-5-21-
2148495376-676708989-2532157484-500", "ViewName" : "audit_filer_actions",
"Fsid" : 6642, "RequestType" : "SetAttributes", "Result" : "Ok",
"EntityType" : "File", "EntityPath" : "audit-actions/d1/e1/f1.txt",
"EntityInode" : 6214}
```


Check Audit Logs Status for Views

You can enable or disable the status of the Audit Logs option for views.

You can check the audit status for views to know whether audit logs are being logged for required views. The GUI displays all views with the status and it allows you to filter the views to see the audit log enabled or disabled list of views.

To check the audit status for views, complete the following:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Audit Status of Views** section, click the **Audit Logs** drop-down menu and select **Enabled** or **Disabled**, and click **Apply**

Based on your selection, the list of views with **Audit Log enabled** or **disabled** status appears.

Enable or Disable Audit Logs for Views

You can check the audit status of a view to know whether audit logs are being logged for required the view and enable or disable the **Audit Log** option.

You can also enable or disable the Audit Logs option for a view from: **SmartFiles > Views**, click the actions menu (⋮) and select **Edit** > **Audit Logs.**

To enable or disable audit logs status for a view, complete the following:

1. Navigate to **System > Audit Logs** > **File Services**.
2. In the **Audit Status of Views** section, click the **Audit Logs** drop-down menu and select **Enabled** or **Disabled**, and click **Apply**.
3. Based on your selection, the list of views with **Audit Log** enabled or disabled status appears.
4. To enable or disable audit logs, click the toggle corresponding to the view under the **Audit Logs** column.

Add a Syslog Server

You can add a Syslog server to the Cohesity cluster and send cluster audit logs, filer audit logs, alert audit logs, and data protection events entries to that server. These logs are

identified by their respective audit tags as listed below:

- Cluster audit log entries are identified with the tag `cluster_audit`.
- Filer audit log entries are identified with the tag `filesystem_audit`.
- Alert audit log entries are identified by the tag `cohesity_alerts`.
- Data protection events are identified by the tag `dataprotection_events`.

Optionally, you can set custom tags or token IDs for these Syslog messages to aggregate and parse them based on the set tags and tokens.

Add a Syslog server to the Cohesity cluster

There are two ways you can add a Syslog server to the Cohesity cluster.

- Adding a Syslog server using the Cohesity CLI
- Adding a Syslog server using the Cohesity UI

Adding a Syslog server using the Cohesity CLI

1. Start the Cohesity CLI remotely or locally as described in the Cohesity CLI Reference Guide.

   Specify the static IPv4 or IPv6 address of one of the Cohesity nodes. For example, if the CLI was downloaded to a Linux system, you might use the following command. Provide a Cohesity cluster username and password.

   ```
   ./iris_cli -server 172.16.3.101 -username=admin
   ```

2. Add a syslog server to the Cohesity cluster using the `syslog create` CLI command:

   ```
   admin>syslog create address=<IPAddressorHostname> port=<Port>
   ```

CLI Option	Example	Description
`address=`	`198.51.100.1`	Specify the hostname or IP address of the syslog server.
`name=`	`MySyslogServer`	Optional. Specify a name or description for the syslog server.
`port=`	`100`	Specify the port where the syslog server listens.
`protocol=`	`udp`	Optional. Specify the protocol used by the syslog server. The default is TCP, but UDP is also supported.
`program-names`	`cluster_audit`	Optional. The logs of specific programs sent to the syslog server.

CLI Option	Example	Description
`facilities`	`auth,authpriv`	The logs of specific facilities sent to the syslog server.
`permitted-peer`	-	Optional. Specify the syslog server certificate common name.
`message-patterns`	`denied`	The message with the patterns will be sent to the syslog server.
`raw-message-patterns`	`session`	The raw message with the patterns will be sent to the syslog server.
`ca-certificate-path`		Specify the syslog server CA certificate absolute file path.
`enable`	`true`	Enable or disable the syslog server to the Cohesity cluster.
cluster-audit	-	Deprecated. Same as program name “cluster_audit”
filer-audit	-	Deprecated. Same as program name “filesystem_audit”
dataprotection-events	-	Deprecated. Same as program name “dataprotection_events”
ssh-events-log	-	Deprecated. Same as program name “sshd”

> **Tip:** You can also use the CLI Help feature to view the `syslog create` options by entering the following command: `syslog create help`.

Adding a Syslog server using the Cohesity UI

1. Navigate to **Settings** > **Summary** and then select the **Syslog** tab.
2. Click **Add** located at the top-right of the page.
3. **IP Address/Hostname**: Specify the hostname or IP address of the syslog server.
4. **Protocol**: Specify the protocol used by the syslog server. The default is TCP, but UDP is also supported.
5. **Port**: Specify the port where the syslog server listens.
6. **Stream**: Select the logs to be sent to the syslog server from the drop-down list.

7. **Facilities (Optional)**: Enable this option to send the logs of specific facilities to the syslog server.
8. **Secure Connection**: Toggle to configure the authentication method to securely connect to the syslog server.
    1. **Authentication Mode**: Select one of two modes from the drop-down list.
        1. **anon**: This method uses TLS encryption without server identity validation, suitable for trusted networks where authentication is not required.
        2. **x509/name**: This method uses TLS with certificate-based authentication. The server's identity will be verified using a trusted CA certificate and an expected subject name.

           In **Permitted Peer Identity** field, enter the expected subject name of the syslog server's TLS certificate. For example, syslog.example.com.

           > **Important:** The connection will only be established when the server presents a certificate with this identity.
    2. **Trusted CA Certificate**: Paste the PEM-encoded certificate text or click **Select Certificate File** to select the file containing the certificate text.
9. **Enable Syslog Forwarding**: Toggle to forward logs and events to the syslog server.
10. Click **Add**.

Optional. Click **Manage Streams** to set custom tags or token IDs for cluster audit logs, file system audit logs, alert audit logs, and data protection event entries to aggregate and parse them based on the set tags and tokens.

You must set up an alert rule to get the logs for the Alert Audit Logs stream sent to the Syslog server. For more information, see Add Alert Notification Rule.

Enable or Disable Syslog Server

You can enable or disable the syslog server to the Cohesity cluster.

To enable a syslog server:

1. Navigate to **Settings** > **Summary** and then select the **Syslog** tab.
2. Find the syslog server in the list, click the **Edit** button, and toggle on or off the **Syslog Server** button to enable or disable the syslog server to the Cohesity cluster.

Delete a Syslog Server

You can delete a syslog server when you do not want the Cohesity cluster to send logs to the syslog server.

To delete a syslog server:

1. Navigate to **Settings** > **Summary** and then select the **Syslog** tab.
2. Find the syslog server in the list, click the **Delete** button and then click **Confirm** to delete the Syslog server.

Syslog Format and Keys Published to the Syslog Server

This topic provides the actions logged and published to the Syslog server.

Cohesity_Alerts

The following table describes the actions that are logged for the event `cluster_alerts`.

Field	Description
ClusterName	The cluster name on which the alert is triggered.
AlertCode	The alert code.
AlertName	The alert name.
AlertSeverity	The severity rating of the alert.
AlertDescription	A brief description of the problem that triggered the alert.
AlertCause	A brief description of the cause of the problem.

Data Protection_Events

The following table describes the actions that are logged for the event `dataprotection_events`.

Field	Description
EventMessage	A brief description of the event
Timestamp	The time the event occurs
ClusterInfo {ClusterId, ClusterName}	The cluster name and cluster id
EventType	The type of the event
EnvironmentType	Type of the adapter
EntityType	The type of entity

Field	Description
EntityId	The entity Id on which the event is generated.
EntityName	Entity Name on which the event is generated.
RegisteredSource	The source under which the entity is present.
RestoreTarget	The target entity for restore
BackupJobName	The protection run name generated for the event
BackupJobId	The protection run Id generated for the event
Entities	The entities for which the event is generated
Error {ErrorCode, ErrorMessage}	The error code and error message associated with the event
TaskId	The Id of the task
ReplicationTarget	The target cluster for replication
ArchivalTarget {Type, Name}	The archival target type and name
AttributeMap {_key, _value}	The key-value associated with the entity activities

Cluster_Audit

The following table describes the actions that are logged for the event `cluster_audit`.

Field	Description
Time	The time the audit log is generated
User	The username who created the audit log
Entity Name	The entity name
Category	The entity type
Action	The action-type for the audit log
Details	The change details of the audit log

For more information, see Cluster Audit Logs.

File Services_Audit

The following table describes the actions that are logged for the event `filesystem_audit`.

Field	Description
Timestamp	The time the event occurs
Protocol	The protocol for the specific event
Client Ip	The client's IP address
User name/User Id	The user name and the user ID of the user who performed the action
Domain name	The AD domain name
Share name	The name of the share the user is accessing
Entity Name(s)/Entity Id(s)	The name of the share and the share Id.
Entity attributes	The attributes of the entity for NFS

For more information, see File Services Audit Logs.

API_Audit

The following table describes the actions that are logged for the event `api_audit`.

Field	Description
Username	Cohesity login username
Domain	Active Directory or Local
Method	HTTP method
UrlPath	API URL
StatusCode	HTTP status code
ResponseHeader	HTTP response
ResponseTime	HTTP response time stamp
RequestTimestamp	HTTP request time stamp

CLI_Audit

The following table describes the actions that are logged for the event `cli_audit`.

Field	Description
User	Cohesity login username
Domain	Active Directory Domain or Local
Timestamp	The time the event occurs
Command	Iris_cli command

Shell_Commands_Audit

The following table describes the actions that are logged for the event `shell_commands_audit`.

Field	Description
IP Address	The Cohesity cluster IP address.
tty	The path of the terminal device.
User	Cohesity login username.
Command	Shell command

Console_login_audit

The following table describes the actions that are logged for the event `console_login_audit`.

Field	Description
User	The username attempting to log in For example: `cohesity` or `root`
tty	The path of the terminal device For example: `/dev/tty1`, `/dev/pts/0`

Field	Description
Node ID	The Cohesity cluster node ID where the console login event occurred
Action	The type of console event: login succeeded, login failed, or logout

Capture Logs

Use the `timecapsule` command to capture a comprehensive log bundle of Cohesity services across a single or multiple Nodes in a Cohesity cluster for a specified time range.

To capture logs:

Use any of the following `timecapsule` command depending on your requirement:

- To capture logs from the cluster (all Nodes) for all components for a time range:

  ```
  timecapsule --clusterip=<Cluster_Ip_address> -start_time_secs=<Seconds_in_Unix_Time> -end_time_secs=<Seconds_in_Unix_Time>
  ```

  For example, `timecapsule --clusterip=198.51.100.1 -start_time_secs=1549001025 -end_time_secs=1549173873`

  By default, the `timecapsule` command captures logs for all services.

- To capture logs from the cluster (all Nodes) for specific components for a time range:

  ```
  timecapsule --clusterip=<Cluster_Ip_address> -start_time_secs=<Seconds_in_Unix_Time> -end_time_secs=<Seconds_in_Unix_Time> -service_names=<Service_names_list>
  ```

  For example, `timecapsule --clusterip=198.51.100.1 -start_time_secs=1549001025 -end_time_secs=1549173873 -service_names=genie, magneto`

- To capture logs from a single or multiple nodes for all components for a time range:

  ```
  timecapsule --hostips=<Nodes_IP_addresses> -start_time_secs=<Seconds_in_Unix_Time> -end_time_secs=<Seconds_in_Unix_Time>
  ```

  For example, `timecapsule --hostips=198.51.100.1,198.51.100.2,198.51.100.3 -start_time_secs=1549001025 -end_time_secs=1549173873`

- To capture logs from a single or multiple nodes for specific components for a time range:

```
timecapsule --hostips=<Nodes_IP_addresses> -start_time_secs=<Seconds_in_Unix_Time> -end_time_secs=<SSeconds_in_Unix_Time> -service_names=<Service_names_list>
```

For example, `timecapsule --hostips=198.51.100.1 -start_time_secs=1549001025 -end_time_secs=1549173873 -service_names=bridge,apollo`

CLI Option	Placeholder Name	Example	Description
`--hostips=`	<Nodes_IP_addresses>	`198.51.100.1, 198.51.100.2, 198.51.100.3`	Specify the IP addresses of Nodes in a comma-separated list, from which you want to capture logs.
`--clusterip=`	<Cluster_Ip_address>	`198.51.100.2`	Specify the IP address of any of the nodes in the cluster from which you want to capture logs.
`-start_time_secs`	<Seconds_in_Unix_Time>	1549001025	Specify the start time in seconds (Unix timestamp).
`-end_time_secs`	<Seconds_in_Unix_Time>	1549001025	Specify the end time in seconds (Unix timestamp).
`-service_names`	<Service_names_list>	all, logwatcher, statscollector, physical_agents, windows_agent, and so on.	Specify the service names in a comma-separated list for which you to want parse logs. If the value provided is 'all', then log files of all services are parsed.
`-msg_type`	<log_file_level>	INFO, WARNING, ERROR, or FATAL	*Optional*. Specify the log file level to parse.
`-output_dir`	<output_directory>		*Optional* Specify the output directory for collected data. This directory should not already exist. The `timecapsule` command will create this directory.

CLI Option	Placeholder Name	Example	Description
`-skip_cores`		true	*Optional*. If true, skip copying the core files that are generated when a service crashes.
`-tar_gz_node_dumps`		true	*Optional*. If true, tar gzips the dump directories on each node before aggregating.
`-tar_gz_timecapsule`		true	*Optional*. If true, tar gzips the final **timecapsule** directory after aggregating the dumps from each node.
`-service_pages_only`		false	*Optional*. If true, it enables the collection of only Cohesity service pages (varz, flagz, tracez, tec). The default value is false.

Advanced Diagnostics

Cohesity provides common pre-canned statistics in charts and graphs with pre-determined X and Y axes at the following locations in the Cohesity UI:

- **Performance Statistics**—Select **System > Performance**. For more information, see Performance.
- **Storage Statistics**—Select **System > Storage**. For more information, see Storage Statistics.

In addition, Cohesity provides an extensive set of reports. For more information, see Reports.

Cohesity also provides advanced diagnostics that allow you to customize the plotted statistics.

To use Advanced Diagnostics:

1. Navigate to **System > Advanced Diagnostics**.
2. From the **Schema** drop-down, select a Cohesity component that you are interested in viewing statistics about. For example, if you want to learn more about physical disk

statistics on the Cohesity cluster, select **Disk Physical Stats** in the **Schema** drop-down. After selecting a Schema, the Metric and Entity drop-downs are populated with options specific to this Schema.

> **Note:** Entity is equivalent to an Object in the Cohesity Dashboard.

3. Select a Metric to plot. The options available in the Metric drop-down are specific to the chosen Schema. The Metric drop-down provides data that can be plotted on a graph. For example, after selecting the **Node Physical Stats** Schema, you could select the **Bytes Written** Metric.
4. Select an Entity to report on. The options available in the Entity drop-down are specific to the chosen Schema. The Entity drop-down defines the specific Cohesity Object to report on such as a disk, node or storage domain. For example, after selecting the **Node Physical Stats** Schema and the **Bytes Written** Metric, you can select one of the nodes in the cluster from the **Entity** drop-down.
5. Select a Rollup Function to apply to the data points returned by the selected Metric and Entity for the specified time period. For example, you could select the **Node Physical Stats** Schema, the **Bytes Written** Metric, the **Node:1** Entity and the **Average** Rollup Function.

   The available Rollup Functions are listed in the following table.

Rollup Function Name	Function Description
Raw	The data is not altered by a Rollup Function. All data points in the specified time interval are reported.
Sum	The sum of all the data points over the specified time interval. Add all the data points together for the specified time period.
Average	The mean of the data points over the specified time period. Sum of the data points divided by the number of data points.
Count	The number of data points for the specified time interval.
Max	The largest value of all the data points for the specified time interval.
Min	The smallest value of all the data points for the specified time interval.

Rollup Function Name	Function Description
Median	The middle value of all the data points if all the data points in the specified time interval are sorted by value. (50 Percentile)
Rate	The sum of the data points divided by the specified time interval in seconds. For example if the **Bytes Read** Metric is specified for the **Cluster Physical** Schema, the **Rate** Rollup Function and an interval of 10 Minutes, the number of bytes per second read by the cluster is reported. This is the throughput of the cluster. In this example, the rate of bytes per second is calculated by dividing the number of bytes read in 600 seconds by 600 seconds (10 Minutes).
95 Percentile	The data point that is 95% greater than all the other data points if all the data points in the specified time interval are sorted by value.
Latest	Is the last data point that was recorded during the specified interval.

> **Warning:** Not all Rollup Functions are appropriate for all metrics. For example, reporting the bytes read every hour on the cluster using the **Latest** Function is probably not very informative. However, reporting cumulative statics about the number of bytes read during a time period on the cluster is more informative. For example, you can report on the average number of bytes read on the cluster in one hour using the **Average** function or you can report on the total number of bytes read on the cluster during one day using the **Sum** function.

6. Select a time interval from the drop-down or choose **Custom** to specify your own time interval in seconds. For example, you could select the **Node Physical Stats** Schema, the **Bytes Written** Metric, the **Node:1** Entity, using the **Average** Rollup Function over a **10 Minutes** Interval.
7. Click **Plot** and either a new graph displays with the plotted data or a line of the plotted data is added to an existing graph. For this example, a new graph showing the average number of bytes written per ten-minute interval to a node is plotted as shown in the following figure.

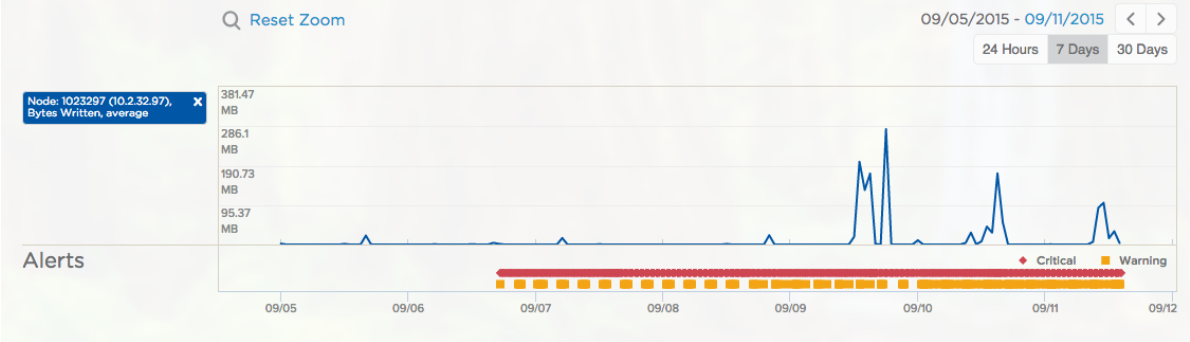

Plot lines that have the same Y-axis are plotted on the same graph with different colors. For example, all plot lines that report on bytes are plotted on the same graph. (The X-axis is always time.) A plot line with different Y-axis are plotted in a new graph.

A graph showing Critical (Red) and Warning (Yellow) Alerts can be found along the bottom of the page. You can click an Alert in the graphic to drill down to more information about the Alert. A detail page about the Alert is displayed in a new Browser tab.

8. You can delete a plot line by clicking the X in the top right corner of the colored label.
9. You can optionally further customize the data displayed on the Advanced Diagnostics page:
    - You can customize the data range reported by clicking **24 Hours**, **7 Days** or **30 Days**.
    - You can also change the starting date that is reported by clicking the arrow buttons or on the date.
    - You can zoom in on part of the graph by dragging the mouse across a section of the graph. The graph is zoomed in. To restore the original view, click **Reset Zoom**.
    - **Runs** - Toggle this setting on to view the type of protection run within the specific interval. On the graph, mouse hover on a specific protection run to view the start and end time of the protection run

Related Topic

Schema and Metric.

Schema and Metric

This topic provides details on the different types of schema and its corresponding metrics that helps you to understand the performance of each Cohesity component.

Object Level Backup Performance

The table below lists the metrics for the following schema:

- Backup Job Stats
- Backup Source Stats

> **Note:** The **Backup Job Stats** and **Backup Source Stats** schemas are not supported for Oracle sources.

Metric	Description
Bytes backed up	Based on the Rollup Function and Interval, the graph plots the total bytes backed up on the specified Entity.
Read latency from backup source	Based on the Rollup Function and Interval, the graph plots the read latency from the backup source. Preferred Rollup Function is **Average**.

Cloud Spinned VMs

The table below lists the metrics for the **Cloud Spin Stats** schema.

Metric	Description
Count of cloud spinned VMs	Based on the Rollup Function and Interval, the graph plots the total number of Virtual machines created on the cloud.

Cluster Performance

The table below lists the metrics for the following schema:

- Cluster CPU/MemoryUsage Metrics
- Node CPU/Memory Usage Metrics
- Process CPU/Memory Usage

Metric	Description
CPU Usage	Based on the Rollup Function and Interval, the graph plots the CPU usage of the specified Entity. Preferred Rollup Function is **Average**.
Memory Usage	Based on the Rollup Function and Interval, the graph plots the memory usage of the specified Entity. Preferred Rollup Function is **Average**.

Cluster Backup Performance

The table below lists the metrics for the **Cluster Backup and Restore Stats** schema.

Metric	Description
Bytes backed up	Based on the Rollup Function and Interval, the graph plots the total bytes backed up in the specified Entity.
Bytes restored	Based on the Rollup Function and Interval, the graph plots the total bytes restored in the specified Entity.
Num active backup tasks	Based on the Rollup Function and Interval, the graph plots the total number of active backup tasks running on the specified Entity.
Num active restore tasks	Based on the Rollup Function and Interval, the graph plots the total number of active restore tasks running on the specified Entity.
Read latency from backup source	Based on the Rollup Function and Interval, the graph plots the read latency from the backup source. Preferred Rollup Function is **Average**.
Write latency to restore target	Based on the Rollup Function and Interval, the graph plots the write latency to restore target. Preferred Rollup Function is **Average**.

External Target Space Utilization

The table below lists the metrics for the following schema:

- Cluster External Target Stats
- External Target Job Stats
- External Target Stats
- Node Physical Stats Storage
- Domain Logical Stats
- Domain Tier Physical Stats

Metric	Description
Cloud Direct Downtier Op Errors	Based on the Rollup Function and Interval, the graph plots the number of errors occurred while downtiering chunkfiles to cloud target. Preferred Rollup Function is either **Rate** or **Sum**.
Cloud Direct Downtier Op Latency	Based on the Rollup Function and Interval, the graph plots the latency of downtiering chunkfiles to cloud targets. Preferred Rollup Function is **Average**.
Cloud Direct Downtier Ops	Based on the Rollup Function and Interval, the graph plots the number of chunkfiles downtier ops triggered.
Cloud Direct Downtiered bytes	Based on the Rollup Function and Interval, the graph plots the number of bytes downtiered from the specified Entity.
Downtier Op Errors	Based on the Rollup Function and Interval, the graph plots the number of errors occurred while executing the downtier ops. Preferred Rollup Function is either **Rate** or **Sum**.
Downtier Op Latency	Based on the Rollup Function and Interval, the graph plots the latency of downtier ops. Preferred Rollup Function is **Average**.
Downtier Ops	Based on the Rollup Function and Interval, the graph plots the number of downtier ops triggered.
Downtiered bytes	Based on the Rollup Function and Interval, the graph plots the number of bytes downtiered.
Journal Bandwidth	Based on the Rollup Function and Interval, the graph plots the write bandwidth of the specified Entity. Preferred Rollup Function is **Rate**.
Journal Flush Bandwidth	Based on the Rollup Function and Interval, the graph plots the flush bandwidth of the specified Entity. Preferred Rollup Function is **Rate**.
Journal Flush Latency	Based on the Rollup Function and Interval, the graph plots the latency of flush operation observed on the specified Entity. Preferred Rollup Function is **Average**.

Metric	Description
Journal IOs	Based on the Rollup Function and Interval, the graph plots the number of write IOs. Preferred Rollup Function is **Rate**.
Random IOs	Based on the Rollup Function and Interval, the graph plots the number of random IOs. Preferred Rollup Function is **Rate**.
Read Bandwidth	Based on the Rollup Function and Interval, the graph plots the read bandwidth of the entity chosen. Preferred Rollup Function is **Rate**.
Read Errors	Based on the Rollup Function and Interval, the graph plots the number of read errors occurred on the specified Entity. Preferred Rollup Function is either **Rate** or **Sum**.
Read IOs	Based on the Rollup Function and Interval, the graph plots the number of Read IOs performed on the specified Entity. Preferred Rollup Function is **Rate**.
Read Latency	Based on the Rollup Function and Interval, the graph plots the latency of read operation observed by the entity chosen. Preferred Rollup Function is **Average**.
Sequential IOs	Based on the Rollup Function and Interval, the graph plots the number of sequential IOs performed on the specified Entity. Preferred Rollup Function is **Rate**.
Total Data Archived (Morphed)	Based on the Rollup Function and Interval, the graph plots the number of bytes archived to the external target after compression and encryption.
Total Data Archived (Unmorphed)	Based on the Rollup Function and Interval, the graph plots the number of bytes archived to the external target after decryption and decompression.
Total Data Restored (Morphed)	Based on the Rollup Function and Interval, the graph plots the number of bytes restored from the external target after compression and encryption.

Metric	Description
Total Data Restored (Unmorphed)	Based on the Rollup Function and Interval, the graph plots the number of bytes restored to the external target after decryption and decompression.
Total Data Uptiered (Unmorphed)	Based on the Rollup Function and Interval, the graph plots the number of bytes uptiered in the external target after decryption and decompression.
Total Garbage Collected (Morphed)	Based on the Rollup Function and Interval, the graph plots the number of garbage bytes collected from the external target after compression and encryption.
Total Garbage Collected (Unmorphed)	Based on the Rollup Function and Interval, the graph plots the number of garbage bytes collected from the external target after decryption and decompression.
Total IOs	Based on the Rollup Function and Interval, the graph plots the total IOs involved during the selected interval. Preferred Rollup Function is **Rate**.
Total Storage Used	Based on the Rollup Function and Interval, the graph plots the total storage used by the specified Entity. Preferred Rollup Function is **Sum**.
Write Bandwidth	Based on the Rollup Function and Interval, the graph plots the write bandwidth of the specified Entity. Preferred Rollup Function is **Rate**.
Write Errors	Based on the Rollup Function and Interval, the graph plots the write errors of the specified Entity. Preferred Rollup Function is either **Rate** or **Sum**.
Write IOs	Based on the Rollup Function and Interval, the graph plots the number of write IOs performed on the specified Entity. Preferred Rollup Function is **Rate**.
Write Latency	Based on the Rollup Function and Interval, the graph plots the write latency of the specified Entity. Preferred Rollup Function is **Average**.

Metric	Description
Dedup Morphed Usage	Based on the Rollup Function and Interval, the graph plots the estimated used bytes count (space utilized) in the specified Entity. This bytes count is post-encryption /post-compression/post-dedup. Preferred Rollup Function is either **Average** or **Latest**.
Dedup Unmorphed Usage	Based on the Rollup Function and Interval, the graph plots the estimated used bytes count (space utilized) in the specified Entity. This bytes count is pre-encryption /pre-compression/post-dedup. Preferred Rollup Function is either **Average** or **Latest**.
Non Dedup Morphed Usage	Based on the Rollup Function and Interval, the graph plots the estimated used bytes count (space utilized) in the specified Entity. This bytes count is post-encryption /post-compression/pre-dedup. Preferred Rollup Function is either **Average**.
Non Dedup Unmorphed Usage	Based on the Rollup Function and Interval, the graph plots the estimated used bytes count (space utilized) in the specified Entity. This bytes count is pre-encryption /pre-compression/post-dedup. Preferred Rollup Function is **Average**.

Space Consumption by Workloads

The table below lists the metrics for the following schema:

- Cluster Logical Stats
- Cluster Logical Stats by Workload
- Cluster Physical Stats
- Cluster Tier Physical Stats
- Disk Physical Stats
- Storage Domain Logical Stats by Workload Storage

Metric	Description
Cloud Direct Downtier Op Errors	Based on the Rollup Function and Interval, the graph plots the number of errors occurred while downtiering chunkfiles to cloud target. Preferred Rollup Function is either **Rate** or **Sum**.
Cloud Direct Downtier Op Latency	Based on the Rollup Function and Interval, the graph plots the latency of downtiering chunkfiles to cloud targets. Preferred Rollup Function is **Average**.

Metric	Description
Cloud Direct Downtier Ops	Based on the Rollup Function and Interval, the graph plots the number of chunkfiles downtier ops triggered.
Cloud Direct Downtiered bytes	Based on the Rollup Function and Interval, the graph plots the number of bytes downtiered from the specified Entity.
Cluster Capacity	Based on the Rollup Function and Interval, the graph plots the cluster capacity of the specified Entity.
Downtier Op Errors	Based on the Rollup Function and Interval, the graph plots the total errors occurred in the downtier op of the specified Entity. Preferred Rollup Function is either **Rate** or **Sum**.
Downtier Op Latency	Based on the Rollup Function and Interval, the graph plots the downtier latency of the specified Entity. Preferred Rollup Function is **Average**.
Downtier Ops	Based on the Rollup Function and Interval, the graph plots the number of downtier ops triggered.
Downtiered bytes	Based on the Rollup Function and Interval, the graph plots the total number of bytes downtiered.
Journal Bandwidth	Based on the Rollup Function and Interval, the graph plots the write bandwidth of the specified Entity. Preferred Rollup Function is **Rate**.
Journal Flush Bandwidth	Based on the Rollup Function and Interval, the graph plots the flush bandwidth of the specified Entity. Preferred Rollup Function is **Rate**.
Journal Flush Latency	Based on the Rollup Function and Interval, the graph plots the latency of flush operation observed on the specified Entity. Preferred Rollup Function is **Average**.
Journal IOs	Based on the Rollup Function and Interval, the graph plots the number of write IOs. Preferred Rollup Function is **Rate**.
Morphed Garbage	Based on the Rollup Function and Interval, the graph plots the estimated garbage bytes count in the specified Entity. This bytes count is post-encryption / post-compression.

Metric	Description
Morphed Usage	Based on the Rollup Function and Interval, the graph plots the estimated used bytes count (space utilized) in the specified Entity. This bytes count is post-encryption / post-compression. Preferred Rollup Function is **Average**.
Read Bandwidth	Based on the Rollup Function and Interval, the graph plots the read bandwidth of the specified Entity. Preferred Rollup Function is **Rate**.
Read Errors	Based on the Rollup Function and Interval, the graph plots the number of read errors performed on the specified Entity. Preferred Rollup Function is either **Rate** or **Sum**.
Read IOs	Based on the Rollup Function and Interval, the graph plots the number of reads IOs performed on the specified Entity. Preferred Rollup Function is **Rate**.
Read Latency	Based on the Rollup Function and Interval, the graph plots the read latency of the specified Entity. Preferred Rollup Function is **Average**.
Sequential IOs	Based on the Rollup Function and Interval, the graph plots the number of sequential IOs performed on the specified Entity. Preferred Rollup Function is **Rate**.
System Capacity	Based on the Rollup Function and Interval, the graph plots the system capacity of the specified Entity.
Total IOs	Based on the Rollup Function and Interval, the graph plots the number of write IOs performed on the specified Entity. Preferred Rollup Function is **Rate**.
Unmorphed Garbage	Based on the Rollup Function and Interval, the graph plots the estimated garbage bytes count in the specified Entity. This bytes count is the logical count that is pre-encryption / pre-compression.
Unmorphed Usage	Based on the Rollup Function and Interval, the graph plots the estimated used bytes count (space utilized) in the specified entity. This bytes count is the logical count pre-encryption / pre-compression. Preferred Rollup Function is **Average**.

Metric	Description
Write Bandwidth	Based on the Rollup Function and Interval, the graph plots the write bandwidth of the specified Entity. Preferred Rollup Function is **Rate**.
Write Errors	Based on the Rollup Function and Interval, the graph plots the write errors of the specified Entity. Preferred Rollup Function is either **Rate** or **Sum**.
Write IOs	Based on the Rollup Function and Interval, the graph plots the number of write IOs performed on the specified Entity. Preferred Rollup Function is **Rate**.
Write Latency	Based on the Rollup Function and Interval, the graph plots the write latency of the specified Entity. Preferred Rollup Function is **Average**.
Brick Referenced Usage	Based on the Rollup Function and Interval, the graph plots the total usage of all bricks in bytes.
Pre Replication Morphed Usage	Based on the Rollup Function and Interval, the graph plots the usage after dedup and compression but before replication or erasure coding.

CPU Performance

The table below lists the metrics for the **CPU Temperature Stats** schema.

Metric	Description
CPU Temperature	Based on the Rollup Function and Interval, the graph plots the CPU temperature of the specified Entity. Preferred Rollup Function is **Max**.

Disk Performance

The table below lists the metrics for the **Disk Health Metrics** schema.

Metric	Description
Disk Health	Based on the Rollup Function and Interval, the graph plots the disk health of the specified Entity. Preferred Rollup Function is **Average**.
Percentage Used (Wearouts)	Based on the Rollup Function and Interval, the graph plots the percentage of ssd disks used, measuring SSD wearout on the specified Entity. Preferred Rollup Function is **Average**.

Node Backup Performance

The table below lists the metrics for the following schema:

- Node Backup Stats
- Node Restore Stats
- Restore Source Stats
- Restore Task Stats

Metric	Description
Bytes restored	Based on the Rollup Function and Interval, the graph plots the total bytes restored on the specified Entity.
Num active backup tasks	Based on the Rollup Function and Interval, the graph plots the total number of active backup tasks running on the specified Entity.
Read Latency from backup source	Based on the Rollup Function and Interval, the graph plots the read latency from the backup source. Preferred Rollup Function is **Average**.
Write latency to restore target	Based on the Rollup Function and Interval, the graph plots the write latency to restore the target. Preferred Rollup Function is **Average**.

Power Supply Performance

The table below lists the metrics for the **Power supply temperature stats** schema.

Metric	Description
Power supply temperature	Based on the Rollup Function and Interval, the graph plots the power supply temperature of the specified Entity. Preferred Rollup Function is **Average**.

Replication Performance

The table below lists the metrics for the following schema:

- BackupJobReplicationStats
- BackupJobTargetNodeReplicationStats
- Target Replication Stats

Metric	Description
Logical bytes received	Based on the Rollup Function and Interval, the graph plots the logical bytes received on the specified Entity.
Logical bytes transmitted	Based on the Rollup Function and Interval, the graph plots the logical bytes transmitted by the specified Entity.
Number of metadata actions transmitted	Based on the Rollup Function and Interval, the graph plots the total number of metadata transmitted by the specified Entity.
Physical bytes received	Based on the Rollup Function and Interval, the graph plots the physical bytes received by the specified Entity.
Physical bytes transmitted	Based on the Rollup Function and Interval, the graph plots the physical bytes transmitted by the specified Entity.
Physical bytes transmitted on wire	Based on the Rollup Function and Interval, the graph plots the physical bytes transmitted on wire by the specified Entity.
Tx RPC round trip latencies	Based on the Rollup Function and Interval, the graph plots the Remote Procedure Call (RPC) round trip latency of the specified Entity. Preferred Rollup Function is **Average**.
Tx WAN Latencies	Based on the Rollup Function and Interval, the graph plots the WAN latency of the specified Entity. Preferred Rollup Function is **Average**.

Chapter 33: Virtual Edition

This chapter includes the following topics:

- Cohesity Virtual Edition 4910
- Initial Virtual Edition Cluster Setup 4911
- Configure vSphere Availability for Virtual Edition for VMware 4911
- Extend the Capacities of a Virtual Edition for VMware 4912

Cohesity Virtual Edition

Cohesity Virtual Edition is a single-node Cohesity cluster (ROBO) or multi-node Cohesity cluster that is hosted on a Virtual Machine that is hosted on VMware vCenter Server, Microsoft Hyper-V server, Nutanix AHV, RHEL KVM, AlmaLinux, or OCI Linux KVM. You can use a single-node Virtual Edition if your environment has smaller offices with reduced workloads that do not require the full computing power of a Cohesity cluster running on multiple nodes. In addition, Cohesity Clustered Virtual Edition is a multiple node Cohesity cluster that can support larger workflows and is hosted on multiple Virtual Machines in VMware vCenter Server, Nutanix AHV, RHEL KVM, AlmaLinux, or OCI Linux KVM. For information on the hypervisors and their support for single-node and multi-node clusters, see Virtual Edition.

A single Node Virtual Edition is appropriate for the following reduced workflow example. A company has San Jose and Las Vegas remote offices that require backing up of a small number of VMware Virtual Machines. However, the company wants to centralize the backups on a central Cohesity cluster located in the Datacenter of the Corporate Headquarters. These two smaller remote offices can use Cohesity cluster Virtual Editions to capture Snapshots of the VMs and replicate the Snapshots to the central Cohesity cluster located in the Chicago Corporate Datacenter. This central Cohesity cluster is running directly on hardware supported by Cohesity.

In addition, other remote offices can also replicate Snapshots to the central Cohesity cluster using Cohesity clusters running directly on supported hardware as shown by the following examples:

- At the Cincinnati remote office, a Cohesity cluster captures and then stores Snapshots of a physical server.
- At the Des Moines remote office, a Cohesity cluster captures and then stores Snapshots of Virtual Machines.

These Snapshots are then replicated to the central Cohesity cluster located in Chicago.

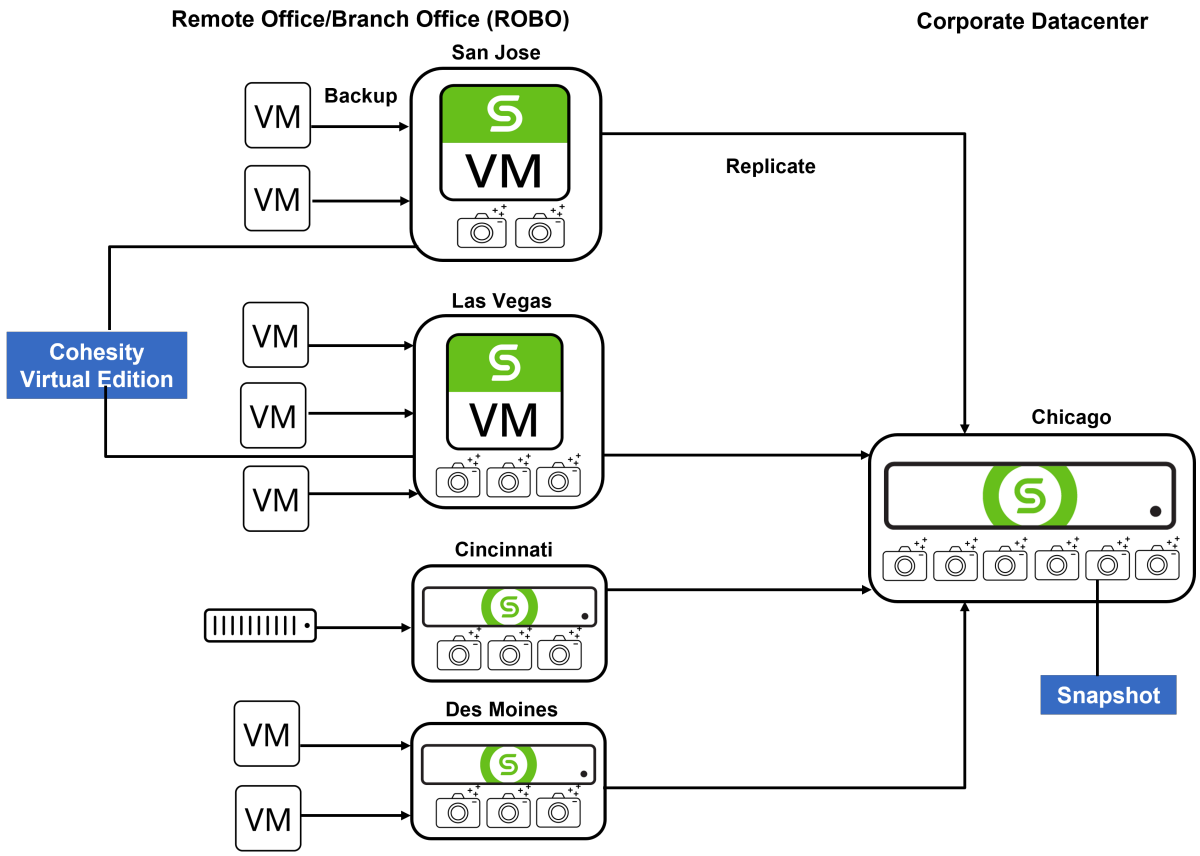

Initial Virtual Edition Cluster Setup

The initial Cohesity cluster Virtual Edition setup instructions are available at the Initial Cluster Setup.

Configure vSphere Availability for Virtual Edition for VMware

- Turn On vSphere HA
- Set the Host Failure Response to Restart VMs
- Enable VM Monitoring of the heartbeat

To enable these vSphere Availability Settings, follow the **Configure vSphere Availability for Virtual Edition for VMware** procedure in the Setup Guide (Cohesity Virtual Edition for VMware) and complete the appropriate procedure for your vSphere version.

Extend the Capacities of a Virtual Edition for VMware

Extending the node capacities of a Virtual Edition(VE) for VMware is predefined by the cluster configuration, either SMALL or Large.

> **Note:** For multi-node cluster,
> - you can add or extend the disk for individual nodes or all nodes.
> - you can add multiple disks to a node.
> - you must add or extend the disk on only one node at a time.

This section assumes that the node to be extended is part of the cluster and is healthy with no cluster operation running on the node. To check the cluster operations, run `iris_cli cluster status` and ensure the cluster status is as below:

```
SERVICE STATE SYNC: DONE
```

```
CLUSTER ACTIVE OPERATION: <this field must be empty>
```

```
CLUSTER HEAL STATUS: NORMAL
```

Small Configurations (8TB VE)

Small configurations support a maximum of 8TB data disk space on each node accumulating all data disks on the node. The node to be extended must satisfy the following conditions before and after the capacity extend:

Settings	Value
Minimum size for each Metadata disk	256GB (preferred size is 512GB)
Maximum size for each Metadata disk	1TB
Maximum total number of Metadata disks allowed to be attached to the node	2
Maximum total size of all Metadata disks on a node	1TB
Minimum size for each Data disk	512GB (preferred size is 1TB)
Maximum size for each Data disks	8TB
Maximum total number of data disks allowed to be attached to the node	8
Maximum total size of all Data disks on a node	8TB
Data disk size (x) to Metadata disk size (y) maximum ratio	x/y<=20

Settings	Value
Data disk size (x) to Metadata disk size (y) preferred ratio	x/y<=16
Minimum number of CPUs	12
Minimum memory required	32GB

Large Configuration (16TB VE)

Large configurations support a maximum 16TB data disk space on each node accumulating all data disks on the node. The node to be extended must satisfy the following conditions before and after the capacity extend:

Settings	Value
Minimum size for each Metadata disk	256GB (preferred size is 512GB)
Maximum size for each Metadata disk	1TB
Maximum total number of Metadata disks allowed to be attached to the node	4
Maximum total size of all Metadata disks on a node	2TB
Minimum size for each Data disk	512GB (preferred size is 1TB)
Maximum size for each Data disks	8TB
Maximum total number of data disks allowed to be attached to the node	8
Maximum total size of all Data disks on a node	16TB
Data disk size (x) to Metadata disk size (y) maximum ratio	x/y<=20
Data disk size (x) to Metadata disk size (y) preferred ratio	x/y<=16
Minimum number of CPUs	24
Minimum memory required	64GB

During the initial installation and configuration of the Virtual Edition Virtual Machine, you select one of the predefined configuration deployment types, either SMALL or LARGE.

The initial installation and configuration process of the Virtual Edition VM is described in the Virtual Edition for VMware Setup Guide (PDF).

Review the following prerequisites and warnings before starting the process to extend the capacities of the Virtual Edition VMware VM:

- You must have a vCenter Server account that has been granted administrator privileges for the ESXi host that is hosting the Virtual Edition.
- During this process, the Cohesity cluster services must be stopped and the cluster must be powered off. Therefore, during this process, the cluster is unable to run workflow tasks such as Protection Groups, Recover tasks, Replicating Data tasks, etc. Cohesity recommends stopping all tasks on the cluster before starting this process.
- If the ESXI host reboots, power cycles or powers off during this extension process, data loss may occur. If this occurs, contact Cohesity Support.
- The ESXi Datastores that are hosting the Virtual Edition VM must have enough capacity to support the disk size expansion of the VM. A new Datastore cannot be used to extend the size of the existing disks on the VM. It can be used only to expand the capacities of an existing VM by adding a new disk with a restriction on the total number of disks attached to the VM.
- All disks must be attached to different SCSI controllers.

You can extend the node capacities of a Virtual Edition VMware in two ways:

- Add a new disk to Virtual Edition VMware.
- Extend an existing disk on the Virtual Edition VMware.

> **Note:** Cohesity recommends adding multiple disks each from different datastore to the Virtual Edition VMware as it provides better disk bandwidth.

Add a New Data Disk or Metadata Disk

1. Run the Cohesity CLI locally or remotely on the Virtual Edition. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

   ```
   iris_cli
   admin>
   ```

2. Verify the capacities of the existing Metadata and Disk Tier disks. For a Metadata disk, STORAGE TIER is set to PCIeSSD. For a Data Tier disk, STORAGE TIER is set to SATA-HDD.

   ```
   admin> disk ls
   ```

   Sample Output

   In the following sample output, the capacity of the Data disk is 7.99 TiB and the capacity of the Metadata disk is 500 GiB.

```
DISK ID                   : 7
CURRENT NODE ID           : 1508041826908565
CURRENT NODE IP           : 10.2.43.160
REMOVAL STATE             :
REMOVAL REASON            :
STORAGE TIER              : SATA-HDD
DISK SERIAL               :
LOCATION IN CHASSIS       : N/A
MOUNT PATH                : /home/cohesity/data/disks/sdc2_b1c65849-
40ad-4e59-86e5-51c0134f73bc
MAXIMUM PHYSICAL CAPACITY : 7.99 TiB
MAXIMUM USABLE CAPACITY   : 7.99 TiB
DISK RESERVATIONS         :
librarian                 : 163.64 GiB
apollo                    : 114.55 GiB
hydra                     : 81.82 GiB
groot                     : 81.82 GiB

DISK ID                   : 6
CURRENT NODE ID           : 1508041826908565
CURRENT NODE IP           : 10.2.43.160
REMOVAL STATE             :
REMOVAL REASON            :
STORAGE TIER              : PCIeSSD
DISK SERIAL               :
LOCATION IN CHASSIS       : N/A
MOUNT PATH                : /home/cohesity/data/disks/sdb1_5a6ccde4-
4600-4b50-9ae6-ab441c2fe42b
MAXIMUM PHYSICAL CAPACITY : 500.00 GiB
MAXIMUM USABLE CAPACITY   : 500.00 GiB
DISK RESERVATIONS         :
scribe                    : 163.64 GiB
hydra                     : 35.00 GiB
yoda                      : 5.00 GiB
```

3. Using a VMware vSphere Web Client, log in to a vCenter Server that hosts Virtual Edition VM:
    1. Enter the URL of your vCenter Server in the address bar of a web browser in the following format: https://<vcenter_server_ip_address_or_fqdn>.
    2. On the vSphere Welcome page, click **Launch vSphere Client (HTML5)**.
    3. On the vSphere login page, enter the username and password of the vCenter Server and click **Login**.
4. In the left pane, browse for the Virtual Edition VM to extend capacity.

5. If the VM is running, right-click the VM and select **Power Off**.
6. Right-click the VM and select **Edit Settings**.
7. Click the **Add hard disk** tab > **New hard disk**.
8. Update the size of the new hard disk.

   > **Warning:** Do not change the size of existing data disks.

   > **Note:** If the newly added disk is for the data tier, the data tier size after extension must be at most 16 times than the size of the metadata tier. If this is not the case, you need to extend the metadata disk tier capacity either by updating the size of the existing metadata disk or adding one more disk for the metadata tier. For example if the Virtual Edition VM was originally a configuration with a 500GB Metadata disk and a 8000GB Data Tier disk and you increase the disk sizes to a 1000 GB Metadata disk and a 16000 GB Data Tier disk, you must also increase the number of vCPUs to 8 and the memory to 64000. Ensure that the Data Tier disk size does not exceed 16 times the Metadata disk size. In addition, the total Data Tier disk capacity per node must not exceed the supported upper limits of 8 TB or 16 TB, depending on the configuration. The Metadata drive size must be smaller than the Data Tier drive size.

9. Click **OK** to save the changes. Wait until this operation is completed.
10. Right-click the VM and select **Power On**.
11. Run the Cohesity CLI locally or remotely on the Virtual Edition. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

    ```
    iris_cli
    admin>
    ```

12. Stop the cluster services.

    ```
    admin> cluster stop
    ```

    Wait for 2 minutes for the cluster services to stop.

13. Find and copy the node ID of the cluster.

    ```
    admin> node list
    ```

14. Based on the requirement, perform the following:

    Adding a Data Disk to the Node

> **Note:** You must be on the same VM node where the new disk is added. If the data-to-metadata tier ratio is not maintained after adding a new data disk, you must add a new metadata disk or extend the existing one before adding a new data disk to maintain the ratio.

Perform the following steps:

1. Run `lsblk` on bash on the node to see the new disk added and copy its disk path.
2. Run the following command:

```
iris_cli
```

```
admin> disk assimilate disk-paths=<disk_path_of_new_disk_on_node>
node-id=<node ID>
```

Adding a Metadata Disk to the Node

Perform the following steps:

1. Run `lsblk` on bash on the node to see the new disk added and copy its disk path.
2. Run the following command:

```
iris_cli
```

```
admin> disk assimilate disk-paths=<disk_path_of_new_disk_on_node>
node-id=<node ID> is-metadata=true
```

15. View the current disk size

```
admin> disk ls
```

Keep repeating this command until the MAXIMUM PHYSICAL CAPACITY field settings for both the Metadata disk and the Data Tier disk are equal to the new sizes you specified. For a Metadata disk, STORAGE TIER is set to PCIeSSD. For a Data Tier disk, STORAGE TIER is set to SATA-HDD.

> **Note:** It can take up to 10 minutes for the size to be increased.

16. After the sizes are equal to the new sizes you specified, start the cluster services.

```
admin> cluster start
```

Wait for 2 minutes for the cluster services to start.

17. Log in to the Cohesity Dashboard:

a. Start a browser and enter the Cohesity cluster domain name in the browser address bar.
b. Log in to the Cohesity Dashboard. The default System Admin account is `admin` and the default password is `administartor`.

18. Wait for the Cohesity Dashboard to report the updated storage capacity.

    It can take about 10 minutes for the Cohesity Dashboard to report the updated storage capacity.

19. Resume running the workflow tasks on the cluster such as Protection Groups, Recover tasks, Replicating Data tasks, etc.

Extend Existing Data Disk or Metadata Disk

1. Run the Cohesity CLI locally or remotely on the Virtual Edition. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

```
iris_cli
admin>
```

2. Verify the capacities of the existing Metadata and Disk Tier disks. For a Metadata disk, STORAGE TIER is set to PCIeSSD. For a Data Tier disk, STORAGE TIER is set to SATA-HDD.

```
admin> disk ls
```

Sample Output

In the following sample output, the capacity of the Data disk is 7.99 TiB and the capacity of the Metadata disk is 500 GiB.

```
DISK ID                   : 7
CURRENT NODE ID           : 1508041826908565
CURRENT NODE IP           : 10.2.43.160
REMOVAL STATE             :
REMOVAL REASON            :
STORAGE TIER              : SATA-HDD
DISK SERIAL               :
LOCATION IN CHASSIS       : N/A
MOUNT PATH                : /home/cohesity/data/disks/sdc2_b1c65849-
40ad-4e59-86e5-51c0134f73bc
MAXIMUM PHYSICAL CAPACITY : 7.99 TiB
MAXIMUM USABLE CAPACITY   : 7.99 TiB
DISK RESERVATIONS         :
librarian                 : 163.64 GiB
apollo                    : 114.55 GiB
hydra                     : 81.82 GiB
```

```
groot                       : 81.82 GiB

DISK ID                     : 6
CURRENT NODE ID             : 1508041826908565
CURRENT NODE IP             : 10.2.43.160
REMOVAL STATE               :
REMOVAL REASON              :
STORAGE TIER                : PCIeSSD
DISK SERIAL                 :
LOCATION IN CHASSIS         : N/A
MOUNT PATH                  : /home/cohesity/data/disks/sdb1_5a6ccde4-
4600-4b50-9ae6-ab441c2fe42b
MAXIMUM PHYSICAL CAPACITY : 500.00 GiB
MAXIMUM USABLE CAPACITY   : 500.00 GiB
DISK RESERVATIONS           :
scribe                      : 163.64 GiB
hydra                       : 35.00 GiB
yoda                        : 5.00 GiB
```

3. Using a VMware vSphere Web Client, log in to a vCenter Server that hosts Virtual Edition VM:
    1. Enter the URL of your vCenter Server in the address bar of a web browser in the following format: https://<vcenter_server_ip_address_or_fqdn>.
    2. On the vSphere Welcome page, click **Launch vSphere Client (HTML5)**.
    3. On the vSphere login page, enter the username and password of the vCenter Server and click **Login**.
4. The vSphere Client home page opens. The left panel displays the hosts, clusters, and VMs in the vSphere inventory. From the left pane, find the Virtual EditionVM to extend capacity.
5. If the VM is running, right-click the VM and select **Power Off**.
6. Right-click the VM and select **Edit Settings**.
7. Select the hard disk that you want to extend and update its size.

    > **Note:** If you are extending the Data disk, the disk size must be extended to a size that is at most 16 times the size of the Metadata disk. If this requirement is not met, you need to extend the Data disk capacity by increasing the size of the existing Data disk or by adding an additional Data disk. To add a new disk, see Add a New Data Disk or Metadata Disk.

8. Click **OK** to save the changes. Wait until this operation completes.
9. Right-click the VM and select **Power On**.

10. Run the Cohesity CLI locally or remotely on the Virtual Edition. For more information, see Using the Cohesity CLI and List of Cohesity CLI Commands.

```
iris_cli
admin>
```

11. Stop the cluster services.

```
admin> cluster stop
```

Wait for 2 minutes for the cluster services to stop.

12. Disable Secure Shell

```
admin> cluster secure-shell enable=false
```

> **Note:** After disabling the secure shell, the current session will exit, wait a few minutes and log back in again.

13. Stop the containerd service.

```
[support@virtual-node-1 ~]$ sudo systemctl stop containerd
```

14. Check for any Cohesity services that are still actively using the disk. Ensure the output of this command is empty before you proceed. Update the sdc2 value according to the disk that needs to be extended (refer to the output from Step 2).

```
[support@virtual-node-1 ~]$ sudo lsof | grep /home_cohesity_
data/disks/sdc2
```

15. Find and copy the node ID of the cluster.

```
admin> node list
```

16. Extend the size of the disks

```
admin> disk extend node-id=<node ID of node where disks are extended>
```

> **Note:** If you are extending the Data disk or Metadata disk size by adding a new disk, follow the steps in Add a New Data Disk or Metadata Disk.

17. View the current disk size

```
admin> disk ls
```

Keep repeating this command until the MAXIMUM PHYSICAL CAPACITY field settings for both the Metadata disk and the Data Tier disk are equal to the new sizes you specified. For a Metadata disk, STORAGE TIER is set to PCIeSSD. For a Data Tier disk, STORAGE TIER is set to SATA-HDD.

> **Note:** It can take up to 10 minutes for the size to be increased.

18. Start the containerd service.

```
[support@virtual-node-1 ~]$ sudo systemctl start containerd
```

19. Start Secure Shell

```
admin> cluster secure-shell enable=true
```

> **Note:** Enabling the secure shell will disconnect the current session, wait a few minutes and login again.

20. After the sizes are equal to the new sizes you specified, start the cluster services.

```
admin> cluster start
```

Wait for 2 minutes for the cluster services to start.

21. Log in to the Cohesity Dashboard:
    a. Start a browser and enter the Cohesity cluster domain name in the browser address bar.
    b. Log in to the Cohesity Dashboard. The default System Admin account is `admin` and the default password is `administartor`.
22. Wait for the Cohesity Dashboard to report the updated storage capacity.

    It can take about 10 minutes for the Cohesity Dashboard to report the updated storage capacity.
23. Resume running the workflow tasks on the cluster such as Protection Groups, Recover tasks, Replicating Data tasks, etc.

Chapter 34: Cohesity Cloud Edition

Cohesity simplifies data management for organizations, allowing them to experience the benefits of the cloud without giving up control over their data. Cohesity Cloud Edition is a Cohesity cluster deployed in a cloud service, designed to offer high performance and greater cluster capacity.

Cohesity supports the following Cohesity Cloud Edition:

- Cohesity Cloud Edition on AWS
- Cohesity Cloud Edition on Azure
- Cohesity Cloud Edition on GCP
- Cohesity Cloud Edition on OCI

For information on the Cohesity Cloud Edition supportability, seeCohesity Cloud EditionCohesity Cloud Edition Supportability.

Cohesity Cloud Edition on AWS

Cohesity Cloud Edition on AWS is a Cohesity cluster deployed in AWS, designed to offer high performance and greater cluster capacity. This solution employs advanced data deduplication and compression techniques to significantly reduce data footprints and storage costs.

This section provides the following information on Cohesity Cloud Edition on AWS:

Supported Regions for Deploying Cohesity Cloud Edition on AWS

You can deploy Cohesity Cloud Edition on AWS in all regions supported by the AWS cloud (including AWS GovCloud). For the list of supported regions, see Regions and zone.

Supported Cohesity Cloud Edition on AWS Configurations

Based on the performance and cluster capacity per cluster node requirement, Cohesity supports the following cluster configurations:

<table>
<tr><th>Cluster Configuration</th><th>Instance Type (Intel/AMD)</th><th>EBS Volume Type</th><th>CPU and Memory</th><th>MetaData Tier</th><th>Cache Tier</th><th>Boot Disk</th></tr>
<tr><td>Entry</td><td>m6i.2xlarge /m6a.2xlarge</td><td rowspan="4">gp3 for metadata and cache tiers.

gp3 for boot disk.</td><td>8 vCore 32 GB</td><td>2 x 511 GB</td><td>2 x 512 GB</td><td rowspan="4">250 GB: For new deployments of cluster versions 7.3.1 and above.

81 GB: For cluster versions below 7.3.1.</td></tr>
<tr><td>Core</td><td>m6i.4xlarge/m6a.4xlarge</td><td>16 vCore 64 GB</td><td>2 x 1 TB</td><td>2 x 1 TB</td></tr>
<tr><td>Archive</td><td>m6i.4xlarge/m6a.4xlarge</td><td>16 vCore 64 GB</td><td>2 x 2 TB</td><td>2 x 2 TB</td></tr>
<tr><td>Performance</td><td>m6i.8xlarge/m6a.8xlarge</td><td>32 vCore 128 GB</td><td>2 x 1 TB</td><td>2 x 1 TB</td></tr>
</table>

Cohesity Cloud Edition on AWS Permissions

The following table lists the permission that is required for the deploying Cohesity Cloud Edition on AWS:

For information on backup and recovery permissions, see AWS resources.

> **Note:** For a node being added to an existing cluster. Once the cluster is deployed, the IAM permissions are not required and can be removed.

Open the JSON (or right-click and save as a file).

AWS Resource Type	Permissions	Description
IAM	ListAttachedUserPolicies GetPolicyVersion GetPolicy ListGroupsForUser ListUsers ListPolicies ListRolePolicies GetUser GetRole CreateRole TagRole AttachRolePolicy DetachRolePolicy CreateInstanceProfile RemoveRoleFromInstanceProfile DeleteInstanceProfile PassRole AddRoleToInstanceProfile GetInstanceProfile DeleteRole	Cohesity uses the access key and the secret access key provided to validate the IAM user. IAM permissions are required to get the user, the policy for the user and to validate the credentials. To create a cluster by using Helios, Cohesity uses System Service Manager (SSM). A role is created for SSM and passed onto the service. These IAM permissions are required to create a policy, create a role for SSM operation, attach the policy to this role and pass the role onto the SSM service for performing operations to create a cluster using Helios. Delete role and delete role policy permissions are required to delete the role and policy created for the SSM operation after the task is complete.
CloudFormation	UpdateStack ListStackResources CreateStack GetTemplate DeleteStack DescribeStacks DescribeStackEvents ValidateTemplate	Cohesity creates a CloudFormation template using the information provided on the UI for cluster creation, and creates a CloudFormation stack. The stack is deleted when destroy cluster command is used from the control VM. Validate cluster validates the Cloud Formation template which requires the permission "`cloudformation:ValidateTemplate`". "`cloudformation:DeleteStack`" is required while executing destroy cluster.

AWS Resource Type	Permissions	Description
EC2	AttachNetworkInterface AttachVolume AuthorizeSecurityGroupEgress AuthorizeSecurityGroupIngress CreateLaunchTemplate CreateNetworkInterface CreateSecurityGroup CreateTags CreateVolume DeleteNetworkInterface DeleteSecurityGroup DeleteTags DeleteVolume DescribeAvailabilityZones DescribeIamInstanceProfileAssociations DescribeInstanceStatus DescribeInstances DescribeNetworkInterfaces DescribeRegions DescribeSecurityGroups DescribeSubnets DescribeVolumes DescribeVpcs DetachNetworkInterface DetachVolume ModifyNetworkInterfaceAttribute ModifyVolumeAttribute ReplaceIamInstanceProfileAssociation RevokeSecurityGroupEgress RevokeSecurityGroupIngress RunInstances StartInstances StopInstances TerminateInstances	Creating a cluster requires creating AWS EC2 instances. Network Interface and volumes are created and attached to the instance to create a cluster node in a particular region. Tags are created and added to these EC2 services. Permissions to start/stop these instances are required while executing commands to start/stop cluster from the control VM. Permissions to delete all the EC2 services created while creating cluster are required while executing destroy cluster/delete VMs from the control VM. While creating a cluster using Helios, users can create a new security group and deploy the cluster using that security group instead of using the existing security group. Therefore, Cohesity needs permission to create the security group, delete the security group, authorize the security group ingress and egress, and revoke the security group ingress and egress.

AWS Resource Type	Permissions	Description
SSM	GetCommandInvocation SendCommand	Cohesity uses the SSM agent to communicate with the cluster node VMs when they come up. This communication is required to send commands to create the cluster and also to check the cluster creation status. Therefore, SSM commands are required to create a cluster using Helios.
S3	CreateBucket DeleteBucket DeleteObject GetObject PutBucketPublicAccessBlock PutObject ListBucket GetBucketObjectLockConfiguration GetLifecycleConfiguration GetObjectRetention PutObjectRetention	These permissions are required when the CloudFormation template size exceeds 51200 bytes and Cohesity needs to upload the template to an S3 bucket. `S3:GetBucketObjectLockConfiguration`, `S3:GetObjectRetention`, and `S3:PutObjectRetention` are the permissions required if you are enabling object lock. • `S3:GetBucketObjectLockConfiguration`: Gets the Object Lock configuration for a bucket. • `S3:GetObjectRetention`: Retrieves an object's retention settings. • `S3:PutObjectRetention`: Places an Object Retention configuration on an object.

Cohesity Cloud Edition on AWS Setup

The initial setup instructions for Cohesity Cloud Edition on AWS are available at Initial Cluster Setup.

Cohesity Cloud Edition on AWS Supportability

In Cohesity Cloud Edition on AWS, backed-up data is down-tiered to object storage configured with the Standard, Cold, or Archive storage tier.

Supported Sources and Workloads for Standard Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on AWS when the object storage is configured with the Standard tier.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Relational Database Service for MySQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both snapshot-based and ingest-based protection. Cross-cloud restore is not supported.
Amazon Aurora MySQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Cluster-level (With snapshot-based protection.)	Supports both snapshot-based and ingest-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service for PostgreSQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (with snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Relational Database Service for PostgreSQL	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Cluster-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service for Microsoft SQL Server	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service for Oracle	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Instance-level	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Relational Database Service Custom for Oracle			Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon Relational Database Service Custom for SQL Server	Yes	Yes	Ingest based protection: Yes Snapshot based protection: No	Ingest based protection: Yes (Only to remote cluster target.) Snapshot based protection: No	Database-level (With ingest-based protection.) Instance-level (With snapshot-based protection.)	Supports both ingest based and snapshot-based protection. Cross-cloud restore is not supported.
Amazon RDS and Amazon Aurora	Yes	Yes	No	No	Instance-level	Supports only snapshot-based protection
Amazon Redshift Cluster	Yes	Yes	Yes	Yes	Database-level	Supports only ingest-based protection.
Amazon DocumentDB	Yes	Yes	Yes	Yes	Database-level	Supports only ingest-based protection.
Amazon DynamoDB	Yes	Yes	Yes	Yes	Table-level	Supports only ingest-based protection.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Elastic Compute Cloud Instances	Yes	Yes	Yes	Yes	File-level	File-level recovery is not supported for EC2 instances with XFSv5 file system.
Amazon Simple Storage Service Bucket	Yes	Yes	Yes	Yes	File-level	Granular recovery of S3 objects with prefixes is supported.
VMware Cloud (VMC)	Yes	Yes	Yes	No	File-level	Multi-tenancy should be enabled and hybrid extender needs to be configured.
Amazon Elastic VMware Service	Yes	Yes	Yes	No	File-level	
Nutanix AHV Prism Central	Yes	Yes	Yes	Yes	File-level	
Acropolis: Acropolis Standalone Cluster	Yes	Yes	Yes	Yes	File-level	
Physical Servers	Yes	Yes	Yes	Yes	File-level	Supports both file-based and block-based backups.
Microsoft SQL Server	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based and VDI-based backups.
Microsoft 365	Yes	Yes	Yes	Yes	File-level	

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
NAS	Yes	Yes	Yes	Yes	File-level	• Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions. • NetApp: Supports only File-Runner-based backup. SnapDiff-based backup is not supported. • Isilon: Supports both File-runner and CFT-based backups.
SAP HANA on Linux (x86)	Yes	Yes	Yes	Yes	Not Applicable	
MongoDB	Yes	Yes	Yes	Yes	Not Applicable	
Cassandra	Yes	Yes	Yes	Yes	Not Applicable	

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Hadoop	Yes	Yes	Yes	Yes	File-level (Only Hadoop HDFS)	Supports protection of HDFS, Hive, and HBase objects.
Oracle Database	Yes	Yes	Yes	Yes	Not Applicable	Supports both Oracle Adapter and Oracle SBT Plug-in.
IBM DB2	Yes	Yes	Yes	Yes	Not Applicable	
PostgreSQL	Yes	Yes	Yes	Yes	Not Applicable	Supports instant restore.
SAP on Oracle for Linux	Yes	Yes	Yes	Yes	Not Applicable	
SAP on Oracle for Microsoft Windows	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Linux	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Windows	Yes	Yes	Yes	Yes	Not Applicable	
MySQL EE	Yes	Yes	No	No	Database-level	

*Cohesity Cloud Edition on AWS supports only Cloud Archive with **Incremental Forever** archival format.

- Cohesity Cloud Edition on AWS supports Write Once Read Many (WORM) for Cloud Archive with incremental forever archival format. The WORM model ensures that the objects written to the WORM-supported targets cannot be tampered with (deleted or overwritten). The archives created are immutable for the duration configured. For the list of external targets with WORM support,

For the list of supported external targets for Cloud Archive, see Supported Workflows and External Targets.

Supported Sources and Workloads for Archive and Cold Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on AWS when the object storage is configured with the Archive and Cold Tier tiers.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Amazon Elastic Compute Cloud Instances	Yes	Yes	No	No	File-level	• File-level recovery is not supported for EC2 instances with XFSv5 file system. • VM indexing disabled.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
NAS	Yes	Yes	No	No	File-level	• Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions. • NetApp: Supports only File-Runner-based backup. SnapDiff-based backup is not supported. • Isilon: Supports both File-runner and CFT-based backups.
VMware VM	Yes	Yes	No	No	File-level	VM indexing disabled.

Supported Features and Services

Features and Services	Standard Tier	Cold Tier	Archive Tier	Notes
SmartFiles Support	Yes	Yes	No	Cohesity Cloud Edition on AWS configured with Standard Tier and Cold Tier support SmartFiles for SMB, NFS, and S3 protocols, with functionality limited to View Snapshots and Replication.

Features and Services	Standard Tier	Cold Tier	Archive Tier	Notes
Marketplace Support	Yes	No	No	You can install Marketplace Apps on Cohesity Cloud Edition on AWS configured with Standard Tier. For more information on Marketplace App, see Apps
FortKnox Support	Yes	No	No	Cohesity FortKnox supports vaulting data from Cohesity Cloud Edition on AWS configured with Standard Tier to a cloud vault. For more information, see Clusters.
Multitenancy Support	Yes	No	No	Cohesity Cloud Edition on AWS configured with Standard Tier supports the multitenancy feature from 7.2.2_u1 onwards. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. For more information, see Organizations (Multitenancy) .

Manage Nodes of Cohesity Cloud Edition on AWS

For information on managing nodes of Cohesity Cloud Edition on AWS, see Manage Nodes.

Cohesity Cloud Edition on AWS Object Lock

For Cohesity Cloud Edition on AWS, you can configure Object Lock to secure the backed-up data. Enabling Object Lock prevents data from being modified, deleted, or overwritten on the object storage. You can configure Object Lock while registering the object storage as an external target to store the data.

Prerequisites for Object Lock

Before registering an object storage as an external target with Object Lock enabled, ensure the following:

- **For AWS external target (AWS S3 Bucket)**: The AWS external target must have versioning and Object Lock enabled. For more information, see Using S3 Object Lock. This property must be set at the time of target creation and cannot be enabled after creation.

> **Note:**
> Avoid setting default retention on AWS target- this is not required as Cohesity locks the objects for the duration required. Large default retentions affect garbage collection and increase costs.

Considerations for Object Lock

- For the AWS target, the objects will be locked in the Governance retention mode by default. To use the Compliance mode, contact Cohesity Support.
- When backing up data to versioned buckets, ensure that you do not edit the storage class of objects directly. Editing the storage class of objects directly will create a new version of the object while the original version remains in the previous storage tier. Object-level locks will not be carried over to the new version.
- The expiry of the object lock will match the retention of the backup snapshots. However, the object lock will not update if you modify the expiry for archived snapshots after the Protection Group has run.
- Once enabled, Object Lock cannot be disabled.
- With Object Lock enabled, the storage utilization of external targets may be higher.
- If Object Lock compliance is not met for a backup, then the backup will fail.

Cohesity Cloud Edition on AWS Considerations

Review the following considerations for Cohesity Cloud Edition on AWS:

- Cohesity does not support the following features on Cohesity Cloud Edition on AWS:
  - Direct migration from Cohesity Cloud Edition (legacy) on AWS to Cohesity Cloud Edition on AWS. As an alternative, you can deploy a new Cohesity Cloud Edition on AWS cluster and then replicate the data from Cohesity Cloud Edition (legacy) on AWS to the newly deployed Cohesity Cloud Edition on AWS.
  - CloudArchive with periodic full.
  - CloudTier workflow.
  - Data movement policy to move data to archive tiers.
  - Policy-based Cloud Spin.
- Cohesity Cloud Edition on AWS supports Write Once Read Many (WORM) for CloudArchive with incremental forever format. The WORM model ensures that the objects written to the WORM-supported AWS targets cannot be tampered with (deleted or overwritten). The archives created are immutable for the duration configured. For the list of WORM supported AWS external targets, see AWS External Targets.

- You cannot unregister an external target associated with a cluster's Storage Domain for down-tiering protected data (primary copy). Only external targets registered for archival (secondary copy) can be unregistered from the cluster.

Cohesity Cloud Edition on Azure

Cohesity Cloud Edition on Azure is a Cohesity cluster deployed in Azure, designed to offer high performance and greater cluster capacity. This solution employs advanced data deduplication and compression techniques to significantly reduce data footprints and storage costs.

This section provides the following information on Cohesity Cloud Edition on Azure:

Supported Regions for Deploying Cohesity Cloud Edition on Azure

You can deploy Cohesity Cloud Edition on Azure in all regions supported by the Azure cloud. For the list of supported regions, see Regions.

Supported Cohesity Cloud Edition on Azure Configurations

Based on the performance and cluster capacity per cluster node requirement, Cohesity supports the following cluster configurations:

Supported configurations for 7.2 and above versions:

> **Important:**
> The v2 cluster configurations (**Core SSDv2**, **Archive SSDv2**, **Performance SSDv2**, and **Performance Archive SSDv2**) are recommended for all new deployments. These configurations provide improved cost efficiency and performance benefits compared to v1 cluster configurations (**Entry**, **Core**, **Archive**, **Performance**, and **Performance Archive**).
>
> The v1 cluster configurations are included for reference purposes only. If you are currently using v1 configurations, we recommend transitioning to v2 configurations to take advantage of cost optimizations and newer disk capabilities.

Cluster Configuration	Managed Disk Type	CPU and Memory	Metadata Tier	Cache Tier	Boot Disk
NGCE Azure Entry - Config 1	Premium_LRS (SSD)	16 Core 64 GB	1 x 1 TB	2 x 1 TB	81 GB (Premium_ LRS (SSD))
NGCE Azure Core - Config 2		16 Core 64 GB	2 x 1 TB	4 x 1TB	
NGCE Azure Archive - Config 3		16 Core 64 GB	2 x 2 TB	4 x 2TB	
NGCE Azure Performance - Config 4		32 Core 128 GB	2 x 1 TB	4 x 1TB	
NGCE Azure Performance - Config 5		32 Core 128 GB	2 x 2 TB	4 x 2TB	
NGCE Azure Core v2 - Config 12	PremiumV2_ LRS (SSDv2) *	16 Core 64 GB	2 x 1 TB	2 x 1TB (ssd v2)	
NGCE Azure Archive v2 - Config 13		16 Core 64 GB	2 x 2 TB	2 x 2TB (ssd v2)	
NGCE Azure Performance v2 - Config 14		32 Core 128 GB	2 x 1 TB	2 x 1TB (ssd v2)	
NGCE Azure Performance v2- Config 15		32 Core 128 GB	2 x 2 TB	2 x 2TB (ssd v2)	

Cohesity Cloud Edition on Azure supports Standard_DSv5 (Intel) and Standard_DASv5 (AMD) series instance types.

***Considerations for PremiumV2_LRS**

- PremiumV2_LRS is supported only in limited Azure regions.
- PremiumV2_LRS requires deployment into a specific zone.

- PremiumV2_LRS has more restrictive quotas than Premium_LRS. Therefore ensure the quota for PremiumV2_LRS is high enough in the region and zone you want to deploy your Cohesity Cloud Edition on Azure.

Supported configurations for 7.1, 7.1.1, and 7.1.2 versions:

Cluster Configuration	Managed Disk Type	CPU and Memory	MetaData Tier	Cache Tier	Boot Disk
Entry	Premium_LRS (SSD)	16 Core 64 GB	1 x 1023	2x 1024	81 GB (Premium_ LRS (SSD))
Core		16 Core 64 GB	2 x 1023	4 x 1024	
Archive		16 Core 64 GB	2 x 2047	4 x 2048	
Core ssdv2	PremiumV2_ LRS (SSDv2) *	16 CPU 64 GB	2 x 1023	4 x 1024	

Cohesity Cloud Edition on Azure supports Standard_DSv5 (Intel) and Standard_DASv5 (AMD) series instance types.

***Considerations for PremiumV2_LRS**

- PremiumV2_LRS is supported only in limited Azure regions.
- PremiumV2_LRS requires deployment into a specific zone.
- PremiumV2_LRS has more restrictive quotas than Premium_LRS. Therefore ensure the quota for PremiumV2_LRS is high enough in the region and zone you want to deploy your Cohesity Cloud Edition on Azure.

Quota Requirement

The quotas required for instance types and disk types to deploy Cohesity Cloud Edition on Azure depend on the node configuration and the number of nodes in the cluster. The following table lists the quotas requirement for instance type and disk types to deploy Cohesity Cloud Edition on Azure.

Type	Quota Name	Quota consumed per node
Instance Type	Standard DSv5 Family vCPUs (Intel) or Standard DASv5 Family vCPUs (AMD)	16 or 32(performance) vCPUs

Type	Quota Name	Quota consumed per node
Premium_LRS Disk Type	Premium Managed Storage	• For 7.1, 7.1.1, and 7.1.2 versions: 1 boot disk, and 3 or 6 data disks • For 7.2 version: 1 boot disk, and 2 or 4 data disks
PremiumV2_LRS Disk Type	PremiumV2StorageDisks	• For 7.1, 7.1.1, and 7.1.2 versions: 3 or 6 data disks • For 7.2 version: 2 or 4 data disks
	PremiumV2TotalDiskSizeGB	• For 7.1, 7.1.1, and 7.1.2 versions: 3072, 6144, or 12288 GiB • For 7.2 version: 2048, 4096, or 8192 GiB

Cohesity Cloud Edition on Azure Permissions

The following table lists the permission required for deploying Cohesity Cloud Edition on Azure:

For information on backup and recovery permissions, see Azure VM.

Open the JSON (or right-click and save as a file).

Action	Resource Provider	Operation Name
Helios based cluster creation	Microsoft.Authorization	Microsoft.Authorization/locks/write

Action	Resource Provider	Operation Name
	Microsoft.Resources	Microsoft.Resources/deployments/validate/action Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/subscriptions/locations/read Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read Microsoft.Resources/deployments/operations/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/delete Microsoft.Storage/storageAccounts/listAccountSas/action Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/blobServices/containers/read
	Microsoft.Network	Microsoft.Network/networkInterfaces/delete Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkSecurityGroups/write Microsoft.Network/networkSecurityGroups/read Microsoft.Network/networkSecurityGroups/join/action

Action	Resource Provider	Operation Name
	Microsoft.Compute	Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/delete Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/runCommand/action Microsoft.Compute/disks/write Microsoft.Compute/disks/delete Microsoft.Compute/locations/runCommands/read

Action	Resource Provider	Operation Name
Cluster creation using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write

Action	Resource Provider	Operation Name
	Microsoft.Resources	Microsoft.Resources/deployments/validate/action Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/subscriptions/locations/read Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/delete Microsoft.Storage/storageAccounts/listAccountSas/action Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/blobServices/containers/read
	Microsoft.Network	Microsoft.Network/networkInterfaces/delete Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkSecurityGroups/write Microsoft.Network/networkSecurityGroups/read Microsoft.Network/networkSecurityGroups/join/action

Action	Resource Provider	Operation Name
	Microsoft.Compute	Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/delete Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/runCommand/action Microsoft.Compute/disks/write Microsoft.Compute/disks/delete Microsoft.Compute/locations/runCommands/read

Action	Resource Provider	Operation Name
Add or remove a node using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write

Action	Resource Provider	Operation Name
	Microsoft.Resources	Microsoft.Resources/deployments/validate/action Microsoft.Resources/deployments/read Microsoft.Resources/deployments/write Microsoft.Resources/subscriptions/locations/read Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/subscriptions/resourcegroups/write
	Microsoft.Storage	Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/delete Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/listKeys/action
	Microsoft.Network	Microsoft.Network/networkInterfaces/delete Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/read Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkSecurityGroups/read Microsoft.Network/networkSecurityGroups/write Microsoft.Network/networkSecurityGroups/join/action

Action	Resource Provider	Operation Name
	Microsoft.Compute	Microsoft.Compute/disks/write Microsoft.Compute/disks/delete Microsoft.Compute/virtualMachines/read Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/delete
Add or remove a node using an existing node in the cluster.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/deployments/read Microsoft.Resources/deployments/write Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/blobServices/containers/write Microsoft.Storage/storageAccounts/blobServices/containers/read Microsoft.Storage/storageAccounts/listKeys/action
	Microsoft.Network	Microsoft.Network/virtualNetworks/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action

Action	Resource Provider	Operation Name
Power on the cluster using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/blobServices/containers/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read

Action	Resource Provider	Operation Name
Power off the cluster using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/blobServices/containers/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read

Action	Resource Provider	Operation Name
Destroy the cluster using the setup utility on the Control VM.	Microsoft.Authorization	Microsoft.Authorization/locks/write
	Microsoft.Resources	Microsoft.Resources/subscriptions/resourcegroups/delete Microsoft.Resources/subscriptions/resourcegroups/read Microsoft.Resources/subscriptions/resourcegroups/write Microsoft.Resources/deployments/write Microsoft.Resources/deployments/read
	Microsoft.Storage	Microsoft.Storage/storageAccounts/write Microsoft.Storage/storageAccounts/listKeys/action Microsoft.Storage/storageAccounts/read Microsoft.Storage/storageAccounts/blobServices/containers/write
	Microsoft.Network	Microsoft.Network/virtualNetworks/read Microsoft.Network/networkInterfaces/write Microsoft.Network/virtualNetworks/subnets/join/action Microsoft.Network/networkInterfaces/join/action Microsoft.Network/networkInterfaces/read
	Microsoft.Compute	Microsoft.Compute/virtualMachines/start/action Microsoft.Compute/virtualMachines/deallocate/action Microsoft.Compute/virtualMachines/write Microsoft.Compute/virtualMachines/read Microsoft.Compute/disks/delete

Cohesity Cloud Edition on Azure Setup

The initial setup instructions for Cohesity Cloud Edition on Azure are available at Initial Cluster Setup.

Cohesity Cloud Edition on Azure Supportability

In Cohesity Cloud Edition on Azure, backed-up data is down-tiered to Azure storage configured with the Hot, Cold, or Archive storage tier.

Supported Sources and Workloads for Hot Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on Azure when the object storage is configured with the Hot tier.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
Azure Blob Storage and Azure Data Lake Storage	Yes	Yes	Yes	Yes	Yes	Yes	File-level	
Azure SQL Managed Instance	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure SQL Database	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure Table Storage	Yes	No	Yes	Yes	Yes	Yes	Table-level	
Azure Cosmos DB - Table	Yes	No	Yes	Yes	Yes	Yes	Table-level	
Azure Cosmos DB - Cassandra	Yes	No	Yes	Yes	Yes	Yes	Keyspace -level	
Azure Cosmos DB - NoSQL	Yes	No	Yes	Yes	Yes	Yes	Database-level	

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
Azure Cosmos DB - MongoDB	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure PostgreSQL	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure MySQL	Yes	No	Yes	Yes	Yes	Yes	Database-level	
Azure VMware Solution (AVS) Gen1 and Gen2	Yes	Yes	Yes	Yes	Yes	Yes	File-level	
Azure Kubernetes Service (AKS)	Yes	No	Yes	Yes	Yes	Yes	File-level	
Azure	Yes	No	Yes	Yes	Yes	Yes	File-level	
Nutanix AHV Prism Central	Yes	No	Yes	Yes	Yes	Yes	File-level	
Acropolis: Acropolis Standalone Cluster	Yes	No	Yes	Yes	Yes	Yes	File-level	

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
Nutanix Cloud Clusters on Azure	Yes	No	Yes	Yes	Yes	Yes	File-level	
Physical Servers	Yes	No	Yes	Yes	Yes	Yes	File-level	
Microsoft SQL Server (Standalone and FCI)	Yes	Yes	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based and VDI-based backups.
Microsoft 365	Yes	Yes	Yes	Yes	Yes	Yes	File-level	
NAS	Yes		Yes	Yes	Yes	Yes	File-level	Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions.
SAP HANA	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
MongoDB	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
Cassandra	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
IBM DB2	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP on Oracle for Linux	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP on Oracle for Microsoft Windows	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
Hadoop	Yes	No	Yes	Yes	Yes	Yes	Yes (Only Hadoop HDFS)	Supports protection of HDFS, Hive, and HBase objects.
PostgreSQL	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supports instant restore.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archival*	Replication	Granular Recovery	Compatibility
SAP Sybase ASE on Microsoft Windows	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supported from 7.2.2_u1 version onwards.
SAP Sybase ASE on Linux	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supported from 7.2.2_u1 version onwards.
Oracle Database	Yes	Yes	Yes	Yes	Yes	Yes	Not Applicable	Supports both Oracle Adapter and Oracle SBT Plug-in.
SAP MaxDB on Linux	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Windows	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	
SAP Sybase IQ	Yes	No	Yes	Yes	Yes	Yes	Not Applicable	Supported from 7.2.2_u1 version onwards.
MySQL EE	Yes	No	Yes	Yes	No	No	Database-level	

* Cohesity Cloud Edition on Azure supports both **Cloud Archive with periodic full** and **Cloud Archive with Incremental Forever** archival formats.

For the list of supported external targets for Cloud Archive, see Supported Workflows and External Targets.

Supported Sources and Workloads for Archive and Cold Tier

The following table provides the mapping of workflows supported for each workload in Cohesity Cloud Edition on Azure when the object storage is configured with the Archive and Cold Tier tiers.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Azure	Yes	No	Yes	Yes	Yes	Yes	File-level	VM indexing disabled.

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
NAS	Yes	No	Yes	Yes	No	No	File-level	• Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions. • NetApp: Supports only File-Runner-based backup. SnapDiff-based backup is not supported. • Isilon: Supports both File-

Workload	Supported on Cohesity Cloud Edition on Azure	Supported on Cohesity Cloud Edition on Azure GovCloud	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
								runner and CFT-based backups.
VMware VM	Yes	Yes	Yes	Yes	No	No	File-level	VM indexing disabled.

Supported Features and Services

Features and Services	Hot Tier	Cold Tier	Archive Tier	Notes
SmartFiles Support	Yes	Yes	No	Cohesity Cloud Edition on Azure configured with Hot Tier and Cold Tier support SmartFiles for SMB, NFS, and S3 protocols, with functionality limited to View Snapshots and Replication.
Marketplace Support	Yes	No	No	You can install Marketplace Apps on Cohesity Cloud Edition on Azure configured with Hot Tier. For more information on Marketplace App, see Apps
FortKnox Support	Yes	No	No	Cohesity FortKnox supports vaulting data from Cohesity Cloud Edition on Azure configured with Hot Tier to a cloud vault. For more information, see Clusters.
Multitenancy Support	Yes	No	No	Cohesity Cloud Edition on Azure configured with Hot Tier supports the multitenancy feature from 7.2.2_u1 onwards. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. For more information, see Organizations (Multitenancy) .

Manage Nodes of Cohesity Cloud Edition on Azure

For information on managing nodes of Cohesity Cloud Edition on Azure, see Manage Nodes.

Cohesity Cloud Edition on Azure Object Lock

For Cohesity Cloud Edition on Azure, you can configure Object Lock to secure the backed-up data. Enabling Object Lock prevents data from being modified, deleted, or overwritten on the object storage. You can configure Object Lock while registering the object storage as an external target to store the data.

Prerequisites for Object Lock

Before registering an object storage as an external target with Object Lock enabled, ensure the following:

- **For Azure external targets (Azure Hot Blob Storage)**: The Azure external target must have version-level immutability enabled on the storage container, which indicates that the storage account must also have version-level immutability. This property must be set at the time of target creation and cannot be enabled after creation.For more information, see Enable version-level immutability support on a container.

> **Note:**
> Avoid setting default retention on Azure target- this is not required as Cohesity locks the objects for the duration required. Large default retentions affect garbage collection and increase costs.

Considerations for Object Lock

- For the Azure target, the objects will be locked in the Locked mode by default.
- When backing up data to versioned buckets, ensure that you do not edit the storage class of objects directly. Editing the storage class of objects directly will create a new version of the object while the original version remains in the previous storage tier. Object-level locks will not be carried over to the new version.
- The expiry of the object lock will match the retention of the backup snapshots. However, the object lock will not update if you modify the expiry for archived snapshots after the Protection Group has run.
- Once enabled, Object Lock cannot be disabled.
- With Object Lock enabled, the storage utilization of external targets may be higher.
- If Object Lock compliance is not met for a backup, then the backup will fail.

Cohesity Cloud Edition on Azure Considerations

Review the following considerations for Cohesity Cloud Edition on Azure:

- Cohesity does not support the features on Cohesity Cloud Edition on Azure:
  - Direct migration from Cohesity Cloud Edition (legacy) on Azure to Cohesity Cloud Edition on Azure. As an alternative, you can deploy a new Cohesity Cloud Edition on Azure and then replicate the data from Cohesity Cloud Edition (legacy) on Azure to the newly deployed Cohesity Cloud Edition on Azure.
  - CloudTier workflow.
  - Data movement policy to move data to archive tier s.
  - Policy-based Cloud Spin.
- Cohesity Cloud Edition on Azure backend APIs are only V2 compliant; to automate workflows against Cohesity Cloud Edition on Azure ensure to call the V2 APIs and not V1 APIs.
- From 7.2 version onwards, Cohesity Cloud Edition on Azure supports Write Once Read Many (WORM) for CloudArchive with periodic full. The WORM model ensures that the objects written to the WORM-supported Azure targets cannot be tampered with (deleted or overwritten). The archives created are immutable for the duration configured. For the list of WORM supported Azure external targets, see Microsoft Azure External Targets.
- You cannot unregister an external target associated with a cluster's Storage Domain for down-tiering protected data (primary copy). Only external targets registered for archival (secondary copy) can be unregistered from the cluster.

Cohesity Cloud Edition on GCP

Cohesity Cloud Edition on GCP is a Cohesity cluster deployed in GCP, designed to offer high performance and greater cluster capacity. This solution employs advanced data deduplication and compression techniques to significantly reduce data footprints and storage costs.

This section provides the following information on Cohesity Cloud Edition on GCP:

Supported Regions for Deploying Cohesity Cloud Edition on GCP

You can deploy Cohesity Cloud Edition on GCP in all regions supported by the GCP cloud. For the list of supported regions, see Regions.

Supported Cohesity Cloud Edition on GCP Configurations

Based on the performance and cluster capacity per cluster node requirement, Cohesity supports the following cluster configuration:

<table>
<tr><th>Configuration</th><th>Instance Type (Intel/AMD)</th><th>Disk Type</th><th>Cache Tier Disks</th><th>Metadata Tier disks</th><th>Boot Disk</th><th>Reference Guides</th></tr>
<tr><td>Entry</td><td rowspan="2">n2-standard-16/n2d-standard-16</td><td rowspan="3">pd-balanced for metadata and cache tiers
pd-standard for boot disk</td><td>1 x 1 TB</td><td>1 x 1 TB</td><td rowspan="3">250 GB (pd-standard): For new deployments of cluster versions 7.3.1 and above.
81 GB (pd-standard): For cluster versions below 7.3.1.</td><td rowspan="5">DataProtect for Cloud in GCP</td></tr>
<tr><td>Core</td><td>2 x 1 TB</td><td>2 x 1 TB</td></tr>
<tr><td>Archive</td><td>2 x 2 TB</td><td>2 x 2 TB</td></tr>
<tr><td>Performance</td><td>n2-standard-32/n2d-standard-32</td><td>2 x 1 TB</td><td>2 x1 TB</td><td rowspan="2">250 GB (Hyperdisk-balanced)</td></tr>
<tr><td>Hyperdisk Core</td><td>n4-standard-16/c4-standard-16</td><td>Hyperdisk-balanced for metadata, cache, and boot disk</td><td>1 x 512 GB</td><td>2 x 511 GB</td></tr>
<tr><td>Original (7.1.2 and below versions)</td><td>n2-standard-32/n2d-standard-32</td><td>pd-ssd for metadata tier
pd-standard for cache and boot disk</td><td>1 x 2 TB</td><td>1 x 1.9 TB</td><td>81 GB (pd-standard)</td><td></td></tr>
</table>

Quota Requirement

The quotas required for instance types and disk types to deploy Cohesity Cloud Edition on GCP depend on the node configuration and the number of nodes in the cluster. The following table lists the quotas requirement for instance type and disk types to deploy Cohesity Cloud Edition on GCP.

Type	Quota Name	Quota consumed per node
Instance Type	n2-standard-16/n2-standard-16	• 16vCPUs for Entry, Core, Archive configurations. • 32vCPUs for Performance and Original configurations.
	n4-standard-16/c4-standard-16	16vCPUs
Disk Type	pd-balanced	2 to 8 TB depending on the configuration you choose.
	pd-standard	2 TB for Original configuration.
	pd-ssd	1.9 TB for Original configuration.
	hyperdisk-balanced	2 to 8 TB depending on the configuration you choose.

Cohesity Cloud Edition on GCP Permissions

The following table lists the permissions required for the GCP service account that is used for deploying Cohesity Cloud Edition on GCP.

For information on backup and recovery permissions, see GCP VM.

Action	Permissions
Helios based cluster deployment	cloudfunctions.functions.call cloudfunctions.functions.create cloudfunctions.functions.delete cloudfunctions.functions.getIamPolicy cloudfunctions.functions.setIamPolicy cloudfunctions.functions.get cloudfunctions.functions.list cloudfunctions.locations.list compute.addresses.createInternal compute.addresses.deleteInternal compute.addresses.delete compute.addresses.get compute.addresses.useInternal compute.disks.create compute.disks.delete compute.disks.get compute.disks.list compute.disks.use compute.disks.useReadOnly compute.images.get compute.images.useReadOnly compute.instances.attachDisk compute.instances.create compute.instances.delete compute.instances.get compute.instances.list compute.instances.setDeletionProtection compute.instances.setLabels compute.instances.setMetadata compute.instances.setTags compute.instances.start compute.instances.stop compute.instances.update compute.networks.get compute.networks.list compute.projects.get compute.regions.get compute.regions.list compute.subnetworks.get

Action	Permissions
	compute.subnetworks.list compute.subnetworks.use compute.subnetworks.useExternalIp compute.zones.get compute.zones.list iam.serviceAccounts.actAs resourcemanager.projects.get resourcemanager.projects.getIamPolicy run.services.getIamPolicy run.services.setIamPolicy run.services.get run.services.list vpcaccess.connectors.get vpcaccess.connectors.list vpcaccess.locations.list vpcaccess.operations.get vpcaccess.connectors.use storage.buckets.get storage.objects.create storage.objects.delete storage.objects.get storage.objects.list storage.objects.update

Action	Permissions
Cluster creation using the setup utility on the Control VM.	compute.addresses.delete compute.addresses.get compute.addresses.useInternal compute.disks.create compute.disks.delete compute.disks.get compute.disks.list compute.disks.use compute.disks.useReadOnly compute.images.get compute.images.useReadOnly compute.instances.attachDisk compute.instances.create compute.instances.delete compute.instances.get compute.instances.list compute.instances.setDeletionProtection compute.instances.setMetadata compute.instances.setTags compute.instances.start compute.instances.stop compute.networks.get compute.subnetworks.get compute.subnetworks.use storage.multipartUploads.abort storage.multipartUploads.create storage.multipartUploads.list storage.multipartUploads.listParts storage.objects.setIamPolicy storage.objects.update storage.buckets.get storage.objects.create

Action	Permissions

Add or remove a node using the setup utility on the Control VM.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions

Add or remove a node using an existing node in the cluster.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions

Power on the cluster using the setup utility on the Control VM.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions

Power off the cluster using the setup utility on the Control VM.	`compute.addresses.delete`
	`compute.addresses.get`
	`compute.addresses.useInternal`
	`compute.disks.create`
	`compute.disks.delete`
	`compute.disks.get`
	`compute.disks.list`
	`compute.disks.use`
	`compute.disks.useReadOnly`
	`compute.images.get`
	`compute.images.useReadOnly`
	`compute.instances.attachDisk`
	`compute.instances.create`
	`compute.instances.delete`
	`compute.instances.get`
	`compute.instances.list`
	`compute.instances.setDeletionProtection`
	`compute.instances.setMetadata`
	`compute.instances.setTags`
	`compute.instances.start`
	`compute.instances.stop`
	`compute.networks.get`
	`compute.subnetworks.get`
	`compute.subnetworks.use`

Action	Permissions
Destroy the cluster using the setup utility on the Control VM.	compute.addresses.delete compute.addresses.get compute.addresses.useInternal compute.disks.create compute.disks.delete compute.disks.get compute.disks.list compute.disks.use compute.disks.useReadOnly compute.images.get compute.images.useReadOnly compute.instances.attachDisk compute.instances.create compute.instances.delete compute.instances.get compute.instances.list compute.instances.setDeletionProtection compute.instances.setMetadata compute.instances.setTags compute.instances.start compute.instances.stop compute.networks.get compute.subnetworks.get compute.subnetworks.use

Cohesity Cloud Edition on GCP Setup

The initial setup instructions for Cohesity Cloud Edition on GCP are available at Initial Cluster Setup.

Manage Nodes of Cohesity Cloud Edition on GCP

For information on managing nodes of Cohesity Cloud Edition on GCP, see Manage Nodes.

Cohesity Cloud Edition on GCP Supportability

Supported Sources and Workloads

The following table provides the mapping of the workflows that are supported on the sources Cohesity Cloud Edition on GCP supports, as per the latest release:

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Bigtable	Yes	Yes	Yes	Yes	Database-level	
Firestore	Yes	Yes	Yes	Yes	Database-level	
BigQuery	Yes	Yes	Yes	Yes	Database-level	
Spanner	Yes	Yes	Yes	Yes	Database-level	
AlloyDB for PostgreSQL	Yes	Yes	Yes	Yes	Database-level	
Cloud SQL for SQL Server	Yes	Yes	Yes	Yes	Database-level	
Cloud SQL for MySQL	Yes	Yes	Yes	Yes	Database-level	
Cloud SQL for PostgreSQL	Yes	Yes	Yes	Yes	Database-level	
GCP VM	Yes	Yes	Yes	Yes	File-level	File-level recovery is not supported for GCP VMs with XFSv5 file system.
Cassandra	Yes	Yes	Yes	Yes	Not Applicable	
Hadoop	Yes	Yes	Yes	Yes	File-level (Only Hadoop HDFS)	Supports protection of HDFS, Hive, and HBase objects.
Microsoft SQL Server	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based volume-based, and VDI-based backups.

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
MongoDB	Yes	Yes	Yes	Yes	Not Applicable	
Google Cloud VMware Engine (GCVE)	Yes	Yes	Yes	No	File-level	On the GCVE, the default user (CloudOwner@gve.local) is missing the permission to access the Host -> Configuration -> Storage Partition Configuration. As a result, instant recovery fails, with the following error message: `Could not mount datastore on any host. Additional details: SOAP 1.1 fault 'ServerFaultCode [no subcode] - Permission to perform this operation was denied.` As a workaround, ensure that the default user is granted the Host -> Configuration -> Storage Partition Configuration permission.
Physical Servers	Yes	Yes	Yes	Yes	File-level	
Cohesity Views	Yes	Yes	Yes	Yes	File-level	

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
NAS(Generic NAS, NetApp, Isilon, Pure Storage FlashBlade)	Yes	Yes	Yes	Yes	File-level	Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions.
Red Hat OpenShift Virtualization	Yes	Yes	Yes	Yes	File-level	Supports only 4.14 version.
PostgreSQL	Yes	Yes	Yes	Yes	Not Applicable	Supports instant restore.
Oracle Database	Yes	Yes	Yes	Yes	Not Applicable	Supports both Oracle Adapter and Oracle SBT Plug-in.
SAP MaxDB on Linux	Yes	Yes	Yes	Yes	Not Applicable	
SAP MaxDB on Windows	Yes	Yes	Yes	Yes	Not Applicable	
SAP HANA	Yes	Yes	Yes	Yes	Not Applicable	
MySQL EE	Yes	Yes	No	No	Database-level	

*Cohesity Cloud Edition on GCP supports only Cloud Archive with **Incremental Forever** archival format.

For the list of supported external targets for Cloud Archive, see Supported Workflows and External Targets.

SmartFiles Support

Cohesity Cloud Edition on GCP supports SmartFiles for SMB, NFS, and S3 protocols, with functionality limited to View Snapshots and Replication.

> **Note:** NGCE cluster maintains a local cache of snapshots for a maximum of 7 days to optimize recovery. As a result, only snapshot versions from the last 7 days are available.

Marketplace Support

You can install Marketplace Apps on Cohesity Cloud Edition on GCP. For more information on Marketplace App, see Apps.

Multitenancy

Cohesity Cloud Edition on GCP supports the multitenancy feature. The multitenancy feature enables you to configure the multitenant environment on your Cohesity Platform and securely isolate each tenant in your environment. For more information, see Organizations (Multitenancy) .

Cohesity Cloud Edition on GCP Object Lock

On Cohesity Cloud Edition on GCP, you can enable Object Lock for GCP cloud storage buckets while registering them as external targets. A GCP external target with Object Lock enabled can be:

- Used as the main backup location by associating with the Storage Domain (primary copy).
- Used for archiving (secondary copy).

Object Lock protects backup or archive data from accidental or malicious deletion or modification. It helps meet compliance requirements by making backed-up or archived data read-only for a specified period. You can configure Object Lock while registering the cloud storage bucket as an external target.

Prerequisites for Object Lock

Before registering a GCP cloud storage bucket as an external target with Object Lock enabled, ensure the GCP cloud storage bucket has versioning and Object Retention Lock enabled. For more information, see Object Rentention Lock. The Object Retention Lock property must be set at the time of the GCP cloud storage bucket creation and cannot be enabled after creation.

Considerations for Object Lock

- Object Lock is supported only if the GCP Cloud Storage bucket and the cluster are located in the same region.
- When creating the Protection Policy, avoid setting a long retention period, as Cohesity locks the objects for the entire duration. Longer retention periods can delay garbage

collection and increase costs.

- For the GCP target, the objects will be locked in the Governance retention mode by default. To use the Compliance mode, contact Cohesity Support.
- When backing up data to versioned buckets, ensure that you do not edit the storage class of objects directly. Editing the storage class of objects directly will create a new version of the object while the original version remains in the previous storage tier. Object-level locks will not be carried over to the new version.
- The expiry of the object lock will match the retention of the backup snapshots. However, the object lock will not update if you modify the expiry for archived snapshots after the Protection Group has run.
- Once enabled, Object Lock cannot be disabled.
- With Object Lock enabled, the storage utilization of external targets may be higher.
- If Object Lock compliance is not met for a backup, then the backup will fail.

Cohesity Cloud Edition on GCP Considerations

Cohesity does not support the following features on Cohesity Cloud Edition on GCP:

- Direct migration from Cohesity Cloud Edition (legacy) on GCP to Cohesity Cloud Edition on GCP. As an alternative, you can deploy a new Cohesity Cloud Edition on GCP and then replicate the data from Cohesity Cloud Edition (legacy) on GCP to the newly deployed Cohesity Cloud Edition on GCP.
- CloudArchive with periodic full.
- Write Once Read Many (WORM) for CloudArchive.
- CloudTier workflow.
- Data movement policy to move data to archive tiers.
- Policy-based Cloud Spin.
- You cannot unregister an external target associated with a cluster's Storage Domain for down-tiering protected data (primary copy). Only external targets registered for archival (secondary copy) can be unregistered from the cluster.

Cohesity Cloud Edition on OCI

Cohesity Cloud Edition on OCI is a Cohesity cluster deployed in Oracle Cloud Infrastructure, designed to offer high performance and greater cluster capacity. This solution employs advanced data deduplication and compression techniques to significantly reduce data footprints and storage costs.

This section provides the following information on Cohesity Cloud Edition on OCI:

Supported Regions for Deploying Cohesity Cloud Edition on OCI

You can deploy Cohesity Cloud Edition on OCI in all regions supported by the Oracle Cloud Infrastructure. For the list of supported regions, see Regions.

Supported Cohesity Cloud Edition on OCI Configurations

Based on the performance and cluster capacity per cluster node requirement, Cohesity supports the following cluster configuration:

Cluster Configuration	Instance Type	Block Storage Volume Type	CPU and Memory	MetaData Tier	Cache Tier	Boot Disk
Core	VM.Standard.E4.Flex	Balanced	8 OCPUs 64 GB	2 x 511 GB	2 x 512 GB	250 GB

Cohesity Cloud Edition on OCI Permissions

Create a policy at the tenancy level with the following statements for deploying and managing Cloud Edition cluster resources.

```
oci iam policy create \ --compartment-id $TENANCY_OCID \ --name
"DPCDeploymentPolicy" \ --description "Policy for Cohesity Cloud Edition
cluster deployment automation" \ --statements '[

"Allow group DPCDeploymentAutomation to manage instances in tenancy",

"Allow group DPCDeploymentAutomation to manage instance-configurations in
tenancy",

"Allow group DPCDeploymentAutomation to manage volumes in tenancy",

"Allow group DPCDeploymentAutomation to manage volume-attachments in
tenancy",

"Allow group DPCDeploymentAutomation to read virtual-network-family in
tenancy",

"Allow group DPCDeploymentAutomation to use subnets in tenancy",

"Allow group DPCDeploymentAutomation to use network-security-groups in
tenancy",

"Allow group DPCDeploymentAutomation to use vnics in tenancy",

"Allow group DPCDeploymentAutomation to read instance-images in tenancy",
```

```
"Allow group DPCDeploymentAutomation to read app-catalog-listing in
tenancy",
"Allow group DPCDeploymentAutomation to inspect identity-providers in
tenancy"
```

Cohesity Cloud Edition on OCI Setup

The initial setup instructions for Cohesity Cloud Edition on OCI are available at Initial Cluster Setup.

Cohesity Cloud Edition on OCI Supportability

Supported Sources and Workloads

The following table provides the mapping of the workflows that are supported on the sources supports, as per the latest release:

Workload	Backup	Recovery	Cloud Archive*	Replication	Recovery Level	Compatibility
Physical Servers	Yes	Yes	Yes	Yes	File-level	
NAS (Generic NAS, NetApp, Isilon, Pure Storage FlashBlade)	Yes	Yes	Yes	Yes	File-level	Supports NAS volumes with NFSv3, NFSv4.1, SMB v1, v2.x, and v3 versions.
Microsoft 365	Yes	Yes	Yes	Yes	File-level	
Microsoft SQL Server	Yes	Yes	Yes	Yes	File-level	Supports MS SQL File-based and VDI-based backups.
Oracle Database	Yes	Yes	Yes	Yes	Not Applicable	Supports Oracle Adapter-based protection.

Chapter 35: Tools

This chapter includes the following topics:

Using the Cohesity CLI 4981
Cluster Health Check 5137
Support Toolbox 5140
Using the Secure Shell 5140
Cohesity REST API 5143

Using the Cohesity CLI

You can use the Cohesity Command Line Interface (CLI) locally or remotely to manage your Cohesity cluster.

Run the Cohesity CLI Locally

Before you begin:

Cohesity restricts SSH access to the host operating system shell, which means you can run custom scripts and private tools only on Secure Shell and not on Cohesity CLI (iris CLI). Cohesity recommends evaluating the security enhancements mentioned in the Secure Shell Considerations section and deciding whether to enable Secure Shell or leverage the Cohesity Cluster and Helios APIs.

Follow these steps to run the CLI commands on Secure Shell:

1. Get the IP address for any node of the Cohesity cluster. You can get the list of IP addresses from the Cohesity Cluster UI by selecting **Settings > Summary** and then selecting the **Nodes** tab.
2. Use the IP address to log into the node as the support user.

> **Note:** Make sure that you have enabled the Support user on the Cohesity cluster. For more information, see Enable Support User for Local Shell Access.

Command

```
ssh support@<node IP address>
```

Example

```
ssh support@111.1.1.1
```

3. When prompted, enter the password of the support user account.

```
support@111.1.1.1's password:
```

Once your password is successfully authenticated, the prompt changes, and the `support@restricted` prompt is displayed.

```
[support@restricted-cluster-1-node-4 ~]\>
```

4. Run the command to start the Cohesity CLI console as a local user.

```
[support@restricted-cluster-1-node-4 ~]$ iris_cli
```

You are prompted to enter the username and password. Enter the credentials of the local user account you use to log into the Cohesity cluster's User Interface.

Sample Output

```
[support@restricted-cluster-1-node-4 ~]$ iris_cli
```

```
Username: john
```

```
Password: .....
```

5. Once the local user account's password is successfully authenticated, the Cohesity CLI console opens. You can now run any CLI command.

Sample Output

```
john@111.1.1.1>cluster status
```

```
Username: john
```

```
Password: .....
```

If you only enter the CLI commands at the prompt and then enter your credentials when prompted (as shown in the example above), you're in the Interactive Mode.

Login with Active Directory User

You can log in to `iris_cli` with an Active Directory (AD) user by specifying the AD domain using the `-domain` flag.

Prerequisites

The cluster must be joined to one or more Active Directory domains. For more information, see Active Directory (AD) and Join Active Directory.

Steps to Log In with an AD User

1. After starting the `iris_cli`, use the `-domain` switch to specify the AD domain:

   ```
   iris_cli -domain <AD-domain-name>
   ```

2. When prompted, provide the username (in the format `<username>`) and the password for the AD user.

   Example

   ```
   Username: john
   ```

   ```
   Password: .....
   ```

> **Note:** Usernames must be provided as-is, as `iris_cli` does not support the NetBIOS (DOMAIN\username) or UPN format (username@domain). `iris_cli` also does not support login for SSO-based users (SAML).

Enable Support User for Local Shell Access

You need to use the **support** user account to log into the Cohesity bash shell using SSH.

By default, the support user account is disabled on the Cohesity node that is part of a Cohesity cluster. You need to enable the support user account by setting a password for the user account in the Cohesity Dashboard. Also, you can enable Linux sudo access for the support user account.

Set the Support User Password for the First Time

Using the Cohesity Dashboard or Cohesity CLI, you can set the password for the support user. You cannot retrieve the password once it is set. Therefore, ensure that you set a strong password that is longer and easy to remember. In case, if you want to reset the support user password and have forgotten the existing password, then contact Cohesity Support.

To set the support user password:

In the Cohesity Dashboard, select **Settings** > **Access Management**.

In the **Support user** section, under **Password**, click **Set Password** and define the new password for the support user account.

Also, from the Cohesity CLI, you can use the following `iris_cli` command to change the support user password:

```
iris_cli user linux-user linux-username=support linux-password=<password>
```

Enable Linux Sudo Access

When operating within the Host OS, without sudo access, the support user has read-only access privileges and limited read-write access privileges. To perform advanced operations within the Host OS that require privileged access, you need to enable Linux sudo access for the support user account.

Use the following `iris_cli` command to enable Linux sudo access for the support user:

```
iris_cli user linux-sudo-access enable=true duration-in-hours=xxx
```

> **Note:** You can enter a duration within the range of 1 to 720 hours.

To disable Linux sudo access for the support user, use the following `iris_cli` command:

```
iris_cli user linux-sudo-access enable=false
```

> **Note:** Apart from the `admin` user, if you use a local user account to enable Linux sudo access for the support user, ensure that you configure the local user account with the `Manage Linux user sudo access` privilege.

Update cohesity_console User Password

By default, the root user account is disabled on the Cohesity cluster. For operations that require root access, you can use the cohesity_console user account.

To update the cohesity_console user password, use the following iris_cli command:

```
iris_cli user linux-user linux-username=cohesity_console linux-pass-
word=<password>
```

Explore the Cohesity CLI

To view the top-level Cohesity CLI commands and help about the commands, enter `help`.

```
admin>help
```

To view the operations available for a command and help about the operations, enter the command, a space, and `help`.

```
admin>node help
```

To view the parameters of an operation and get help about the parameters, enter the command, a space, the operation, and `help`.

```
admin>node list help
```

To list the top-level commands that start with specified characters using command-line completion, enter some characters, and a tab. For example:

```
admin>vie<tab>
view view-box
```

To list the operations available for a command using command-line completion, enter the command, a space, and a tab. For example:

```
admin>node <tab>
```

To list the operations using command-line completion, enter the command, a space, one or more characters of the sub-command, and a tab.

```
admin>node l<tab>
list list-packages ls ls-free-nodes ls-pkgs
```

To list the parameters using command-line completion, enter the command, a space, the operation, a space, and a tab. For example:

```
admin>node list<tab>
list list-packages
admin>node list
```

To see a list of available Nodes on the Cohesity cluster.

```
admin>node list
```

To see the status of the Cohesity cluster.

```
admin>cluster status
```

To get details about the Cohesity cluster.

```
admin>cluster info
```

> **Note:** This is a small sample of the commands available from the Cohesity CLI. Explore the Cohesity CLI help for more options.

Key Sequences to Edit the Cohesity CLI Text

The following key sequences can be used to edit text in the Cohesity CLI prompt.

Key Sequence	Action
↑ (Up Arrow)	Show the previous command in the Cohesity CLI history.
↓ (Down Arrow)	Show the next command in the Cohesity CLI history.
Tab	Attempt command completion.
Control-u	Erase the current line of text.
Control-w	Delete the last word in the current line of text.
Control-a	Go to the first character in the current line of text.
Control-e	Go to the last character in the current line of text.
Control-b	Go back one word in the current line of text.
Control-k	Delete the rest of the text in the current line.

Use the Bash Shell on a Cohesity Node

This section describes how to access the Cohesity bash shell for troubleshooting and configuration.

1. To log into the Cohesity cluster, use the DNS Name or IP that is used to browse into the Cohesity Dashboard. The bash shell can also be accessed from the IPMI remote console (IPMI, iLO, CIMC, or iDRAC, depending on the hardware vendor) of the node.
2. Use PuTTY (shown below), an SSH client, or IPMI remote console to log in to the cluster. Use the username `support` to log in to the cluster and specify the IP address.

   On a Mac, Unix, or Linux computer, use the ssh command:

   ```
   ssh support@<hostname>
   ```

   Example:

   ```
   ssh support@192.168.1.100
   ```

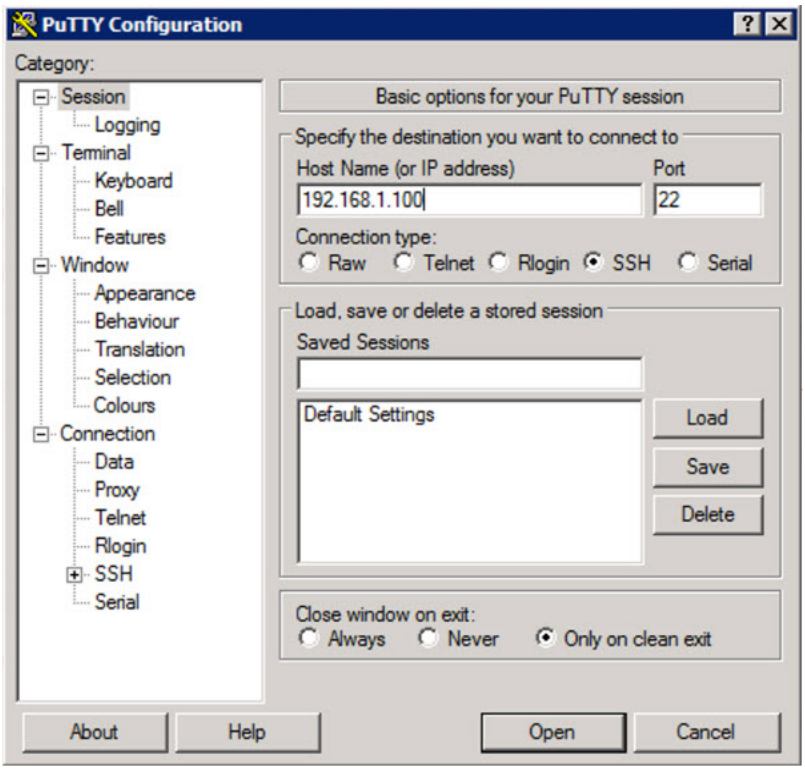

3. When prompted, enter the password.

> **Note:** This password refers to the console password and not the Cohesity UI credentials. You can provide the default password if the node is not part of a Cohesity cluster. For the default password of the support user account, contact Cohesity Support.

```
[support@192.168.1.100's password:
```

The `bash shell` prompt is displayed.

Run the Cohesity CLI Remotely

You can also run the Cohesity CLI on a Windows, macOS, or Linux system and remotely connect to the Cohesity cluster.

> **Note:** The user using a Mac with Apple Silicon needs to install Rosetta 2 to run the Cohesity CLI.

To download and start the remote version of the Cohesity CLI:

1. Download the appropriate CLI executable for your Operating System using one of the following methods:
   - Download from the Cohesity UI:
     a. Click **Help** (?) in the top right corner and click **Download CLI**.
     b. Find the appropriate CLI for your Operating System and click **Download** (↓).
   - Download over HTTP:

     Download a Cohesity CLI executable from one of the following URLs on the Cohesity cluster:

Windows	On Windows, you can download the Cohesity CLI using the following command in PowerShell or CMD: wget http://<*nodeIP*>/files/bin/cli/windows-x86_64/iris_cli.exe, where *<nodeIP>* is the IP address of one of the Nodes in the Cohesity cluster. Example: wget http://192.168.1.100/files/bin/cli/windows-x86_64/iris_cli.exe
macOS	On Mac, perform the following steps to download the Cohesity CLI: 1. curl http://<*nodeIP*>/files/bin/cli/darwin-x86_64/iris_cli -o iris_cli, where *<NodeIP>* is the IP address of one of the nodes on the Cohesity cluster. Example: curl http://198.51.100.12/files/bin/cli/darwin-x86_64/iris_cli -o iris_cli 2. chmod 755 iris_cli
Linux	On Linux, perform the following steps to download the Cohesity CLI: 1. wget http://<*nodeIP*>/files/bin/cli/linux-x86_64/iris_cli, where *<NodeIP>* is the IP address of one of the nodes on the Cohesity cluster. Example: wget http://198.51.100.12/files/bin/cli/linux-x86_64/iris_cli 2. chmod 755 iris_cli

2. Start an Operating System terminal or command prompt.
3. Change directories to the location of the CLI executable.
4. Start the CLI with the `-help` option to view all the start-up options.

Windows	`.\iris_cli.exe -help`
macOS	`./iris_cli -help`
Linux	`./iris_cli -help`

5. Specify the options to connect the Cohesity CLI remotely to the Cohesity cluster, as shown in the following example. Provide a cluster username and password.

   Specify one of the IP addresses for a Node in the Cohesity cluster with the `server` option:

   ```
   ./iris_cli -server=198.51.100.12
   ```

   When prompted, enter the username and password.

   Sample Output:

   ```
   ./iris_cli -server=198.51.100.12
   Username: admin
   Password: *****
   ```

   > **Note:** If you change the default port number (443) to connect to the Cohesity cluster, you must also specify the port number on the command line, for example: `-port="445"`

   The Cohesity CLI prompt appears. You can now enter the Cohesity CLI commands. For more information, see Explore the Cohesity CLI.

List of Cohesity CLI Commands

Command Name	Short Name	Description
active-directory	ad	To run any Active Directory related operations.
airgap-config		To run any Airgap config related operations.
alert		To run any alert related operations.
apps		To run any apps related operations.

Command Name	Short Name	Description
audit-logs		To run audit logs related operations
baseos-patch		To run any base OS patch related operations.
bond		To run any bond related operations.
chassis		To manage chassis in cluster.
cipher		To run any cipher related operations.
clientcsr		To run any clientcsr related operations.
cluster		To run any Cluster related operations.
csr		To run any csr related operations.
custom-reporting		To run any custom reporting related operations.
disk		To run any Disk related operations.
firewall-ipset		To run any Firewall ipset related operations.
firewall-profile		To run any Firewall profile related operations.
group		To run any Group related operations.
hyx-vlan		To run any Bifrost Vlan related operations.
interface-group		To run network interface group related operations.
interface		To run interface related operations.
ip		To run ip related operations.
ipmi		To run IPMI related operations on local and other nodes of the cluster.
kms		To run any KMS related operations.
ldap		To run any LDAP related operations.

Command Name	Short Name	Description
mfa-config		To run any cluster MFA config related operations.
network-reset		network configurations after doing a "set-to-factory-state" operation. These operations would not affect the data in the cluster.
nfs-share		To run any NFS share related operations.
node-csr		To run any node csr related operations.
node-group		To run node group related operations.
node		To run any Node related operations.
patch		To run any Patch related operations.
privilege		To run any Privilege related operations.
protection-job		To run any Protection Job related operations.
qos-profile		To run any Qos profile related operations.
rack		To manage racks in cluster.
remote-cluster	rc	To run any remote Cluster registration related operations.
role		To run any Role related operations.
route		To run any Static Route related operations. A Static Route designates which VLAN and Network Interfaces to use for communication with a particular destination network.
routing		Routing commands.
security-config		To run any security config related operations.
service-identity		To run any Service Identity related operations.

Command Name	Short Name	Description
set-option	set	To set or override existing CLI options at runtime.
snmp		To run any SNMP related operations
storage-domain	sd	To run any Storage Domain related operations.
support-bundle	bundle	To generate support bundle.
support-server	support	To run Support server and Support channel related operations.
swift		To run any Swift related operations.
sysctl		Command to update sysctl configuration for nodes.
syslog-server	syslog	To run any Syslog Server related operations.
system-app		To run any System app related operations.
trusted-cas		To run any trusted CA related operations.
uda-connector		To run any UDA connector related operations
user-quotas		To run any user quotas related commands.
user		To run any User related operations.
vault	vt	To run any Vault related operations.
view		To run any View related operations.
vlan		To run any Vlan related operations.

active-directory

To run any Active Directory related operations.

Operation	Short Name	Description
add-ca-path		To add Kerberos capath.
create	add	To create an account for the cluster in Active Directory.
delete	rm	To delete the cluster account from Active Directory.
enable-allowed-trusted-domains		Enable allowed trusted domains to only use them for authentication.
enable-trusted-domains-discovery		Enables or disables the discovery of Trusted Domains of an Active Directory.
list	ls	To list the Active Directory information.
list-ca-paths		To list the existing Kerberos capaths.
list-centrify-zones		To list the centrify zones of an active directory.
list-kerberos-closest-dc-locator-plugin		Lists the Windbind based closest Kerberos server dc locator plugin setting.
rediscover-trusted-domains		To re-discover trusted domains of a primary Active Directory domain.
remove-ca-path		To remove Kerberos capath.
set-allowed-trusted-domains		Sets the allowed trusted domains for an active directory
set-id-mapping-info		Sets the user id mapping info for an active directory domain.
set-ignored-trusted-domains		Sets the trusted domains to be ignored during trusted domain discovery.
set-kerberos-closest-dc-locator-plugin		Sets the Windbind based closest Kerberos server dc locator plugin setting. The change will be applied after bridge service is restarted.
set-ldap-provider		Sets the LDAP provider to Active Directory.
set-preferred-domain-controllers		Sets the preferred domain controllers to an Active Directory.
update-ldap-bind-credentials		To update the Ldap bind credentials of an Active Directory.
update-machine-accounts		To update the machine accounts of an active directory.

active-directory add-ca-path

To add Kerberos capath.

Parameter	Type	Required	Description
auth-domain	string	yes	FQDN of the client Auth domain.
cluster-domain	string	yes	FQDN of the cluster Active Directory domain.
intermediate-domain	string	yes	FQDN of the intermediate Active Directory domain.

active-directory create

To create an account for the cluster in Active Directory.

Parameter	Type	Required	Description
domain-controllers	list of strings	no	Comma separated values for the domain controllers of an Active Directory domain, these domain controllers will be mapped to Active Directory's domain name.
domain-name	string	yes	Domain name of the Active Directory.
enable-trusted-domain	bool	no	Enable Trusted Domain discovery.
machine-accounts	list of strings	yes	Comma separated values for Machine accounts to use with Active Directory.
ou	string	no	Org Unit for Active Directory
overwrite-existing-accounts	bool	no	Overwrites existing machine accounts in this active directory domain.
password	string	yes	Password of the User.
user-name	string	yes	Username of the User.
workgroup	string	no	Workgroup supported in Active Directory

active-directory delete

To delete the cluster account from Active Directory.

Parameter	Type	Required	Description
domain-name	string	yes	Domain name of the Active Directory.
password	string	yes	Password of the User.
user-name	string	yes	Username of the User.

active-directory enable-allowed-trusted-domains

Enable allowed trusted domains to only use them for authentication.

Parameter	Type	Required	Description
domain-name	string	yes	Domain name of the Active Directory.
enable	bool	yes	Whether to enable allowed trusted domains.

active-directory enable-trusted-domains-discovery

Enables or disables the discovery of Trusted Domains of an Active Directory.

Parameter	Type	Required	Description
domain-name	string	yes	Domain name of the Active Directory.
enable-trusted-domain	bool	yes	Enable Trusted Domain discovery.

active-directory list-centrify-zones

To list the centrify zones of an active directory.

Parameter	Type	Required	Description
domain-name	string	yes	Domain name of the Active Directory.

active-directory rediscover-trusted-domains

To re-discover trusted domains of a primary Active Directory domain.

Parameter	Type	Required	Description
domain-name	string	yes	Domain name of the Active Directory.

active-directory remove-ca-path

To remove Kerberos capath.

Parameter	Type	Required	Description
auth-domain	string	yes	FQDN of the client Auth domain.
cluster-domain	string	yes	FQDN of the cluster Active Directory domain.
intermediate-domain	string	yes	FQDN of the intermediate Active Directory domain.

active-directory set-allowed-trusted-domains

Sets the allowed trusted domains for an active directory.

Parameter	Type	Required	Description
allowed-trusted-domains	list of strings	no	Comma separated values for existing and new trusted domain to be allowed.
domain-name	string	yes	Domain name of the Active Directory.

active-directory set-id-mapping-info

Sets the user id mapping info for an active directory domain.

Parameter	Type	Required	Description
domain-name	string	yes	Domain name of the Active Directory.
fallback-gid	int	no	Specify the fallback fixed unix Gid.
fallback-gid-attr-name	string	no	Specify the fallback custom gid attribute name.

Parameter	Type	Required	Description
fallback-mapping-type	string	no	Specify the mapping type. Can be one of the following: kRid, kFixed.
fallback-uid	int	no	Specify the fallback fixed unix Uid.
fallback-uid-attr-name	string	no	Specify the fallback custom uid attribute name.
fallback-zone-centrify-schema	string	no	Specify the schema type of the fallback centrify zone. Should be one of: kCentrifyDynamicSchema_1_0, kCentrifyDynamicSchema_2_0, kCentrifyDynamicSchema_3_0, kCentrifyDynamicSchema_5_0, kCentrifySfu_3_0, kCentrifySfu_3_0_V5, kCentrifySfu_4_0, kCentrifyCdcRfc2307, kCentrifyCdcRfc2307_2, kCentrifyCdcRfc2307_3.
fallback-zone-description	string	no	Specify a description for this fallback centrify zone.
fallback-zone-distinguished-name	string	no	Specify a distinguished name for this fallback centrify zone.
gid	int	no	Specify the fixed Unix Group ID.
gid-attr-name	string	no	Specify the custom Group ID attribute name.
mapping-type	string	no	Specify the mapping type. Can be one of the following: kRid, kRfc2307, kSfu30, kCentrify, kFixed, kCustomAttributes, kLdapProvider, kNis.
remove-user-id-mapping-info	bool	no	Removes the current mapping between this Active Directory domain user to the user Id mapping information it is currently mapped to.
uid	int	no	Specify the fixed Unix User ID.
uid-attr-name	string	no	Specify the custom Unix User ID attribute name.
unix-root-user-sid	string	no	The Sid of the Active Directory domain user mapped to Unix root user.

Parameter	Type	Required	Description
zone-centrify-schema	string	no	Specify the schema type of the centrify zone. Should be one of: kCentrifyDynamicSchema_1_0, kCentrifyDynamicSchema_2_0, kCentrifyDynamicSchema_3_0, kCentrifyDynamicSchema_5_0, kCentrifySfu_3_0, kCentrifySfu_3_0_V5, kCentrifySfu_4_0, kCentrifyCdcRfc2307, kCentrifyCdcRfc2307_2, kCentrifyCdcRfc2307_3.
zone-description	string	no	Specify a description for this centrify zone.
zone-distinguished-name	string	no	Specify a distinguished name for this centrify zone.

active-directory set-ignored-trusted-domains

Sets the trusted domains to be ignored during trusted domain discovery.

Parameter	Type	Required	Description
clear-ignored-trusted-domains	bool	no	Removes the current mapping between this Active Directory to excluded trusted domains it is currently mapped to.
domain-name	string	yes	Domain name of the Active Directory.
ignored-trusted-domains	list of strings	no	Comma separated values for existing and new trusted domain to be ignored during trusted domain discovery.

active-directory set-kerberos-closest-dc-locator-plugin

Sets the Windbind based closest Kerberos server dc locator plugin setting. The change will be applied after bridge service is restarted.

Parameter	Type	Required	Description
enable-plugin	bool	yes	To set the Kerberos closest dc locator plugin on the cluster.

active-directory set-ldap-provider

Sets the LDAP provider to Active Directory.

Parameter	Type	Required	Description
clear-ldap-mapping	bool	no	Removes the current mapping between this Active Directory to the LDAP provider it is currently mapped to.
domain-name	string	yes	Domain name of the Active Directory.
ldap-id	int	no	Id of the LDAP provider.

active-directory set-preferred-domain-controllers

Sets the preferred domain controllers to an Active Directory.

Parameter	Type	Required	Description
active-directory-domain-name	string	yes	Domain name of the Active Directory.
clear-preferred-domain-controller	bool	no	Removes the current mapping between this Active Directory to the Preferred domain controller it is currently mapped to.
domain-controllers	list of strings	no	Comma separated values for the domain controllers of an Active Directory domain.
trusted-domain-name	string	no	Trusted domain name of an Active Directory to which preferred domain controllers are mapped to.If trusted domain name is not specified, then preferred domain controllers are mapped to Active Directory domain name.

active-directory update-ldap-bind-credentials

To update the Ldap bind credentials of an Active Directory.

Parameter	Type	Required	Description
clear-credentials	bool	no	To clear the Ldap bind credentials of an Active Directory.
domain-name	string	yes	Domain name of the Active Directory.
password	string	no	Password of the User.
user-name	string	no	Username of the User.

active-directory update-machine-accounts

To update the machine accounts of an active directory.

Parameter	Type	Required	Description
domain-name	string	yes	Domain name of the Active Directory.
machine-accounts	list of strings	yes	Comma separated values for Machine accounts to use with Active Directory.
overwrite-existing-accounts	bool	no	Overwrites existing machine accounts in this active directory domain.
password	string	yes	Password of the User.
user-name	string	yes	Username of the User.

airgap-config

To run any Airgap config related operations.

Operation	Short Name	Description
add-exception		To add exception firewall-profile for Airgap config.
list	ls	To list the Airgap configuration.
remove-exception		To remove exception firewall-profile for Airgap config.
update		To update an existing configuration for Airgap.

airgap-config add-exception

To add exception firewall-profile for Airgap config.

Parameter	Type	Required	Description
exception-profile	string	yes	Traffic matching these profiles are allowed when cluster is in Airgap mode.

airgap-config remove-exception

To remove exception firewall-profile for Airgap config.

Parameter	Type	Required	Description
exception-profile	string	yes	Traffic matching these profiles are allowed when cluster is in Airgap mode.

airgap-config update

To update an existing configuration for Airgap.

Parameter	Type	Required	Description
airgap-status	string	yes	Airgap status as Enable or Disable.
exception-profile	list of strings	no	Traffic matching these profiles are allowed when cluster is in Airgap mode.

alert

To run any alert related operations.

Operation	Short Name	Description
list	ls	To list the alerts.

alert list

To list the alerts.

Parameter	Type	Required	Description
end-time	datetime	no	End time of the time range in "YYYY-MM-DDTHH:mm" format. If specified, only the alerts that were raised before the end time will be fetched.
ids	list of strings	no	Comma separated list of alert IDs.
max-num	int	no	Maximum number of alerts to fetch.
resolution-ids	list of ints	no	Comma separated list of alert resolution IDs.

Parameter	Type	Required	Description
start-time	datetime	no	Start time of the time range in "YYYY-MM-DDTHH:mm" format. If specified, only the alerts that were raised after the start time will be fetched.
state	string	no	State of the alert. Valid values are NOTE, OPENand RESOLVED.
types	list of ints	no	Comma separated list of alert type IDs.

apps

To run any apps related operations.

Operation	Short Name	Description
update-apps-settings		Update apps settings.

apps update-apps-settings

Update apps settings.

Parameter	Type	Required	Description
allow-external-traffic	bool	no	Allow external traffic for apps.
apps-mode	"bare", "vm", "disable"	no	Edit apps mode. It can be either one of ["bare", "vm", "disable"]
marketplace-apps-mode	"bare", "vm", "disable"	no	Edit apps mode. It can be either one of ["bare", "vm", "disable"]

audit-logs

To run audit logs related operations

Operation	Short Name	Description
list	ls	List cluster audit logs

audit-logs list

List cluster audit logs

Parameter	Type	Required	Description
action	list of strings	no	Specifies a list of actions, only audit logs containing these actions will be returned.
category	list of strings	no	Specifies a list of categories/entity types, only audit logs containing these entity types will be returned.
end-time	int	no	Specifies a unix timestamp in microseconds, return only audit logs made defore the end time.
start-time	int	no	Specifies a unix timestamp in microseconds, return only audit logs made after the start time.

baseos patch

To run any base OS patch related operations.

Operation	Short Name	Description
apply		To apply the base OS patch on cluster.
download		To download the base OS patch on cluster.
get-list		To get the list of the base OS patches.
get-logs		To get the status and logs of the given patch operations.
remove		To cleanup the base OS patch files on cluster.

baseos-patch apply

To apply the base OS patch on cluster.

Parameter	Type	Required	Description
name	string	yes	Name of the base OS patch.

baseos-patch download

To download the base OS patch on cluster.

Parameter	Type	Required	Description
url	string	yes	Download the URL of the base OS patch.

baseos-patch get-logs

To get the status and logs of the given patch operations.

Parameter	Type	Required	Description
name	string	yes	Name of the base OS patch.

baseos-patch remove

To cleanup the base OS patch files on cluster.

Parameter	Type	Required	Description
force-remove	bool	no	Remove the base OS patch files even for the in-progress patch.
name	string	yes	Name of the base OS patch.

bond

To run any bond related operations.

Operation	Short Name	Description
create	add	To create a bond interface.
delete	rm	To delete a bond interface.
update	edit	To update a bond interface.

bond create

To create a bond interface.

Parameter	Type	Required	Description
bond-interface	string	yes	Name of the bond interface.
bond-slave-interface	list of strings	yes	Name of the bond slave interfaces separated by commas

bond delete

To delete a bond interface.

Parameter	Type	Required	Description
bond-interface	string	yes	Name of the bond interface.

bond update

To update a bond interface.

Parameter	Type	Required	Description
bond-interface	string	yes	Name of the bond interface.
bond-slave-interface	list of strings	yes	Name of the bond slave interfaces separated by commas

chassis

To manage chassis in cluster.

Operation	Short Name	Description
list	ls	To get all chassis or one chassis information in cluster.
update	edit	To update a chassis information. Set rack id with 0 will unassign the chassis to a rack.

chassis list

To get all chassis or one chassis information in cluster.

Parameter	Type	Required	Description
id	int	no	The id of a chassis.

chassis update

To update a chassis information. Set rack id with 0 will unassign the chassis to a rack.

Parameter	Type	Required	Description
id	int	yes	The id of a chassis.
name	string	no	Name of the chassis for the Node.
rack-id	string	no	The id of a rack which the chassis belongs to.

cipher

To run any cipher related operations.

Operation	Short Name	Description
list	ls	To list the enabled and disabled ciphers on the cluster.
list-object-store-ciphers		To list object store ciphers on the cluster.
update	edit	To modify the enabled and disabled ciphers on the cluster.
update-object-store-ciphers		To enable or disable a list of object store ciphers on the cluster. Bridge needs to be restarted to make the change in effect.
update-gflag		To enable or disable a list of object store ciphers on the cluster using cluster gflags. Bridge needs to be restarted to make the change in effect.

The list of ciphers on the clusters are separated by single space in the response. You can refer to the sample response below.

```
admin@127.0.0.1> cipher list
ENABLED CIPHERS             : [TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_RSA_WITH_AES
_256_GCM_SHA384 TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA TLS_ECD
HE_RSA_WITH_AES_256_CBC_SHA TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA]
DISABLED CIPHERS            : []
```

cipher update

To modify the enabled and disabled ciphers on the cluster.

Parameter	Type	Required	Description
ciphers	list of strings	yes	A list of ciphers to be set as enabled or disabled on the cluster. The 'enable' parameter specifies whether to enable or disable these ciphers.
enable	bool	yes	Boolean value to specify whether to enable or disable a list of ciphers. True = enable, False = disable.

The list of ciphers on the clusters are separated by single space in the response. You can refer to the sample response below.

```
admin@127.0.0.1> cipher list
ENABLED CIPHERS              : [TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_RSA_WITH_AES
_256_GCM_SHA384 TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA TLS_ECD
HE_RSA_WITH_AES_256_CBC_SHA TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA]
DISABLED CIPHERS             : []
```

cipher update-object-store-ciphers

To enable or disable a list of object store ciphers on the cluster. Bridge needs to be restarted to make the change in effect.

Parameter	Type	Required	Description
ciphers	list of strings	yes	A list of ciphers to be set as enabled or disabledon the cluster. The 'enable' parameter speciifies whether to enable ordisable these ciphers.
enable	bool	yes	Boolean value to specify whether to enable ordisable a list of ciphers. True = enable, False = disable.

The list of ciphers on the clusters are separated by single space in the response. You can refer to the sample response below.

```
admin@127.0.0.1> cipher list
ENABLED CIPHERS              : [TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_RSA_WITH_AES
_256_GCM_SHA384 TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA TLS_ECD
HE_RSA_WITH_AES_256_CBC_SHA TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA]
DISABLED CIPHERS             : []
```

cipher update-gflag

To enable or disable a list of object store ciphers on the cluster using cluster gflags. Bridge needs to be restarted to make the change in effect.

Parameter	Type	Required	Description
service-name	string	yes	The name of the cluster service whose gflag is to be updated. To configure TLS ciphers on the S3-compatible object store (ports 3000 and 3001), specify bridge.
gflag-name	string	yes	The name of the gflag to update. To disable TLS 1.3 cipher suites on the S3-compatible object store, specify bridge_s3_disabled_cipher.
gflag-name	string	yes	The value to assign to the specified gflag. For bridge_s3_disabled_cipher, provide a comma-separated list of cipher names to disable. TLS 1.3 cipher names use the standard format (for example, TLS_AES_128_GCM_SHA256). TLS 1.2 cipher names use the OpenSSL short-name format (for example, ECDHE-RSA-AES128-GCM-SHA256).
reason	string	yes	A brief description of why the gflag is being changed. This value is recorded in the cluster audit log.
effective-now	bool	no	When set to true, the gflag value is applied to the running service instance immediately. Default is false. For bridge_s3_disabled_cipher, a Bridge service restart is still required after setting this flag for the updated cipher list to take effect on active TLS connections.

The list of ciphers on the clusters are separated by single space in the response. You can refer to the sample response below.

```
admin@127.0.0.1> cipher list
ENABLED CIPHERS              : [TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_RSA_WITH_AES
_256_GCM_SHA384 TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA TLS_ECD
HE_RSA_WITH_AES_256_CBC_SHA TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA]
DISABLED CIPHERS             : []
```

clientcsr

To run any clientcsr related operations.

Operation	Short Name	Description
create	add	To create csrs for client and server.
import		To import a certificate by csr.

clientcsr create

To create csrs for client and server.

Parameter	Type	Required	Description
city	string	yes	The locality attribute, which is part of the distinguished name definition. It is used to identify the city where the company is located or the Cluster is installed.
country	string	yes	The country attribute, which is part of the distinguished name definition. It is used to identify the country where the state is located. It is specified as two letter code defined by the ISO standard.
dns-names	list of strings	no	An alternative subject name component to be included in the certificate. It is used to identify the ways the Cluster will be accessed. It is given as a comma separated list of FQDNs. The default value is the Cluster's VIP hostname.
email-address	string	no	An alternative subject name component to be included in the certificate. Format is a standard e-mail address, for example joe@company.com.
host-ips	list of strings	no	An alternative subject name component to be included in the certificate. It is used to identify the ways the Cluster will be accessed. It is given as a comma separated list of IP addresses. The default value is the Cluster's VIP addresses.
key-size	int	no	The size of the keys in bits. The default is 2048 bits for the RSA keys and 256 bits for ECDSA.
key-type	string	no	The algorithm to be used to generate the key pair.RSA is the default value.
organization	string	yes	The organization attribute, which is part of the distinguished name definition. It is used to specify the name of the company.
organization-unit	string	yes	Specifies the organization unit attribute, which is part of the distinguished name definition. It is used to identify the specific department or business unit in the company that is owning the Cluster.
state	string	yes	The state attribute, which is part of the distinguished name definition. It is used to identify the state where the city is located.

clientcsr import

To import a certificate by csr.

Parameter	Type	Required	Description
client-certificate-path	string	yes	Path of client certificate file
server-certificate-path	string	yes	Path of server certificate file

cluster

To run any Cluster related operations.

Operation	Short Name	Description
add-etc-hosts		To add the info of a host to the cluster.
add-nodes	join	To add free nodes to the cluster.
add-ssh-keys		To add public SSH key to authorize access to the cluster.
add-subnet-to-whitelist		To allow a client subnet access to Views.
clear-etc-hosts		To clear the hosts file info stored on the Cluster.
create	add	To create a new cluster. Same number of node IPs, node IPMI IPs, and node IDs are expected. If node IDs are not specified, then node IPs are assumed to be configured on the desired nodes and their IDs will be queried.
delete-etc-hosts		To delete the info for one or more hosts of the cluster.
delete-ssl-certificate	rm-ssl-cert	To delete the SSL certificate from the Cluster.
fru-list	fru-ls	To list the available FRUs in the cluster
fru-locate		To Locate the Fru on Hardware.

Operation	Short Name	Description
fru-locate-list		To list all the Fru's being located currently on the cluster.
generate-source-certificate	gen-src-cert	To create a new certificate for agent authentication.
get-dns-server	ls-dns	To get the IP addresses of the DNS servers for the Cluster.
get-domain-names	ls-domains	To get the domain names of the cluster.
get-etc-hosts		To get the hosts info of the Cluster.
get-info	info	To get information about the cluster.
get-io-pref-tier	ls-io-tier	To get the preferred IO tier of the Cluster.
get-nfs-export-paths	ls-nfs-export-paths	To list NFS export paths accessible in the Cluster.
get-ntp-servers	ls-ntp	To get the list of the NTP servers for the Cluster.
get-proxy-servers	ls-proxies	To get the proxy servers for the Cluster.
get-subnet-whitelist	ls-sn-wl	To list client subnets with permissions to access Views.
get-subnets	ls-subnets	To get services reserved subnets in the cluster.
get-upgrade-status		To get the Cluster software upgrade status.
list-interfaces	ls-ifaces	To list the network interfaces.
list-packages	ls-pkgs	To list the available Cohesity software packages on a Cluster.
list-source-certificates	list-src-cert	To list all global certificates deployed from the cluster.
list-ssl-cert-details	ls-ssl-cert	To get the SSL certificate details of the Cluster.
ls-ssh-keys		List public SSH Keys.
query-replacement-procedure	qrp	To query the fru replacement steps.

Operation	Short Name	Description
register	claim	To register the cluster on helios by token.
remove-proxy-server	rm-proxy	To remove a proxy server from the cluster.
remove-ssh-keys		To remove public SSH key that authorizes access to the cluster.
remove-subnet-from-whitelist		To revoke a client subnet's permissions to access Views.
rotate-ssh-key		Run the above command to rotate cluster ssh keys on demand, this command will only work when there is no active cluster management operation in progress.
route-snapfs-traffic-broker		To route snapfs traffic through broker.
status		To get the status of the Cluster.
sync-ntp		To sync cluster-wide nodes to NTP server.
toggle-protrpc-encryption	tpenc	To either enable or disable protorpc encryption on the cluster.
update	edit	To update editable information of the cluster.
update-active-monitoring	edit-active-mon	To update active monitoring configuration: whether to enable/disable it and what authentication key to use.
update-dns-servers	edit-dns	To update the IP addresses of the DNS Servers for the cluster. Please use empty string to reset DNS Server IPs.
update-domain-names	edit-domains	To update the domain names of the cluster.
update-interfaces	edit-ifaces	To update the interfaces of the cluster.
update-ipmi-credential	edit-ipmi-cred	To update the IPMI username and password.
update-ntp-servers	edit-ntp	To edit the list of NTP servers for the cluster.

Operation	Short Name	Description
update-proxy-server	edit-proxy	To edit a proxy server of the cluster. Proxy servers are used for getting upgrade packages and proactive monitoring.
update-ssl-certificate	edit-ssl-cert	To update the SSL certificate to be used by the cluster. A self-signed SSL certificate is used by default.
update-subnet	edit-subnet	To update services reserved subnets in the cluster.
upgrade		To upgrade the software version on the cluster.
upgrade-check-list-tests		To list available upgrade check tests.
upgrade-check-run-tests		To run upgrade check tests.
upgrade-check-show-result		To show results of upgrade check tests.
upload-etc-hosts		To upload the hosts info to the cluster, the content of the hosts file will be uploaded.
upload-license		To upload valid license string.
virtual-robo-create		To create a new Virtual ROBO cluster.

cluster add-etc-hosts

To add the info of a host to the Cluster.

Parameter	Type	Required	Description
description	string	no	Description for the host.
host-domain-name	list of strings	yes	Comma separated list of host domain names.
host-ip	list of strings	yes	Comma separated list of host IPs.

cluster add-nodes

To add free nodes to the cluster.

Parameter	Type	Required	Description
chassis-serials	list of strings	no	Comma separated list of chassis serials.
dns-groups	list of strings	no	Comma separated list of DNS groups of the new nodes. Each entry should be specified in the following format: <group name1>/<node IP1>:<node IP2>:.../<DNS server IP1>...,<group name2>/<node IP3>/<DNS server IP2>,... For example: dns1/10.1.1.1:10.1.1.2/10.1.0.1:10.1.1.1,dns2/10.2.1.1/10.2.0.1
in-same-rack	bool	no	Specify whether the new nodes should belong to a rack.
node-ids	list of ints	no	Comma separated list of Node IDs.
node-ipmi-ips	list of strings	no	Comma separated list of Node IPMI IP addresses (IPv4 only).
node-ipmi-subnet-gateways	list of strings	no	Comma separated list of IPMI Subnetgateways.
node-ipmi-subnet-masks	list of strings	no	Comma separated list of IPMI Subnet masks. Mask should be a CIDR prefix or a dotted decimal for IPv4.
node-ips	list of strings	yes	Comma separated list of Node IP addresses (IPv4 or IPv6).
rack-ids	list of ints	no	Comma separated list of rack Ids.
subnet-gateways	list of strings	no	Comma separated list of gateways for subnet groups of the new nodes. This is used for multi subnets cluster.
subnet-mask-groups	list of strings	no	Subnet mask groups of the new nodes. Subnet should be either in CIDR prefix or in dotted decimal format. This is used for multiple subnets cluster.Each entry should be specified in the following format: <group name1>/<node IP1>:<node IP2>.../<subnet mask1>,...For example: zone1/10.1.1.1:10.1.1.2/255.255.255.0,zone2/10.2.1.1/20
use-as-compute-nodes	list of bools	no	Comma separated list of true or false values to indicate whether the corresponding node should be used only for compute.

Parameter	Type	Required	Description
vips	list of strings	no	Comma separated list of Virtual IP addresses.

cluster add-ssh-keys

To add public SSH key to authorize access to the cluster.

Parameter	Type	Required	Description
ssh-key-file-path	absolute file path	yes	Absolute file path for the public SSH keys that allows access to the Cluster.

cluster add-subnet-to-whitelist

To allow a client subnet access to Views.

Parameter	Type	Required	Description
subnet-ip	string	yes	IP of the client subnet.
subnet-mask-bits	int	no	Mask bits for the client subnet.
subnet-mask-ip4	string	no	Netmask for the IPv4 client subnet.
subnet-nfs-permission	string	no	Clients from this subnet can mount using NFS protocol. This can be kDisabled, kReadOnly, kReadWrite.
subnet-nfs-root-squash	bool	no	Clients from this subnet can mount as root on NFS. This can be Disabled, ReadOnly, ReadWrite.
subnet-s3-permission	string	no	Clients from this subnet can access using S3 protocol. This can be kDisabled, kReadOnly, kReadWrite.
subnet-smb-permission	string	no	Clients from this subnet can mount using SMB protocol. This can be kDisabled, kReadOnly, kReadWrite.

cluster create

To create a new Cluster. Same number of node IPs, node IPMI IPs, and node IDs are expected. If node IPs are already assigned, don't specify the node ID. If node IDs are not specified, then node IPs are assumed to be configured on the desired nodes and their IDs will be queried.

Parameter	Type	Required	Description
apps-subnet	string	yes	Configure the subnet IP for apps configuration
apps-subnet-mask	string	yes	Configure the subnet Mask for apps configuration
dns-groups	list of strings	no	Comma separated list of DNS groups of the cluster. Each entry should be specified in the following format: <group name1>/<node IP1>:<node IP2>:.../<DNS server IP1>:...,<group name2>/<node IP3>/<DNS server IP2>,... For example: dns1/10.1.1.1:10.1.1.2/10.1.0.1|10.1.1.1,dns2/10.2.1.1/10.2.0.1
dns-server-ips	list of strings	no	Comma separated list of DNS server IP addresses.
domain-names	list of strings	no	Comma separated list of search domain names. First name is also the Cluster domain name which is mandatory.
enable-software-encryption	bool	no	The cluster encryption is ON by default during the initial cluster setup and cannot be disabled later. Encryption will be OFF only if enable-software-encryption is explicitly set to false during cluster creation.
hostname	string	yes	Hostname for the Virtual IPs of the cluster partition.
ip-preference	int	no	Param to specify IP address preference, 1=IpV4(default) or 2=IpV6.
ipmi-password	string	no	IPMI password for the Cluster.
ipmi-subnet-gateway	string	no	IPMI Subnet gateway.
ipmi-subnet-mask	string	no	IPMI Subnet mask should be either in CIDR prefix or in dotted decimal for IPv4.
ipmi-username	string	no	IPMI username for the Cluster.
name	string	yes	Name of the Cluster.
node-ids	list of ints	no	Comma separated list of Node IDs.

Parameter	Type	Required	Description
node-ipmi-ips	list of strings	no	Comma separated list of Node IPMI IP addresses (IPv4 only).
node-ips	list of strings	yes	Comma separated list of Node IP addresses (IPv4 or IPv6).
ntp-auth-enable	bool	no	Flag to enable NTP with authentication.
ntp-auth-key-ids	list of ints	no	Comma separated list of NTP server auth key IDs corresponding to each NTP server auth key.
ntp-auth-keys	list of strings	no	Comma separated list of NTP server auth keys corresponding to each NTP server.
ntp-servers	list of strings	yes	Comma separated list of NTP servers IPs/hostnames.
rotation-period-days	int	no	Period in days when the data encryption key (DEK) is rotated. Default rotation period is 90 days.
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-gateways	list of strings	no	Comma separated list of gateways for subnet groups of the cluster. This is used for multi subnets cluster.
subnet-mask	string	no	Subnet mask for the Cluster. Should be either in CIDR prefix or in dotted decimal.
subnet-mask-groups	list of strings	no	Subnet mask groups of the cluster. Subnet should be either in CIDR prefix or in dotted decimal format. This is used for multiple subnets cluster.Each entry should be specified in the following format: <group name1>/<node IP1>:<node IP2>:.../<subnet mask1>,...For example: zone1/10.1.1.1:10.1.1.2/255.255.255.0,zone2/10.2.1.1/20
trust-domain	string	no	Trust domain name of FQDN format.
use-as-compute-nodes	list of bools	no	Comma separated list of true or false values to indicate whether the corresponding node should be used only as compute node.
vips	list of strings	no	Comma separated list of Virtual IP addresses.

cluster delete-etc-hosts

To delete the info for one or more hosts of the Cluster.

Parameter	Type	Required	Description
host-domain-name	list of strings	no	Comma separated list of host domain names.
host-ip	list of strings	no	Comma separated list of host IPs.

cluster fru-list

To list the available FRUs in the cluster

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.

cluster fru-locate

To Locate the Fru on Hardware.

Parameter	Type	Required	Description
fru-id	string	yes	Fru Id of Fru to be located/replaced
fru-type	string	yes	Type (disk,fan) of Fru to be located/replaced
node-id	int	yes	ID of a Node.
status-on	bool	yes	Fru locate status (on/off)

cluster generate-source-certificate

To create a new certificate for agent authentication.

Parameter Name	Description	Parameter Name	Description
cert-name	Name of the certificate file. Its naming format is ${cert-name}_<random_number>.cfg.	cert-name	Name of the certificate file. Its naming format is ${cert-name}_<random_number>.cfg.
cert-validity	The validity of the source certificate generated. The time is displayed in days.	cert-validity	The validity of the source certificate generated. The time is displayed in days.
host-name	The IP address/host name of the target system where you want to move the certificate. This is a mandatory parameter because the certificate generation process needs details of the host machine to generate public and private keys.	host-name	The IP address/host name of the target system where you want to move the certificate. This is a mandatory parameter because the certificate generation process needs details of the host machine to generate public and private keys.
host-target	The file directory path on the host machine where the certificate will be stored.	host-target	The file directory path on the host machine where the certificate will be stored.

Parameter Name	Description	Parameter Name	Description
host-uname and host-password	The username and password of any user on the host machine that can facilitate the transfer of certificate over an SSH tunnel.	host-uname and host-password	The username and password of any user on the host machine that can facilitate the transfer of certificate over an SSH tunnel.
Parameter Name	Description	Parameter Name	Description

cluster get-info

To get information about the Cluster.

Parameter	Type	Required	Description
show-stats	bool	no	To display the computed statistics.

cluster list-interfaces

To list the network interfaces.

Parameter	Type	Required	Description
cache	bool	no	To show interface cached information or not. Default is to read interface information from cache.

cluster query-replacement-procedure

To query the fru replacement steps

Parameter	Type	Required	Description
fru-id	string	no	Fru Id of Fru to be located/replaced
fru-type	string	no	Type (disk,fan) of Fru to be located/replaced

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
task-id	string	no	Task Id of Any Breakfix Task running in cluster

cluster register

To register the cluster on helios by token.

Parameter	Type	Required	Description
rigel-guid	int	no	To specify the guid to register the rigel with as cluster on helios. If not specified, node id is used as rigel guid.
token	string	yes	To specify the token to register the cluster on helios.

cluster remove-proxy-server

To remove a proxy server from the Cluster.

Parameter	Type	Required	Description
name	string	yes	Name of the proxy server.

cluster remove-ssh-keys

To remove public SSH key that authorizes access to the cluster.

Parameter	Type	Required	Description
ssh-key-file-path	absolute file path	yes	Absolute file path for the public SSH keys that allows access to the Cluster.

cluster remove-subnet-from-whitelist

To revoke a client subnet's permissions to access Views.

Parameter	Type	Required	Description
subnet-ip	string	yes	IP of the client subnet.

cluster route-snapfs-traffic-broker

To route snapfs traffic through broker.

Parameter	Type	Required	Description
enable	bool	yes	To enable or disable snapfs traffic through broker. If set to true, enables the snapfs traffic through broker in rigel. If set to false, disables it.

cluster toggle-protrpc-encryption

To either enable or disable protorpc encryption on the cluster.

Parameter	Type	Required	Description
enable-encryption	bool	yes	To enable or disable internode protorpc encryption, if the set to true the commannd shall enable protorpc encryption if set to false it shall disable.

cluster update

To update editable information of the Cluster.

Parameter	Type	Required	Description
authentication-type	string	no	The authentication type for the cluster. This can be password-only, certificate-only or password-and-certificate
cluster-audit-enabled	bool	no	Whether to enable Cluster Audit Logs.
cluster-audit-retention-days	int	no	Number of days to retain the Cluster Audit Logs.
cluster-audit-verbosity	bool	no	Whether to enable audit verbosity(audit log prev value and new value).
enable-local-groups	bool	no	Enable/Disable creating local groups.

Parameter	Type	Required	Description
enc-key-rotation-period-secs	int	no	Encryption key rotation period in seconds. It should be greater than or equal to 86400 seconds (a day).
filer-audit-enabled	bool	no	Whether to enable Filer Audit Logs.
filer-audit-retention-days	int	no	Number of days to retain the Filer Audit Logs.
local-auth-domain-name	string	no	Domain name for SMB local authentication.
multi-tenancy	bool	no	To enable or disable Multi-Tenancy.
new-name	string	no	New name of the Cluster for update.
ntp-servers-internal	bool	no	Flag to specify if the NTP servers are on internal network.
pcie-ssd-rebalance-delay-secs	int	no	Delay time in seconds for cluster PcieSSD storage tier.
sata-hdd-admission-control	int	no	Admission control for cluster SATA HDD storage tier.
smb-ad-disabled	bool	no	Disable SMB access for Active Directory Users.
tiering-audit-enabled	bool	no	Whether to enable Tiering Audit Logs.
tiering-audit-retention-days	int	no	Number of days to retain the Tiering Audit Logs.
time-zone	string	no	Timezone for displaying Time in Reports download.
turbo-mode	bool	no	To set or clear Turbo mode.
use-heimdall	bool	no	To enable or disable Heimdall.
viewbox-sharing	bool	no	To enable ViewBox Sharing.

cluster update-active-monitoring

To update active monitoring configuration: whether to enable/disable it and what authentication key to use.

Parameter	Type	Required	Description
enable	bool	yes	To enable/disable active monitoring.

cluster update-dns-servers

To update the IP addresses of the DNS Servers for the Cluster. Please use empty string to reset DNS Server IPs.

Parameter	Type	Required	Description
dns-server-ips	string	yes	Comma separated list of DNS server IP addresses.

cluster update-domain-names

To update the domain names of the cluster.

Parameter	Type	Required	Description
domain-names	list of strings	no	Comma separated list of search domain names. First name is also the Cluster domain name which is mandatory.

cluster update-interfaces

To update the interfaces of the cluster.

Parameter	Type	Required	Description
family	int	no	Param to specify IP address family, 1=IpV4(default) or 2=IpV6.
ids	list of ints	yes	Comma separated list of IDs.
interface-ips	string	no	Comma separated list of interface IP addresses (IPv4 only).
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.

Parameter	Type	Required	Description
unconfigure	bool	no	Whether to unconfigure the interfaces.

cluster update-ipmi-credential

To update the IPMI username and password.

Parameter	Type	Required	Description
ipmi-password	string	yes	IPMI password for the Cluster.
ipmi-username	string	yes	IPMI username for the Cluster.

cluster update-ntp-servers

To edit the list of NTP servers for the Cluster.

Parameter	Type	Required	Description
ntp-auth-enable	bool	no	Flag to enable NTP with authentication.
ntp-auth-key-ids	list of ints	no	Comma separated list of NTP server auth key IDs corresponding to each NTP server auth key.
ntp-auth-keys	list of strings	no	Comma separated list of NTP server auth keys corresponding to each NTP server.
ntp-servers	list of strings	yes	Comma separated list of NTP severs IPs/hostnames. cluster edit-ntp ntp-servers="xx.x1, xx.x2, xx.x3" Example: cluster edit-ntp ntp-servers="time.google.com, time2.google.com, time3.google.com"

cluster update-proxy-server

To edit a proxy server of the Cluster. Proxy servers are used for getting upgrade packages and proactive monitoring.

Parameter	Type	Required	Description
address	string	yes	FQDN or IP address of the proxy server.
name	string	yes	Name of the proxy server.
password	string	no	Password to use for the proxy server.
port	int	yes	Port to use for the proxy server.
schemes	list of strings	no	Comma separated list of schemes supported by the proxy server. Choose from http, or https.
username	string	no	Username to use for the proxy server.

cluster update-ssl-certificate

To update the SSL certificate to be used by the Cluster. A self-signed SSL certificate is used by default.

Parameter	Type	Required	Description
ssl-cert-private-key	absolute file path	yes	Absolute file path for the private key to be used in conjunction with the SSL certificate for the Cluster.
ssl-certificate	absolute file path	yes	Absolute file path for the SSL certificate to update the Cluster.

cluster update-subnet

To update services reserved subnets in the cluster.

Parameter	Type	Required	Description
desc	string	no	Description.
service-name	string	no	Name of the service. You can get the list of services by running "cluster status" command.
subnet-id	int	no	Subnet ID.
subnet-ip	string	yes	Subnet IP for the Cluster.

Parameter	Type	Required	Description
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.
subnet-mask-ip4	string	no	Subnet mask for IPv4 type subnet of the Cluster.

cluster upgrade

To upgrade the software version on the Cluster.

Parameter	Type	Required	Description
force-validate-health-checks	bool	no	Force running health check tests as part of upgrade
id	int	no	ID of the Cluster.
skip-nodes	string	no	Comma separated list of node IDs on which upgrade checks will not run
skip-tests	string	no	Comma separated list of test case ids which will not run on all cluster nodes. Use command 'cluster upgrade-list-tests' to get the list of tests.
software-pkg-md5	string	no	MD5 Hash of the software package to be uploaded to the Node.
software-pkg-url	string	no	URL where Cohesity Software package is available.
target-version	string	no	Target software version to be used for upgrade.

cluster upgrade-check-list-tests

To list available upgrade check tests.

Parameter	Type	Required	Description
type	string	yes	Type of upgrade check tests, possible inputs are preupgrade or postupgrade

cluster upgrade-check-run-tests

To run upgrade check tests.

Parameter	Type	Required	Description
skip-nodes	list of strings	no	Comma separated list of node IDs on which upgrade checks will not run
skip-tests	list of strings	no	Comma separated list of test case ids which will not run on all cluster nodes. Use command 'cluster upgrade-list-tests' to get the list of tests.
type	string	yes	Type of upgrade tests, possible inputs are preupgrade or postupgrade

cluster upload-etc-hosts

To upload the hosts info to the cluster, the content of the hosts file will be uploaded.

Parameter	Type	Required	Description
file-path	string	no	Path of the hosts file to upload for cluster level /etc/hosts settings.

cluster upload-license

To upload valid license string.

Parameter	Type	Required	Description
license	string	yes	To upload valid license string.

cluster virtual-robo-create

To create a new Virtual ROBO cluster.

Parameter	Type	Required	Description
apps-subnet	string	no	Configure the subnet IP for apps configuration
apps-subnet-mask	string	no	Configure the subnet Mask for apps configuration
dns-server-ips	list of strings	yes	Comma separated list of DNS server IP addresses.

Parameter	Type	Required	Description
domain-names	list of strings	yes	Comma separated list of search domain names. First name is also the Cluster domain name which is mandatory.
enable-software-encryption	bool	no	Enable software encryption at Cluster level. WARNING: This option cannot be disabled later.
hostname	string	yes	Hostname for the Virtual IPs of the Cluster Partition.
ip-preference	int	no	Param to specify IP address preference, 1=IpV4(default) or 2=IpV6.
name	string	yes	Name of the Cluster.
node-ids	list of ints	no	Comma separated list of Node IDs.
node-ips	string	yes	Comma separated list of Node IP addresses (IPv4 or IPv6).
ntp-auth-enable	bool	no	Flag to enable NTP with authentication.
ntp-auth-key-ids	list of ints	no	Comma separated list of NTP server auth key IDs corresponding to each NTP server auth key.
ntp-auth-keys	list of strings	no	Comma separated list of NTP server auth keys corresponding to each NTP server.
ntp-servers	list of strings	yes	Comma separated list of NTP severs IPs/hostnames.
rotation-period-days	int	no	Period in days when data encryption key (DEK) is rotated. Default rotation period is 90 days.
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-mask	string	yes	Subnet mask for the Cluster. Should be either in CIDR prefix or in dotted decimal.
trust-domain	string	no	Trust domain name of FQDN format.
vips	list of strings	no	Comma separated list of Virtual IP addresses. This is applicable only for Virtual Edition cluster create.

csr

To run any csr related operations.

Operation	Short Name	Description
create	add	To create a csr.
delete	rm	To delete a csr.
import		To import a certificate by csr.
list	ls	To list the csrs.

csr create

To create a csr.

Parameter	Type	Required	Description
city	string	yes	The locality attribute, which is part of the distinguished name definition. It is used to identify the city where the company is located or the Cluster is installed.
common-name	string	no	The common name attribute, which is part of the distinguished name definition. Common name is used to specify a context for the certificate, for example, the name of the Cluster to which the certificate is to be assigned. Default value is the name of the Cluster.
country	string	yes	The country attribute, which is part of the distinguished name definition. It is used to identify the country where the state is located. It is specified as two letter code defined by the ISO standard.
dns-names	list of strings	no	An alternative subject name component to be included in the certificate. It is used to identify the ways the Cluster will be accessed. It is given as a comma separated list of FQDNs. The default value is the Cluster's VIP hostname.
email-address	string	no	An alternative subject name component to be included in the certificate. Format is a standard e-mail address, for example joe@company.com.

Parameter	Type	Required	Description
host-ips	list of strings	no	An alternative subject name component to be included in the certificate. It is used to identify the ways the Cluster will be accessed. It is given as a comma separated list of IP addresses. The default value is the Cluster's VIP addresses.
key-size	int	no	The size of the keys in bits. The default is 2048 bits for the RSA keys and 256 bits for ECDSA.
key-type	string	no	The algorithm to be used to generate the key pair.RSA is the default value.
organization	string	yes	The organization attribute, which is part of the distinguished name definition. It is used to specify the name of the company.
organization-unit	string	yes	Specifies the organization unit attribute, which is part of the distinguished name definition. It is used to identify the specific department or business unit in the company that is owning the Cluster.
service-name	string	no	Service for which the csr is created.
state	string	yes	The state attribute, which is part of the distinguished name definition. It is used to identify the state where the city is located.

csr delete

To delete a csr.

Parameter	Type	Required	Description
id	string	yes	Csr ID.

csr import

To import a certificate by csr.

Parameter	Type	Required	Description
certificate-path	string	yes	The path of the file of the certificate to be imported.
id	string	yes	Csr ID.

csr list

To list the csrs.

Parameter	Type	Required	Description
ids	list of strings	no	Comma separated list of csr IDs.
service-name	string	no	Service for which the csr is created.

custom-reporting

To run any custom reporting related operations.

Operation	Short Name	Description
database	db	To get details for running database information.

disk

To run any Disk related operations.

Operation	Short Name	Description
identify	ident	To identify a disk on a node. Either a disk ID in the cluster, or a disk serial and the ID of the node that the disk is attached to need to be specified
list	ls	To list the disk details.

disk identify

To identify a disk on a node. Either a disk ID in the cluster, or a disk serial and the ID of the node that the disk is attached to need to be specified

Parameter	Type	Required	Description
cluster-id	int	no	ID of the Cluster.

Parameter	Type	Required	Description
disk-id	int	no	ID of the Disk.
identify	bool	yes	Flag to toggle identification LED status.

disk list

To list the disk details.

Parameter	Type	Required	Description
fetch-disk-usage-details	bool	no	Enable/Disable display of disk usage details.
ids	list of ints	no	Comma separated list of IDs.
include-marked-for-removal	bool	no	Include the nodes which have been marked for removal.
node-ids	list of ints	no	Comma separated list of Node IDs.
only-ready-for-removal	bool	no	Display disks which are ready for removal.
show-stats	bool	no	To display the computed statistics.
sum-disk-reservations	bool	no	To sum the disk reservations by different components across all the disks.

firewall-ipset

To run any Firewall ipset related operations.

Operation	Short Name	Description
list	`ls`	To list the IP sets for the firewall.

firewall-profile

To run any Firewall profile related operations.

Operation	Short Name	Description
activate		To activates a profile on a list of interfaces with a list of ipsets.
create	add	To create a collection of service ports with the protocol of each port for firewall.
deactivate		To deactivate a list of active firewall profiles, Each one will stop a profile on a list of interfaces with a list of ipsets.
delete	rm	To remove a list of profiles for firewall.

firewall-profile activate

To activates a profile on a list of interfaces with a list of ipsets.

Parameter	Type	Required	Description
action	string	no	To define the action of the rule set. It can be either "allow" or "deny".
desc	string	no	Description of the firewall rule sets.
interface-groups	list of strings	no	Comma separated list of network interface groups.
ipset-entries	list of strings	no	Comma separated list of subnets/ips for the firewall Ipsets.
profile-name	string	no	Name of the firewall profile.
revert-to-default	bool	no	To revert all firewall rules to default settings.

firewall-profile create

To create a collection of service ports with the protocol of each port for firewall.

Parameter	Type	Required	Description
directions	list of strings	no	Comma separated list of traffic directions. Direction can be "INPUT", "OUTPUT", or "FORWARD". "INPUT" is the default direction.

Parameter	Type	Required	Description
force	bool	no	Set to true to unlock the ports restriction.
name	string	yes	Name of the firewall profile.
ports	list of strings	yes	Comma separated list of service ports.

firewall-profile deactivate

To deactivate a list of active firewall profiles, Each one will stop a profile on a list of interfaces with a list of ipsets.

Parameter	Type	Required	Description
profile-names	list of strings	yes	Comma separated list of names for the firewall profile.

firewall-profile delete

To remove a list of profiles for firewall.

Parameter	Type	Required	Description
names	list of strings	yes	Comma separated list of names for the firewall profile.

group

To run any Group related operations.

Operation	Short Name	Description
add-smb-principal		To add SMB principals as member to the group. SMB principals can be added only if "smb-only" set to true for a group.
create	add	To create a new Group.
delete	rm	To delete a Group.
list	ls	To list Groups.

Operation	Short Name	Description
remove-smb-principal		To delete SMB principal from a group.
update	edit	To update a Group.

group add-smb-principal

To add SMB principals as member to the group. SMB principals can be added only if "smb-only" set to true for a group.

Parameter	Type	Required	Description
domain	string	no	Domain name of the principal.
name	string	yes	Name of the Group.
principal-name	string	no	Name of the SMB principal which may be a group or user.
sid	string	no	unique Security ID (SID) of the principal that look similar to windows domain SID.
type	string	no	Type can be a user or a group.

group create

To create a new Group.

Parameter	Type	Required	Description
description	string	no	Description of the Group.
domain	string	no	Domain identifies User's logging in domain.
name	string	yes	Name of the Group.
roles	list of strings	no	Comma separated list of Roles for the Group.
users	list of strings	no	Comma separated list of User SIDs in the Group.

group delete

To delete a Group.

Parameter	Type	Required	Description
domain	string	no	Domain identifies User's logging in domain.
names	list of strings	yes	Comma separated list of Group names.

group list

To list Groups.

Parameter	Type	Required	Description
domain	string	no	Domain identifies User's logging in domain.
name	string	no	Name of the Group.

group remove-smb-principal

To delete SMB principal from a group.

Parameter	Type	Required	Description
clear-smb-principals	bool	no	Removes all the SMB principals from the SMB group.
name	string	yes	Name of the Group.
principal-name	string	no	Name of the SMB principal which may be a group or user.

group update

To update a Group.

Parameter	Type	Required	Description
description	string	no	Description of the Group.
domain	string	no	Domain identifies User's logging in domain.
name	string	yes	Name of the Group.
roles	list of strings	no	Comma separated list of Roles for the Group.

Parameter	Type	Required	Description
users	list of strings	no	Comma separated list of User SIDs in the Group.

hyx-vlan

To run any Bifrost Vlan related operations.

Operation	Short Name	Description
create	add	To add a Bifrost VLAN for the cluster
delete	rm	To remove a Bifrost VLAN from the cluster
get	get	To get Bifrost VLAN for given IfaceGroupName for the cluster
tenant-disable		This disables tenant-facing components of Internal Bifrost for the provided tenant ids
tenant-enable		This enables tenant-facing components of Internal Bifrost for the provided tenant ids
update	edit	To update a Bifrost VLAN for the cluster

hyx-vlan create

To add a Bifrost VLAN for the cluster

Parameter	Type	Required	Description
desc	string	no	Description.
id	int	no	ID of the VLAN.
mtu	int	no	MTU for the interfaces.
state	string	no	This specifies the state of Bifrost VLAN. State can be UNKNOWN/ACTIVE/DISABLED/DELETING.
subnet-gateway	string	no	Subnet gateway for the Cluster.

Parameter	Type	Required	Description
subnet-ip-cidr	string	yes	Subnet IP CIDR
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.
subnet-mask-ip4	string	no	Subnet mask for IPv4 type subnet of the Cluster.
tenant-id	string	no	To specify id of the tenant owning a resource.
type	string	yes	This specifies the type of Bifrost Vlan. Type can be EXTERNAL/INTERNAL.
vlan-interface-group	string	yes	Name of a network interface group.
vlan-ips	list of strings	no	Comma separated list of VLAN IP addresses.
vlan-name	string	no	VLAN Name for vlan-id created

hyx-vlan delete

To remove a Bifrost VLAN from the cluster

Parameter	Type	Required	Description
tenant-id	string	no	To specify id of the tenant owning a resource.
vlan-interface-group	string	yes	Name of a network interface group.

hyx-vlan get

To get Bifrost VLAN for given IfaceGroupName for the cluster

Parameter	Type	Required	Description
vlan-interface-group	string	yes	Name of a network interface group.

hyx-vlan tenant-disable

This disables tenant-facing components of Internal Bifrost for the provided tenant ids

Parameter	Type	Required	Description
tenant-ids	list of strings	yes	To specify multiple tenant ids.

hyx-vlan tenant-enable

This enables tenant-facing components of Internal Bifrost for the provided tenant ids

Parameter	Type	Required	Description
tenant-ids	list of strings	yes	To specify multiple tenant ids.

hyx-vlan update

To update a Bifrost VLAN for the cluster

Parameter	Type	Required	Description
desc	string	no	Description.
id	int	no	ID of the VLAN.
mtu	int	no	MTU for the interfaces.
state	string	no	This specifies the state of Bifrost VLAN. State can be UNKNOWN/ACTIVE/DISABLED/DELETING.
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-ip-cidr	string	no	Subnet IP CIDR
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.
subnet-mask-ip4	string	no	Subnet mask for IPv4 type subnet of the Cluster.
tenant-id	string	no	To specify id of the tenant owning a resource.
type	string	no	This specifies the type of Bifrost Vlan. Type can be EXTERNAL/INTERNAL.

Parameter	Type	Required	Description
vlan-interface-group	string	yes	Name of a network interface group.
vlan-ips	list of strings	no	Comma separated list of VLAN IP addresses.
vlan-name	string	no	VLAN Name for vlan-id created

interface-group

To run network interface group related operations.

Operation	Short Name	Description
create	add	To create the interface group in the cluster.
delete	rm	To delete the interface group in the cluster.
update	edit	To update the interface group in the cluster.

interface-group create

To create the interface group in the cluster.

Parameter	Type	Required	Description
group-name	string	yes	Name of a network interface group.
group-type	int	no	Type of a network interface group. 1(Bond Interfaces), 2 (Loopback Interfaces). Default is 1.
interface-name	string	no	Name of the bond interface.
node-ids	list of ints	yes	Comma separated list of Node IDs.

interface-group delete

To delete the interface group in the cluster.

Parameter	Type	Required	Description
group-name	string	yes	Name of a network interface group.

interface-group update

To update the interface group in the cluster.

Parameter	Type	Required	Description
add-node-ids	list of ints	no	Comma separated list of Node IDs to add.
group-name	string	yes	Name of a network interface group.
interface-name	string	no	Name of the bond interface.
remove-node-ids	list of ints	no	Comma separated list of Node IDs to remove.

interface

To run interface related operations.

Operation	Short Name	Description
list	ls	To list the network interfaces.
update	edit	To update interface params.

interface list

To list the network interfaces.

Parameter	Type	Required	Description
bondInterfaceOnly	bool	no	To show only bond interface info.
cache	bool	no	To show interface cached information or not. Default is to read interface information from cache.
include-bond-slave-details	bool	no	To include bond slave details.

Parameter	Type	Required	Description
include-stats	bool	no	To include interface statistics.
node-id	int	no	ID of a Node.

interface update

To update interface params.

Parameter	Type	Required	Description
bonding-mode	int	no	Bonding mode to use when bonding the NICs. Modes 1 (active-backup) and 4 (802.3ad) are supported. Mode 4 may require switch configuration. In cluster, the services need to be stopped before changing on primary interface group.
interface-name	string	yes	Name of the network interface. It can be bond interface, vlan interface or interface group.
lacp-rate	string	no	Rate option to use for link partner to transmit LACPDU packets in 802.3ad mode. Possible values are 0 (slow) and 1 (fast). Default setting is 0 (slow).
mtu	int	no	MTU for the interfaces.
xmit-hash-policy	string	no	Transmit hash policy to use for selection in 802.3ad mode. Possible values are 0 (layer2), 1 (layer3+4), and 2 (layer2+3). Default system setting is 1 (layer3+4).

ip

To run ip related operations.

Operation	Short Name	Description
config		To configure IP parameters on the interface.
delete	rm	To delete IP parameters on the interface.

ip config

To configure IP parameters on the interface.

Parameter	Type	Required	Description
family	int	no	Param to specify IP address family, 1=IpV4(default) or 2=IpV6.
interface-ips	string	no	Comma separated list of interface IP addresses (IPv4 only).
interface-name	string	yes	Name of the network interface. It can be bond interface, vlan interface or interface group.
interface-role	string	no	Role of the network interface. It can be primary, secondary.
node-ids	list of ints	no	Comma separated list of Node IDs.
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.

ip delete

To delete IP parameters on the interface.

Parameter	Type	Required	Description
family	int	no	Param to specify IP address family, 1=IpV4(default) or 2=IpV6.
interface-name	string	yes	Name of the network interface. It can be bond interface, vlan interface or interface group.
node-ids	list of ints	no	Comma separated list of Node IDs.

ipmi

To run IPMI related operations on local and other nodes of the cluster.

Operation	Short Name	Description
cluster-update-lan-info		Update IPMI LAN info in the cluster.
cluster-update-users		Update IPMI users in the cluster.
get-fru-info		To get IPMI FRU info.

Operation	Short Name	Description
get-lan-info		To get IPMI LAN info.
get-sdr-info		To get IPMI SDR info.
get-sel		To get the IPMI SEL list.
get-sel-info		To get IPMI SEL info.
list-users		To list IPMI users.
reset-bmc		To reset the BMC.
show-system-led-info		Show System LED information.

ipmi cluster-update-lan-info

Update IPMI LAN info in the cluster.

Parameter	Type	Required	Description
ipmi-subnet-gateway	string	no	IPMI Subnet gateway.
ipmi-subnet-mask	string	no	IPMI Subnet mask should be either in CIDR prefix or in dotted decimal for IPv4.
node-ipmi-ips	list of strings	no	Comma separated list of Node IPMI IP addresses (IPv4 only).
node-ipmi-subnet-gateways	list of strings	no	Comma separated list of IPMI Subnetgateways.
node-ipmi-subnet-masks	list of strings	no	Comma separated list of IPMI Subnet masks. Mask should be a CIDR prefix or a dotted decimal for IPv4.
node-ips	list of strings	no	Comma separated list of Node IP addresses (IPv4 or IPv6).

ipmi cluster-update-users

Update IPMI users in the cluster. For more information, see Manage IPMI Configuration Using the Cohesity CLI.

Parameter	Type	Required	Description
ipmi-password	string	no	IPMI password for the Cluster.
ipmi-username	string	no	IPMI username for the Cluster.
node-ipmi-passwords	list of strings	no	IPMI passwords for the given nodes.
node-ipmi-usernames	list of strings	no	IPMI usernames for the given nodes.
node-ips	list of strings	no	Comma separated list of Node IP addresses (IPv4 or IPv6).

ipmi get-fru-info

To get IPMI FRU info.

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
node-ip	string	no	Node IP addresses (IPv4 or IPv6).

ipmi get-lan-info

To get IPMI LAN info.

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
node-ip	string	no	Node IP addresses (IPv4 or IPv6).

ipmi get-sdr-info

To get IPMI SDR info.

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
node-ip	string	no	Node IP addresses (IPv4 or IPv6).

ipmi get-sel

To get the IPMI SEL list.

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
node-ip	string	no	Node IP addresses (IPv4 or IPv6).
verbose	bool	no	If set to true, more information is displayed.

ipmi get-sel-info

To get IPMI SEL info.

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
node-ip	string	no	Node IP addresses (IPv4 or IPv6).

ipmi list-users

To list IPMI users.

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
node-ip	string	no	Node IP addresses (IPv4 or IPv6).

ipmi reset-bmc

To reset the BMC.

Parameter	Type	Required	Description
node-id	int	no	ID of a Node.
node-ip	string	no	Node IP addresses (IPv4 or IPv6).

ipmi show-system-led-info

Show System LED information.

Parameter	Type	Required	Description
node-ip	string	no	Node IP address for which system LED information is displayed.

kms

To run any KMS related operations.

Operation	Short Name	Description
create-aws		Create a new aws KMS.
create-azure		Create a new Azure KMS.
create-kmip		Create a new kmip KMS.
delete	rm	To delete a KMS by Id.
list		To list KMS.
update-aws		Update an existing aws KMS.
update-kmip		Update an existing kmip KMS.

kms create-aws

Create a new aws KMS.

Parameter	Type	Required	Description
access-key-id	string	no	The AWS acount access key Id.
arn	string	no	The amazon resource number of the CMK.
auth-method	string	yes	The AWS Auth method.
ca-certificate-path	string	no	File path to ca-certificate.

Parameter	Type	Required	Description
iam-role-arn	string	no	Specify the IamRole ARN for CE. It can either kUseIAMUser or kUseIAMRole.
key-alias	string	no	Key alias of the CMK.
key-id	string	no	The amazon CMK key ID.
key-name	string	yes	Name of the KMS Key.
region	string	no	region
secret-access-key	string	no	The AWS account secret Access key.
vault-ids	list of ints	no	Comma separated list of vault ids.
verify-ssl	bool	no	Specify whether to verify SSL certificate when try to access aws kms server.
view-box-ids	list of ints	no	Comma separated list of Storage Domain IDs.

kms create-azure

Create a new Azure KMS.

Parameter	Type	Required	Description
azure-client-id	string	no	Azure Client Id.
azure-client-secret-key	string	no	Specify Azure client secret key.
azure-kms-auth-method	string	yes	The AWS Auth method.
azure-tenant-id	string	no	Azure Tenant Id.
customer-secret-name	string	no	Specifies the name of client secret stored in cohesity keyvault.
key-name	string	yes	Name of the KMS Key.
keyvault-key-name	string	no	Specify the Key name used to encrypt decrypt DEK.
keyvault-url	string	no	Specify the URL of the Azure key vault.

Parameter	Type	Required	Description
vault-ids	list of ints	no	Comma separated list of vault ids.
vault-owner	string	yes	Specifies if its a cohesity managed or customer managed key vault.
keyvault-key-name	string	no	Specify the Key name used to encrypt decrypt DEK.

kms create-kmip

Create a new kmip KMS.

Parameter	Type	Required	Description
additional-address	list of strings	no	Additional IP addresses or FQDNs of the KMS server.
ca-certificate-path	string	yes	File path to ca-certificate.
client-certificate	string	yes	File path to client-certificate.
client-key	string	yes	File path to client-key.
ip	string	yes	IP address or FQDN of the KMS server.
key-name	string	yes	Name of the KMS Key.
kmip-protocol-version	string	yes	kmip-protocol-version
port	int	yes	KMS Port. Default KMIP port is 5696.
vault-ids	list of ints	no	Comma separated list of vault ids.
view-box-ids	list of ints	no	Comma separated list of Storage Domain IDs.

kms delete

To delete a KMS by Id.

Parameter	Type	Required	Description
id	int	yes	The Id of a KMS server.

kms list

To list KMS.

Parameter	Type	Required	Description
id	int	no	The Id of a KMS server.

kms update-aws

Update an existing aws KMS.

Parameter	Type	Required	Description
access-key-id	string	no	The AWS acount access key Id.
ca-certificate-path	string	no	File path to ca-certificate.
id	int	yes	The Id of a KMS server.
key-name	string	yes	Name of the KMS Key.
secret-access-key	string	no	The AWS account secret Access key.
vault-ids	list of ints	no	Comma separated list of vault ids.
verify-ssl	bool	no	Specify whether to verify SSL certificate when try to access aws kms server.
view-box-ids	list of ints	no	Comma separated list of Storage Domain IDs.

kms update-kmip

Update an existing kmip KMS.

Parameter	Type	Required	Description
additional-address	list of strings	no	Additional IP addresses or FQDNs of the KMS server.
ca-certificate-path	string	no	File path to ca-certificate.
client-certificate	string	no	File path to client-certificate.

Parameter	Type	Required	Description
client-key	string	no	File path to client-key.
confirm	bool	yes	Confirming to update KMS server. Please make sure the new updated information is for thte same KMS server configured before.
id	int	yes	The Id of a KMS server.
ip	string	no	IP address or FQDN of the KMS server.
key-name	string	no	Name of the KMS Key.
kmip-protocol-version	string	no	kmip-protocol-version
port	int	no	KMS Port. Default KMIP port is 5696.
vault-ids	list of ints	no	Comma separated list of vault ids.
view-box-ids	list of ints	no	Comma separated list of Storage Domain IDs.

ldap

To run any LDAP related operations.

Operation	Short Name	Description
create	add	To add an LDAP provider to the cluster.
delete	rm	To delete an LDAP provider from the cluster.
update	edit	To update an LDAP provider to the cluster.

ldap create

To add an LDAP provider to the cluster.

Parameter	Type	Required	Description
active-directory-domain-name	string	no	The active directory domain that should be mapped to this LDAP.
attribute-common-name	string	no	Name of the LDAP attribute used for common name of an object.
attribute-gid	string	no	Name of the attribute used to lookup unix GID of an LDAP user.
attribute-member-of	string	no	Name of the LDAP attribute used to lookup members of a group.
attribute-uid	string	no	Name of the attribute used to lookup unix UID of an LDAP user.
attribute-user-name	string	no	Name of the LDAP attribute used to lookup a user by user ID.
auth-type	string	yes	Authentication type used while connecting to LDAP servers. This can be Simple or Anonymous.
base-distinguished-name	string	yes	Base distinguished name used as the base for LDAP search requests.
domain-name	string	no	Domain name of the LDAP provider.
ldap-name	string	yes	User defined name of the LDAP provider.
object-class-group	string	no	Name of the LDAP group object class for user accounts.
object-class-user	string	no	Name of the LDAP user object class for user accounts.
port	string	no	Port on which LDAP server is listening on.
preferred-ldap-servers	list of strings	no	Preferred LDAP servers in fully qualified domain name (FQDN) format or IP address format.
use-ssl	bool	no	Specifies whether to use SSL for LDAP connections.

Parameter	Type	Required	Description
user-distinguished-name	string	yes	User distinguished name that is used for LDAP authentication.
user-password	string	yes	User password that is used for LDAP authentication.

ldap delete

To delete an LDAP provider from the cluster.

Parameter	Type	Required	Description
ldap-id	int	yes	Id of the LDAP provider.

ldap update

To update an LDAP provider to the cluster.

Parameter	Type	Required	Description
active-directory-domain-name	string	no	The active directory domain that should be mapped to this LDAP.
auth-type	string	no	Authentication type used while connecting to LDAP servers. This can be Simple or Anonymous.
base-distinguished-name	string	no	Base distinguished name used as the base for LDAP search requests.
domain-name	string	no	Domain name of the LDAP provider.
ldap-id	int	yes	Id of the LDAP provider.
ldap-name	string	no	User defined name of the LDAP provider.
port	string	no	Port on which LDAP server is listening on.
preferred-ldap-servers	list of strings	no	Preferred LDAP servers in fully qualified domain name (FQDN) format or IP address format.

Parameter	Type	Required	Description
use-ssl	bool	no	Specifies whether to use SSL for LDAP connections.
user-distinguished-name	string	no	User distinguished name that is used for LDAP authentication.
user-password	string	no	User password that is used for LDAP authentication.

mfa-config

To run any cluster MFA config related operations.

Operation	Short Name	Description
list	ls	To list MFA config on the cluster.
update	edit	To update MFA config on the cluster.

mfa-config update

To update MFA config on the cluster.

Parameter	Type	Mandatory	Description
authentication-types	list of strings	no	Specifies cluster MFA authentication types: totp, email.
is-cluster-mfa-enabled	bool	no	Specifies if enable/disable cluster MFA.

nfs-share

To run any NFS share related operations.

Operation	Short Name	Description
configure		To configure NFS shares on the NFS server and mount them to all node in the cluster.

Operation	Short Name	Description
create-snapshot		To create snapshots of NFS shares on the NFS server.
delete-remote-storage		To unregister/remove a registeredremote storage.
list		To list all NFS shares in the cluster.
list-remote-storages		To list registered remote storages.
migrate		To migrate a NFS share to a different node in the cluster.
mount		To mount NFS shares to a node in the cluster. If the node ID is not specified, the NFS shares will be distributed and mounted to all nodes in the cluster.
register-remote-storage		To register/add a new remote storage.
update-data-vips		To update the registered FlashBlade data VIPs and use these data VIPs to remount all NFS shares.
update-remote-storage		To update info of a registeredremote storage.

nfs-share configure

To configure NFS shares on the NFS server and mount them to all node in the cluster.

Parameter	Type	Required	Description
available-capacity	int	yes	Total available capacity for creating NFS shares.
data-vips	list of strings	yes	Data VIPs for mounting NFS shares.

nfs-share create-snapshot

To create snapshots of NFS shares on the NFS server.

Parameter	Type	Required	Description
cluster-id	int	no	ID of the Cluster to take snapshot of all NFS shares that belong to the cluster.

Parameter	Type	Required	Description
nfs-share-names	list of strings	no	Comma separated list of NFS share names.
node-id	int	no	ID of the Node to take snapshot of all NFS shares that belong to the node.
suffix	string	no	Suffix to add to the snapshot name.

nfs-share delete-remote-storage

To unregister/remove a registeredremote storage.

Parameter	Type	Required	Description
id	int	yes	The id of registered remote storage.

nfs-share list

To list all NFS shares in the cluster.

Parameter	Type	Required	Description
cluster-id	int	no	ID of the Cluster.

nfs-share migrate

To migrate a NFS share to a different node in the cluster.

Parameter	Type	Required	Description
disk-id	int	yes	ID of the Disk.
node-id	int	yes	ID of the Node to which the NFS share is migrated.

nfs-share mount

To mount NFS shares to a node in the cluster. If the node ID is not specified, the NFS shares will be distributed and mounted to all nodes in the cluster.

Parameter	Type	Required	Description
data-vips	list of strings	yes	Data VIPs for mounting NFS shares.
nfs-share-names	list of strings	yes	Comma separated list of NFS share names.
node-id	int	no	ID of the Node to which the NFS share is mounted.
storage-tier	string	yes	Storage Tier to be assigned to the NFS share. It can be 'PCIeSSD', 'SATA-SSD' or 'SATA-HDD'.

nfs-share register-remote-storage

To register/add a new remote storage.

Parameter	Type	Required	Description
api-token	string	yes	Api token of the management server of remote storage.
ip	string	yes	Management ip address of remote storage.

nfs-share update-data-vips

To update the registered FlashBlade data VIPs and use these data VIPs to remount all NFS shares.

Parameter	Type	Required	Description
data-vips	list of strings	no	Data VIPs for mounting NFS shares.

nfs-share update-remote-storage

To update info of a registeredremote storage.

Parameter	Type	Required	Description
api-token	string	yes	Api token of the management server of remote storage.
id	int	yes	The id of registered remote storage.
ip	string	yes	Management ip address of remote storage.

node-csr

To run any node csr related operations.

Operation	Short Name	Description
create	add	To create a csr for the node.
import		To import and install signed csr and ca cert used to sign the certificate signing request.
update-hostnames		Update hostnames in the cluster when node-ids and hostnames are provided.

node-csr create

To create a csr for the node.

Parameter	Type	Required	Description
city	string	yes	The locality attribute, which is part of the distinguished name definition. It is used to identify the city where the company is located or the Cluster is installed.
common-name	string	no	The common name attribute, which is part of the distinguished name definition. Common name is used to specify a context for the certificate, for example, the name of the Node (attributed by node ip) to which the certificate is to be assigned.
country	string	yes	The country attribute, which is part of the distinguished name definition. It is used to identify the country where the state is located. It is specified as two letter code defined by the ISO standard.
dns-names	list of strings	no	An alternative subject name component to be included in the certificate. It is given as a comma separated list of FQDNs. If not passed, it shall not be set in the certificate request.
email-address	string	no	An alternative subject name component to be included in the certificate. Format is a standard e-mail address, for example joe@company.com.

Parameter	Type	Required	Description
host-ip	string	no	An alternative subject name component to be included in the certificate. The value passed should be the valid ip address of the node.
organization	string	yes	The organization attribute, which is part of the distinguished name definition. It is used to specify the name of the company.
organization-unit	string	yes	Specifies the organization unit attribute, which is part of the distinguished name definition. It is used to identify the specific department or business unit in the company that is owning the Node.
state	string	yes	The state attribute, which is part of the distinguished name definition. It is used to identify the state where the city is located.

node-csr import

To import and install signed csr and ca cert used to sign the certificate signing request.

Parameter	Type	Required	Description
cacert-path	string	yes	The path of the file of the ca certificate which was used to sign the csr for the node.
signedcert-path	string	yes	The path of the file of the signed certificate to be imported for the particular node.

node-csr update-hostnames

Update hostnames in the cluster when node-ids and hostnames are provided.

Parameter	Type	Required	Description
host-names	list of strings	yes	Comma separated list of hostnames. Run allsssh.sh hostname to get all the hostnames, or set the desried hostnames, The order should correspond to the same which was provided in node-ids.
node-ids	list of ints	yes	Type of a node group. Possible values are 1 (Zone), 2 (Bgp config), 3 (Ospf config) and 4 (Dns).

node-group

To run node group related operations.

Operation	Short Name	Description
create-bgp		To create a BGP node group.
create-dns		To create a DNS node group.
create-zone		To create a ZONE node group.
delete	rm	To delete the node group in the cluster.
list	ls	To list the node groups in the cluster.
update-bgp		To update a BGP node group.
update-dns		To update a DNS node group.
update-zone		To update a ZONE node group.

node-group create-bgp

To create a BGP node group.

Parameter	Type	Required	Description
group-name	string	yes	Name of a node group.
hold-time	int	no	BGP instance hold time in seconds.
keepalive-time	int	no	BGP instance keep alive interval in seconds.
local-as	int	yes	BGP peer local AS number.
node-ids	list of ints	yes	Comma separated list of Node IDs.
peer-address	list of strings	yes	Comma separated list of BGP IP addresses.
peer-description	list of strings	no	Comma separated list of BGP peer descriptions.
peer-hold-time	list of ints	no	Comma separated list of BGP peer hold times in seconds.

Parameter	Type	Required	Description
peer-keepalive-time	list of ints	no	Comma separated list of BGP peer keep alive intervals in seconds.
peer-remote-as	list of ints	yes	Comma separated list of BGP peer remote AS numbers.

node-group create-dns

To create a DNS node group.

Parameter	Type	Required	Description
dns-server-ips	list of strings	yes	Comma separated list of DNS server IP addresses.
group-name	string	yes	Name of a node group.
node-ids	list of ints	yes	Comma separated list of Node IDs.

node-group create-zone

To create a ZONE node group.

Parameter	Type	Required	Description
group-name	string	yes	Name of a node group.
node-ids	list of ints	yes	Comma separated list of Node IDs.
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-ip	string	yes	Subnet IP for the Cluster.
subnet-mask-bits	int	yes	Subnet mask bits for the Cluster.

node-group delete

To delete the node group in the cluster.

Parameter	Type	Required	Description
group-name	string	yes	Name of a node group.

node-group list

To list the node groups in the cluster.

Parameter	Type	Required	Description
group-names	list of strings	no	Comma separated list of node group names.
group-type	int	no	Type of a node group. Possible values are 1 (Zone), 2 (Bgp config), 3 (Ospf config) and 4 (Dns).

node-group update-bgp

To update a BGP node group.

Parameter	Type	Required	Description
add-node-ids	list of ints	no	Comma separated list of Node IDs to add.
group-name	string	yes	Name of a node group.
hold-time	int	no	BGP instance hold time in seconds.
keepalive-time	int	no	BGP instance keep alive interval in seconds.
peer-address	list of strings	no	Comma separated list of BGP IP addresses.
peer-description	list of strings	no	Comma separated list of BGP peer descriptions.
peer-hold-time	list of ints	no	Comma separated list of BGP peer hold times in seconds.
peer-keepalive-time	list of ints	no	Comma separated list of BGP peer keep alive intervals in seconds.
peer-remote-as	list of ints	no	Comma separated list of BGP peer remote AS numbers.
remove-node-ids	list of ints	no	Comma separated list of Node IDs to remove.

node-group update-dns

To update a DNS node group.

Parameter	Type	Required	Description
add-node-ids	list of ints	no	Comma separated list of Node IDs to add.
dns-server-ips	list of strings	no	Comma separated list of DNS server IP addresses.
group-name	string	yes	Name of a node group.
remove-node-ids	list of ints	no	Comma separated list of Node IDs to remove.

node-group update-zone

To update a ZONE node group.

Parameter	Type	Required	Description
add-node-ids	list of ints	no	Comma separated list of Node IDs to add.
group-name	string	yes	Name of a node group.
remove-node-ids	list of ints	no	Comma separated list of Node IDs to remove.
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-ip	string	no	Subnet IP for the Cluster.
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.

node

To run any Node related operations.

Operation	Short Name	Description
configure-network		To configure network on nodes before cluster is created.
delete	rm	To mark the Node for removal.
discover-free-nodes	ls-free-nodes	To list the free nodes.
full-info		To get whole info of the node.

Operation	Short Name	Description
identify	ident	To identify a node in a chassis.
list	ls	To list the Node details.
list-packages	ls-pkgs	To list the available Cohesity software packages on a Node.
status		To get the status of the Node.
upload-package	upload-pkg	To upload a software package to the Node.

node configure-network

To configure network on nodes before cluster is created.

Parameter	Type	Required	Description
bonding-mode	string	no	Bonding mode to use when bonding the NICs. Modes 1 (active-backup) and 4 (802.3ad) are supported. Mode 4 may require switch configuration. In cluster, the services need to be stopped before changing on primary interface group.
ipmi-password	string	no	IPMI password for the Cluster.
ipmi-subnet-gateway	string	no	IPMI Subnet gateway.
ipmi-subnet-mask	string	no	IPMI Subnet mask should be either in CIDR prefix or in dotted decimal for IPv4.
ipmi-username	string	no	IPMI username for the Cluster.
node-ids	list of ints	yes	Comma separated list of Node IDs.
node-ipmi-ips	list of strings	no	Comma separated list of Node IPMI IP addresses (IPv4 only).
node-ips	list of strings	yes	Comma separated list of Node IP addresses (IPv4 or IPv6).
subnet-gateway	string	yes	Subnet gateway for the Cluster.
subnet-mask	string	yes	Subnet mask for the Cluster. Should be either in CIDR prefix or in dotted decimal.

node delete

To mark the Node for removal.

Parameter	Type	Required	Description
id	int	yes	ID of a Node.

node identify

To identify a node in a chassis.

Parameter	Type	Required	Description
cluster-id	int	yes	ID of the Cluster.
id	int	yes	ID of a Node.
identify	bool	yes	Flag to toggle identification LED status.

node list

To list the Node details.

Parameter	Type	Required	Description
cluster-partition-ids	list of ints	no	Comma separated list of Cluster Partition IDs.
ids	list of ints	no	Comma separated list of Node IDs.
include-marked-for-removal	bool	no	Include the nodes which have been marked for removal.
include-only-unassigned-nodes	bool	no	Include only the nodes which have not been assigned to any Cluster.
show-stats	bool	no	To display the computed statistics.
show-system-disks	bool	no	To display the system disks details.

node upload-package

To upload a software package to the Node.

Parameter	Type	Required	Description
cancel	bool	no	Cancel any ongoing package upload.
software-pkg-file-path	string	no	Local file path where the Cohesity Software package is available.
software-pkg-md5	string	no	MD5 Hash of the software package to be uploaded to the Node.
software-pkg-url	string	no	URL where Cohesity Software package is available.
software-version	string	no	Software version of the package.

patch

To run any Patch related operations.

Operation	Short Name	Description
apply		To apply dependent patches.
history		To list patch operations.
import		To import a patch or a hotfix.
list	ls	To list Patch details.
revert		To revert dependent patches.
status		To get the status of the current or the last patch operations.

patch apply

To apply dependent patches.

Parameter	Type	Required	Description
all	bool	no	Apply all the available patches.

Parameter	Type	Required	Description
service-name	string	no	Apply the patch for the service name.

patch history

To list patch operations.

Parameter	Type	Required	Description
service-name	string	no	If a service name is given, the patch operations for the service are shown. Otherwise, all operations are shown.

patch import

To import a patch or a hotfix.

Parameter	Type	Required	Description
file-path	string	yes	Local path to the patch or the hotfix file.

patch list

To list Patch details.

Parameter	Type	Required	Description
applied	bool	no	If applied is given, or set to true, shows the applied patches. By default, shows the available patches.
details	bool	no	If details is given, or set to true, shows the fixed issues. By default, shows only the number of fixed issues.
service-name	string	no	If a service name is given, only the patch for the service is shown. Otherwise, all patches are shown.

patch revert

To revert dependent patches.

Parameter	Type	Required	Description
service-name	string	yes	Name of the service whose patch should be reverted.

patch status

To get the status of the current or the last patch operations.

Parameter	Type	Required	Description
details	bool	no	If details is given, or set to true, shows the progress of operation for each service on each node.

privilege

To run any Privilege related operations.

Operation	Short Name	Description
list	ls	To list Privileges.

privilege list

To list Privileges.

Parameter	Type	Required	Description
name	string	no	Name of the Privilege to list

protection-job

To run any Protection Job related operations.

Operation	Short Name	Description
cancel		To cancel a running protection job.

protection-job cancel

To cancel a running protection job.

Parameter	Type	Required	Description
id	int	no	Id of the job.
name	string	no	Name of the job.

qos-profile

To run any Qos profile related operations.

Operation	Short Name	Description
activate		To activate a profile on a list of interfaces with a list of ipsets.
create	add	To create a collection of service ports with the protocol of each port for QoS.
deactivate		To deactivate a list of active QoS profiles, Each one will stop a profile on matching packets.
delete	rm	To remove a list of profiles for QoS.
feature		To enable/disable Qos feature.
list	ls	To list the groups of qos profiles.
list-active	ls-active	To list all the active QoS profiles. Each attachment activates a profile on specific ipsets.
revert-to-default		To revert all QoS profile to factory default configuration

qos-profile activate

To activate a profile on a list of interfaces with a list of ipsets.

Parameter	Type	Required	Description
desc	string	no	Description of the QoS rule sets.
ipset-entries	list of strings	no	Comma separated list of subnets/ips for the firewall Ipsets.

Parameter	Type	Required	Description
mark-dscp	int	yes	To mark dscp value in the outgoing packet.
profile-name	string	yes	Name of the QoS profile.

qos-profile create

To create a collection of service ports with the protocol of each port for QoS.

Parameter	Type	Required	Description
direction	string	no	Data Traffic direction that can be only ""OUTPUT""
force	bool	no	Set to true to unlock the ports restriction.
name	string	yes	Name of the QoS profile.
ports	list of strings	no	Comma separated list of service ports.
priority	int	no	Decides the order in which rule is applied. Lower the number higher the priority.

qos-profile deactivate

To deactivate a list of active QoS profiles, Each one will stop a profile on matching packets.

Parameter	Type	Required	Description
profile-names	list of strings	yes	Comma separated list of names for the QoS profile.

qos-profile delete

To remove a list of profiles for QoS.

Parameter	Type	Required	Description
names	list of strings	yes	Comma separated list of names for the QoS profile.

qos-profile feature

To enable/disable Qos feature.

Parameter	Type	Required	Description
enable	bool	yes	To enable/disable Qos feature.

rack

To manage racks in cluster.

Operation	Short Name	Description
create	add	To add a rack to cluster.
list	ls	To get all racks or one rack information in a cluster.
update	edit	To update a rack.

rack create

To add a rack to cluster.

Parameter	Type	Required	Description
chassis-ids	list of ints	no	List of chassis in a rack.
location	string	no	The location of a rack.
name	string	yes	The name of rack.

rack list

To get all racks or one rack information in a cluster.

Parameter	Type	Required	Description
id	int	no	The id of a rack.

rack update

To update a rack.

Parameter	Type	Required	Description
chassis-ids	list of ints	no	List of chassis in a rack.
id	int	yes	The id of a rack.
location	string	no	The location of a rack.
name	string	no	The name of rack.

remote-cluster

To run any remote Cluster registration related operations.

Operation	Short Name	Description
add-bandwidth-limit-override-window	add-blow	To add a bandwidth limit override window.
add-storage-domain-mapping		To add local Storage Domain to remote Storage Domain mapping. Only one to one mapping is allowed between local Storage Domains and remote Storage Domains.
create	add	To register a remote Cluster.
delete	rm	To delete the registration of a remote Cluster.
delete-storage-domain-mapping		To delete the local Storage Domain to remote Storage Domain mapping.
list	ls	To list the registration information for remote Clusters.
list-remote-storage-domains		To list the Storage Domains of the remote Cluster.
update	edit	To update registration of remote Cluster.

remote-cluster add-bandwidth-limit-override-window

To add a bandwidth limit override window.

Parameter	Type	Required	Description
bandwidth-limit	int	yes	Bandwidth limit in bytes/second from the local Cluster to the remote Cluster.
cluster-id	int	yes	ID of the remote Cluster.
end-time	string	yes	Bandwidth limit override window end time in hh:mm format.
start-time	string	yes	Bandwidth limit override window start time in hh:mm format.
timezone	string	no	Timezone used for bandwidth limit enforcement.
week-days	list of strings	no	Comma separated list of week days in the format su,m,tu,w,th,f,sa.

remote-cluster add-storage-domain-mapping

To add local Storage Domain to remote Storage Domain mapping. Only one to one mapping is allowed between local Storage Domains and remote Storage Domains.

Parameter	Type	Required	Description
cluster-id	int	yes	ID of the remote Cluster.
local-storage-domain-id	int	no	ID of the Local Storage Domain.
local-storage-domain-name	string	no	Name of the Local Storage Domain.
remote-storage-domain-id	int	no	ID of the Remote Storage Domain.
remote-storage-domain-name	string	no	Name of the Remote Storage Domain.

remote-cluster create

To register a remote Cluster.

Parameter	Type	Required	Description
all-endpoints-reachable	bool	no	Whether any endpoint on the remote Cluster is reachable from this cluster.
bandwidth-limit	int	no	Bandwidth limit in bytes/second from the local Cluster to the remote Cluster.
compression-enabled	bool	no	Flag to enable or disable the compression.
encryption-key	string	no	Key for encryption. Enter the same key in the local and remote cluster to set up encryption for replication traffic.
interface-name	string	no	Name of the network interface. It can be bond interface, vlan interface or interface group.
password	string	yes	Password of the user to connect to the remote Cluster.
remote-ips	list of strings	yes	Ip addresses to connect to nodes of remote Cluster
user-name	string	yes	Username of the user to connect to the remote Cluster.

remote-cluster delete

To delete the registration of a remote Cluster.

Parameter	Type	Required	Description
cluster-id	int	yes	ID of the remote Cluster.

remote-cluster delete-storage-domain-mapping

To delete the local Storage Domain to remote Storage Domain mapping.

Parameter	Type	Required	Description
cluster-id	int	yes	ID of the remote Cluster.
local-storage-domain-id	int	no	ID of the Local Storage Domain.

Parameter	Type	Required	Description
local-storage-domain-name	string	no	Name of the Local Storage Domain.
remote-storage-domain-id	int	no	ID of the Remote Storage Domain.
remote-storage-domain-name	string	no	Name of the Remote Storage Domain.

remote-cluster list

To list the registration information for remote Clusters.

Parameter	Type	Required	Description
ids	list of ints	no	Comma separated list of IDs.
names	list of strings	no	Comma separated list of remote Cluster names.

remote-cluster list-remote-storage-domains

To list the Storage Domains of the remote Cluster.

Parameter	Type	Required	Description
cluster-id	int	yes	ID of the remote Cluster.
cluster-partition-ids	list of ints	no	Comma separated list of IDs.
ids	list of ints	no	Comma separated list of IDs.
include-marked-for-removal	bool	no	Include the nodes which have been marked for removal.
names	list of strings	no	Comma separated names of the Storage Domains.
show-stats	bool	no	To display the computed statistics.

remote-cluster update

To update registration of remote Cluster.

Parameter	Type	Required	Description
all-endpoints-reachable	bool	no	Whether any endpoint on the remote Cluster is reachable from this cluster.
bandwidth-limit	int	no	Bandwidth limit in bytes/second from the local Cluster to the remote Cluster.
cluster-id	int	yes	ID of the remote Cluster.
compression-enabled	bool	no	Flag to enable or disable the compression.
encryption-key	string	no	Key for encryption. Enter the same key in the local and remote cluster to set up encryption for replication traffic.
interface-name	string	no	Name of the network interface. It can be bond interface, vlan interface or interface group.
password	string	no	Password of the user to connect to the remote Cluster.
remote-ips	list of strings	no	Ip addresses to connect to nodes of remote Cluster
user-name	string	no	Username of the user to connect to the remote Cluster.

role

To run any Role related operations.

Operation	Short Name	Description
create	add	To create a new Role.
delete	rm	To delete a Role.
list	ls	To list Roles.
update	edit	To update a Role.

role create

To create a new Role.

Parameter	Type	Required	Description
description	string	no	Description of the Role.
name	string	yes	Name of the Role.
privileges	list of strings	no	Comma separated list of Privileges.

role delete

To delete a Role.

Parameter	Type	Required	Description
names	list of strings	yes	Name of the Roles to delete.

role list

To list Roles.

Parameter	Type	Required	Description
name	string	no	Name of the Role.

role update

To update a Role.

Parameter	Type	Required	Description
description	string	no	Description of the Role.
name	string	yes	Name of the Role.
privileges	list of strings	no	Comma separated list of Privileges.

route

To run any Static Route related operations. A Static Route designates which VLAN and Network Interfaces to use for communication with a particular destination network.

Operation	Short Name	Description
create	add	To add a new Static Route.
delete	rm	To remove an existing Static Route.
list	ls	To list all the configured routes.

route create

To add a new Static Route.

Parameter	Type	Required	Description
advmss	int	no	TCP MSS (advmss) value exchanged during connection establishment for the clients reachable by this route.
desc	string	no	Description.
dest-network	string	yes	Destination network: ip/prefix.
interface-group	string	no	Name of a network interface group.
interface-name	string	no	Interface Name such as bond0. It is only used on free node.
mtu	int	no	MTU for the static route.
next-hop	string	yes	Next hop to destination network.
node-group	string	no	Name of a node group. This is only required in multiple subnet cluster.

route delete

To remove an existing Static Route.

Parameter	Type	Required	Description
dest-network	string	yes	Destination network: ip/prefix.
interface-group	string	no	Name of a network interface group.

Parameter	Type	Required	Description
interface-name	string	no	Interface Name such as bond0. It is only used on free node.
node-group	string	no	Name of a node group. This is only required in multiple subnet cluster.

routing

Routing commands.

Operation	Short Name	Description
show-bgp-neighbors		Shows neighbors information of BGP.
show-bgp-nexthop		Shows nexthop information of BGP.
show-bgp-statistics		Shows statistics information of BGP.
show-bgp-summary		Shows BGP summary.

routing show-bgp-neighbors

Shows neighbors information of BGP.

Parameter	Type	Required	Description
all-nodes	bool	no	On all nodes.

routing show-bgp-nexthop

Shows nexthop information of BGP.

Parameter	Type	Required	Description
all-nodes	bool	no	On all nodes.

routing show-bgp-statistics

Shows statistics information of BGP.

Parameter	Type	Required	Description
all-nodes	bool	no	On all nodes.

routing show-bgp-summary

Shows BGP summary.

Parameter	Type	Required	Description
all-nodes	bool	no	On all nodes.

security-config

To run any security config related operations.

Operation	Short Name	Description
update	edit	To update the security config on the cluster.

security-config update

To update the security config on the cluster.

Parameter	Type	Required	Description
ad-mapping	string	no	Specifies which field to map in AD user. It can be chosen from "SamAccountName", "UserPrincipalName".
certificate-mapping	string	no	Specifies which field to map in Certificates. It can be chosen from "CommonName", "EmailAddress", "UserPrincipalName".
classified-data-msg	string	no	Specifies the classified data message.
enable-mapping-based-auth	bool	no	Specifies if mapping based certificate authentication is enabled or not
failed-login-lock-duration-mins	int	no	Specifies the time duration within which the consecutive failed login attempts causes a local user account to be locked.

Parameter	Type	Required	Description
inactivity-time-days	int	no	Specifies the lockout inactivity time range in days.
include-lower-letter	bool	no	Specifies if the password needs to have at least one lowercase letter.
include-number	bool	no	Specifies if the password needs to have at least one number.
include-special-char	bool	no	Specifies if the password needs to have at least one special character.
include-upper-letter	bool	no	Specifies if the password needs to have at least one uppercase letter.
is-data-classified	bool	no	Specifies whether to mark the web page data classified/unclassified.
limit-sessions	bool	no	Specifies if limitations on number of active sessions is enabled or not.
login-token-timeout-hours	int	no	Specifies login token timeout in hours. Should be between 1 and 24 hours.
max-failed-login-attempts	int	no	Specifies the maximum number of consecutive fail login attempts.
max-lifetime-days	int	no	Specifies the maximum password lifetime in days.
min-length	int	no	Specifies the password minimum length.
min-lifetime-days	int	no	Specifies the minimum password lifetime in days.
num-diff-chars	int	no	Specifies the number of characters in the new password that needs to be different from the old password (only applicable when changing the user's own password).
num-disallowed-old-pwd	int	no	Specifies the minimum number of old passwords that are not allowed for changing the password.
session-absolute-timeout	int	no	Specifies absolute session expiration time in seconds.

Parameter	Type	Required	Description
session-inactivity-timeout	int	no	Specifies inactivity session expiration time in seconds.
session-limit-per-user	int	no	Specifies maximum number of active sessions allowed per user. This applies only when limit sessions is enabled.
session-limit-system-wide	int	no	Specifies maximum number of active sessions allowed system wide. This applies only when limit sessions is enabled.

service-identity

To run any Service Identity related operations.

Operation	Short Name	Description
add-disabled-ports		To add disabled ports for service identity
rm-disable-ports		To remove disable ports for service identity
status		Show service identity status

service-identity add-disabled-ports

To add disabled ports for service identity.

Parameter	Type	Mandatory	Description
ports	string	yes	Comma seperated list of disabled ports

service-identity rm-disable-ports

To remove disable ports for service identity.

Parameter	Type	Mandatory	Description
ports	string	yes	Comma seperated list of disabled ports

set-option

To set or override existing CLI options at runtime.

Operation	Short Name	Description
output		To control the CLI output format.

set-option output

To control the CLI output format.

Parameter	Type	Required	Description
format	string	yes	Desired CLI output format at runtime, e.g. default, json, prettyjson.

snmp

To run any SNMP related operations

Operation	Short Name	Description
set-option	set	To set the SNMP settings on the cluster.

snmp set-option

To set the SNMP settings on the cluster.

Parameter	Type	Required	Description
auth-password	string	no	SNMP Authorization Password (default:username)
auth-protocol	string	no	SNMP Auth Protocol (for SNMP v3), possible values: MD5 or SHA
nms-server	string	no	NMS IP or Hostname
operation	string	no	Enable/disable the SNMP, possible values: enable or disable
priv-password	string	no	SNMP Privacy Password (default: authPassword)
priv-protocol	string	no	SNMP Priv Protocol (for SNMP v3), possible values: AES or DES in Galois/Counter Mode.

Parameter	Type	Required	Description
security-level	string	no	SNMP Security Level(for SNMP v3), possible values: noAuthNoPriv or authNoPriv or authPriv
sys-contact	string	no	System Contact
sys-desc	string	no	System Description
sys-engine-id-type	int	no	SNMP engineIDType specifies the engineIDshould be built from. Possible values: 0(Random), 1 (IPv4 address, default), 2 (IPv6 address), 3(MAC address).
sys-location	string	no	System Location
sys-name	string	no	System Name
sys-objid	string	no	System Object Id
user	string	no	SNMP User(for v3) or Community(for v2) Name
user-type	string	no	SNMP User type, possible values: read or write or trap
version	string	no	SNMP version, possible values: 2 or 3

storage-domain

To run any Storage Domain related operations.

Operation	Short Name	Description
add-subnet-to-whitelist	add-st-wl	To add a subnet to a view's subnet whitelist
create	add	To create a new Storage Domain.
get-subnet-whitelist	ls-st-wl	To get a view's subnet whitelist
list	ls	To list Storage Domain details.
remove-subnet-from-whitelist	rm-st-wl	To remove a subnet from a view's subnet whitelist
update	edit	To update editable information of a Storage Domain.

storage-domain add-subnet-to-whitelist

To add a subnet to a view's subnet whitelist

Parameter	Type	Required	Description
desc	string	no	Description.
id	int	no	ID of a Storage Domain.
name	string	no	Name of the Storage Domain.
subnet-ip	string	yes	Subnet IP for the Cluster.
subnet-mask	string	yes	Subnet mask for the Cluster. Should be either in CIDR prefix or in dotted decimal.
subnet-nfs-permission	string	no	Clients from this subnet can mount using NFS protocol. This can be kDisabled, kReadOnly, kReadWrite.
subnet-nfs-root-squash	bool	no	Clients from this subnet can mount as root on NFS. This can be Disabled, ReadOnly, ReadWrite.
subnet-s3-permission	string	no	Clients from this subnet can access using S3 protocol. This can be kDisabled, kReadOnly, kReadWrite.
subnet-smb-permission	string	no	Clients from this subnet can mount using SMB protocol. This can be kDisabled, kReadOnly, kReadWrite.

storage-domain create

To create a new Storage Domain.

Parameter	Type	Required	Description
active-directory-domain-name	string	no	The active directory domain this view box should be mapped to. Users from this active directory will be given permission to view contents in this storage domain, when accessing through SMB.
app-marker-detection	bool	no	Whether to support app marker detection.
brick-size	int	no	Brick size in bytes used to create blobs.

Parameter	Type	Required	Description
cluster-partition-id	int	no	ID of a Cluster Partition.
cluster-partition-name	string	no	Name of a Cluster Partition.
compression	string	no	Level of compression to use. Values are low or off.
dedup-enabled	bool	no	Flag to turn on/off post-processing deduplication. Deduplication is on by default. Inline deduplication cannot be turned on, if post-processing deduplication is off.
default-user-quota	int	no	Default user quota in GiB for this Storage Domain.
default-user-quota-alert-limit	int	no	Default user quota limit in GiB for this Storage Domain.
default-user-quota-alert-threshold-percentage	int	no	Default user quota threshold alert percentage for this Storage Domain.
default-view-quota	int	no	Default quota in GiB for views created in the Storage Domain. Set -1 to remove the value.
default-view-quota-alert-limit	int	no	Default alerting threshold in GiB for views created in the Storage Domain. Set -1 to remove the value.
ec-num-data	int	no	Number of data stripes for erasure coding.
ec-num-failures	int	no	Number of failures to tolerate for erasure coding.
encryption-enabled	bool	no	Flag to turn on encryption. If encryption is on, it cannot be turned off later.
erasure-coding-enabled	bool	no	To enable erasure coding.

Parameter	Type	Required	Description
inline-compress-enabled	bool	no	Flag to turn on/off inline compression.If deduplication is done inline, then it is recommended that compression should also be done inline. Also, this parameter will have no effect unless compression is enabled.
inline-dedup-enabled	bool	no	Flag to turn on/off inline deduplication. This parameter will have no effect, unless deduplication is enabled.
inline-erasure-coding	bool	no	Whether erasure coding should be done inline.
ldap-id	int	no	Id of the LDAP provider.
name	string	yes	Name of the Storage Domain.
num-failures	int	no	Disk fault tolerance factor for this storage domain. Specifies maximum number of disk faults tolerated.
num-node-failures-tolerated	int	no	Node fault tolerance factor for this storage domain. Specifies maximum number of node faults tolerated.
purpose	string	yes	To inherit the recommended configurations and tunables optimized for NetBackup DirectIO.
quota	int	no	Maximum allowed quota in GiB. Set -1 to remove the value.
quota-alert-limit	int	no	Threshold in GiB, over which an alert will be sent. Set -1 to remove the value.
s3-buckets-allowed	bool	no	Whether creation of S3 bucket is allowed in the Storage Domain.
subnet-ip	string	no	Subnet IP for the Cluster.
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.
subnet-mask-ip4	string	no	Subnet mask for IPv4 type subnet of the Cluster.
subnet-nfs-permission	string	no	Clients from this subnet can mount using NFS protocol. This can be kDisabled, kReadOnly, kReadWrite.

Parameter	Type	Required	Description
subnet-nfs-root-squash	bool	no	Clients from this subnet can mount as root on NFS. This can be Disabled, ReadOnly, ReadWrite.
subnet-s3-permission	string	no	Clients from this subnet can access using S3 protocol. This can be kDisabled, kReadOnly, kReadWrite.
subnet-smb-permission	string	no	Clients from this subnet can mount using SMB protocol. This can be kDisabled, kReadOnly, kReadWrite.
vault-id	int	no	ID of the Vault.
vault-name	string	no	Name of the Vault.

storage-domain get-subnet-whitelist

To get a view's subnet whitelist

Parameter	Type	Required	Description
id	int	no	ID of a Storage Domain.
name	string	no	Name of the Storage Domain.

storage-domain list

To list Storage Domain details.

Parameter	Type	Required	Description
cluster-partition-ids	list of ints	no	Comma separated list of IDs.
ids	list of ints	no	Comma separated list of IDs.
include-marked-for-removal	bool	no	Include the nodes which have been marked for removal.
names	list of strings	no	Comma separated names of the Storage Domains.
show-stats	bool	no	To display the computed statistics.

storage-domain remove-subnet-from-whitelist

To remove a subnet from a view's subnet whitelist

Parameter	Type	Required	Description
id	int	no	ID of a Storage Domain.
name	string	no	Name of the Storage Domain.
subnet-ip	string	yes	Subnet IP for the Cluster.

storage-domain update

To update editable information of a Storage Domain.

Parameter	Type	Required	Description
active-directory-domain-name	string	no	The active directory domain this view box should be mapped to. Users from this active directory will be given permission to view contents in this storage domain, when accessing through SMB.
app-marker-detection	bool	no	Whether to support app marker detection.
brick-size	int	no	Brick size in bytes used to create blobs.
compression	string	no	Level of compression to use. Values are low or off.
dedup-enabled	bool	no	Flag to turn on/off post-processing deduplication. Deduplication is on by default. Inline deduplication cannot be turned on, if post-processing deduplication is off.
default-user-quota	int	no	Default user quota in GiB for this Storage Domain.
default-user-quota-alert-limit	int	no	Default user quota limit in GiB for this Storage Domain.
default-user-quota-alert-threshold-percentage	int	no	Default user quota threshold alert percentage for this Storage Domain.

Parameter	Type	Required	Description
default-view-quota	int	no	Default quota in GiB for views created in the Storage Domain. Set -1 to remove the value.
default-view-quota-alert-limit	int	no	Default alerting threshold in GiB for views created in the Storage Domain. Set -1 to remove the value.
ec-num-data	int	no	Number of data stripes for erasure coding.
ec-num-failures	int	no	Number of failures to tolerate for erasure coding.
erasure-coding-enabled	bool	no	To enable erasure coding.
id	int	no	ID of a Storage Domain.
inline-compress-enabled	bool	no	Flag to turn on/off inline compression.If deduplication is done inline, then it is recommended that compression should also be done inline. Also, this parameter will have no effect unless compression is enabled.
inline-dedup-enabled	bool	no	Flag to turn on/off inline deduplication. This parameter will have no effect, unless deduplication is enabled.
inline-erasure-coding	bool	no	Whether erasure coding should be done inline.
ldap-id	int	no	Id of the LDAP provider.
name	string	no	Name of the Storage Domain.
new-name	string	no	New name of the Storage Domain to update.
num-failures	int	no	Disk fault tolerance factor for this storage domain. Specifies maximum number of disk faults tolerated.
num-node-failures-tolerated	int	no	Node fault tolerance factor for this storage domain. Specifies maximum number of node faults tolerated.

Parameter	Type	Required	Description
quota	int	no	Maximum allowed quota in GiB. Set -1 to remove the value.
quota-alert-limit	int	no	Threshold in GiB, over which an alert will be sent. Set -1 to remove the value.
remove-active-directory-domain-mapping	bool	no	Removes the current mapping between this storage domain to the Active Directory domain it is currently mapped to.
remove-ldap-provider-mapping	bool	no	Removes the current mapping between this storage domain to the LDAP provider it is currently mapped to.
s3-buckets-allowed	bool	no	Whether creation of S3 bucket is allowed in the Storage Domain.
vault-id	int	no	ID of the Vault.
vault-name	string	no	Name of the Vault.
iris_cli storage-domain update name=NameOfStorageDomain compression=high	string	no	Sets the compression level to high on this Storage Domain.

support-bundle

To generate support bundle.

Operation	Short Name	Description
generate		To generate support bundle.

support-server

To run Support server and Support channel related operations.

Operation	Short Name	Description
add-authorized-user		To add an authorized user to login to cluster via support server.
host-access		To enable/disable host shell access.
host-access-token		To generate access token inorder to get access code from Cohesity for gaining host shell access
list	ls	List details of the Cohesity support server.
ls-host-access		List details of the Cohesity node access.
remove-authorized-user		To remove an authorized user to login to cluster via support server.
support-channel	rt	To enable/disable the support channel.
update	edit	Update the details of the Cohesity support server.

support-server add-authorized-user

To add an authorized user to login to cluster via support server.

Parameter	Type	Required	Description
ssh-key-file-path	string	yes	Absolute file path for the public SSH keys that allows access to the Cluster.
user-name	string	yes	Username of the User.
version	int	no	Version of Cohesity RT server. If set to 1(default), then the config will be applied to server version 1. If set to 2, then the config will be applied to server version 2.

support-host-access

To enable/disable host shell access.

Parameter	Type	Required	Description
access-code	string	yes	Node Acces code supplied by support team.
ssh-enable	bool	yes	Enable/Disable ssh (special port) access.

support-server list

List details of the Cohesity support server.

Parameter	Type	Required	Description
version	int	no	Version of Cohesity RT server. If set to 1(default), then the config will be applied to server version 1. If set to 2, then the config will be applied to server version 2.

support-server remove-authorized-user

To remove an authorized user to login to cluster via support server.

Parameter	Type	Required	Description
clear-all	bool	no	To clear all authorized users of support server.
user-name	string	no	Username of the User.
version	int	no	Version of Cohesity RT server. If set to 1(default), then the config will be applied to server version 1. If set to 2, then the config will be applied to server version 2.

support-server support-channel

To enable/disable the support channel.

For example: To enable support channel, run the following command:

```
iris_cli support support-channel enable=1 duration-in-hours=1
```

Parameter	Type	Required	Description
duration-in-hours	int	no	Duration in hours from current time till when the support channel will be enabled. Valid values are 1 - 720 hours. **Note:** duration-in-hours is mandatory when enable=1 is entered.
enable	bool	yes	Enable/Disable the support channel.

support-server update

Update the details of the Cohesity support server.

Parameter	Type	Required	Description
hostname	string	no	Hostname for the Cohesity support server.
port	int	no	Modify the port number for support channel.
private-key-file-loc	absolute file path	no	Absolute file path of the private key (SSH) for the support channel.
server-config-params	absolute file path	no	Server config params in a JSON file.
session-recording	bool	no	Enable/Disable the support channel session recording.
username	string	no	Username for the support channel.
version	int	no	Version of Cohesity RT server. If set to 1(default), then the config will be applied to server version 1. If set to 2, then the config will be applied to server version 2.

swift

To run any Swift related operations.

Operation	Short Name	Description
register		To register a Swift service on Keystone server.
unregister		To unregister a Swift service on Keystone server.

swift register

To register a Swift service on Keystone server.

Parameter	Type	Required	Description
domain-name	string	no	Specifies the domain name for 'Domain' scope. It only applys when 'scope-type' is 'Domain'.

Parameter	Type	Required	Description
password	string	no	Specifies the password of Keystone user.
project-domain-name	string	no	Spcifies the domain name for 'Project' scope. It only applys when 'scope-type' is 'Project'.
project-name	string	no	Specifies the project name for 'Project' scope. It only applys when 'scope-type' is 'Project'.
scope-type	string	no	Specifies the scope type. It can be 'Project' or 'Domain'.
tenant-id	string	yes	Specifies the id of tenant whose Swift service will be registered.
user-domain	string	no	Specifies the domain of Keystone user.
username	string	no	Specifies the Keystone username.

swift unregister

To unregister a Swift service on Keystone server.

Parameter	Type	Required	Description
domain-name	string	no	Specifies the domain name for 'Domain' scope. It only applys when 'scope-type' is 'Domain'.
password	string	no	Specifies the password of Keystone user.
project-domain-name	string	no	Spcifies the domain name for 'Project' scope. It only applys when 'scope-type' is 'Project'.
project-name	string	no	Specifies the project name for 'Project' scope. It only applys when 'scope-type' is 'Project'.
scope-type	string	no	Specifies the scope type. It can be 'Project' or 'Domain'.
tenant-id	string	yes	Specifies the id of tenant whose Swift service will be unregistered.
user-domain	string	no	Specifies the domain of Keystone user.
username	string	no	Specifies the Keystone username.

sysctl

Command to update sysctl configuration for nodes.

Operation	Short Name	Description
update	edit	Update sysctl configuration to nodes or cluster

sysctl update

Update sysctl configuration to nodes or cluster

Parameter	Type	Required	Description
sysctl-conf-names	list of strings	yes	Comma separated list of sysctl config names.
value	list of strings	yes	Comma separated list of sysctl config values.

syslog-server

To run any Syslog Server related operations.

Operation	Short Name	Description
create	add	To add a Syslog Server to the Cluster. Cohesity audit and some common audits will be added by default
delete	rm	To remove an existing Syslog Server from the Cluster.
list	ls	To list all the Syslog Servers in the Cluster.
status		To check the syslog server status in cluster.
update	edit	To update an existing Syslog Server in the Cluster.
update-audit-tags		To update audit tags used in cluster.

syslog-server create

To add a Syslog Server to the Cluster. Cohesity audit and some common audits will be added by default

Parameter	Type	Required	Description
address	string	yes	Hostname or IP address of the Syslog Server.
ca-certificate-path	string	no	Specifies the syslog server ca certificate file path
cluster-audit	bool	no	Deprecated. Send Cluster Audit Logs to the Syslog Server. Default is true.
dataprotection-events	bool	no	Deprecated. Send Dataprotection Event Logs to the Syslog Server. Default is false.
enable	bool	no	Specifies whether to enable the syslog server.
facilities	list of strings	no	The logs of specified facilites sent to the syslog server.
filer-audit	bool	no	Deprecated. Send Filer Audit Logs to the Syslog Server. Default is true.
message-patterns	list of strings	no	The message with the patterns will be sent to the syslog server.
name	string	no	Name or Description of the Syslog Server.
port	int	yes	Port where the Syslog Server listens.
program-names	list of strings	no	The logs of specified programes sent to the syslog server.
protocol	string	no	Protocol used by the Syslog Server. Default is tcp.
raw-message-patterns	list of strings	no	The raw message with the patterns will be sent to the syslog server.
ssh-events-log	bool	no	Deprecated. Send ssh login event logs to the syslog server. Default is false.

syslog-server delete

To remove an existing Syslog Server from the Cluster.

Parameter	Type	Required	Description
all-syslog-servers	bool	no	Select all syslog servers in cluster.
id	int	no	The Id of a syslog server.

syslog-server list

To list all the Syslog Servers in the Cluster.

Parameter	Type	Required	Description
id	int	no	The Id of a syslog server.

syslog-server status

To check the syslog server status in cluster.

Parameter	Type	Required	Description
ids	list of ints	no	List of syslog server Ids.

syslog-server update

To update an existing Syslog Server in the Cluster.

Parameter	Type	Required	Description
address	string	no	Hostname or IP address of the Syslog Server.
ca-certificate-path	string	no	Specifies the syslog server ca certificate file path
cluster-audit	bool	no	Deprecated. Send Cluster Audit Logs to the Syslog Server. Default is true.
dataprotection-events	bool	no	Deprecated. Send Dataprotection Event Logs to the Syslog Server. Default is false.
enable	bool	no	Specifies whether to enable the syslog server.
facilities	list of strings	no	The logs of specified facilites sent to the syslog server.
filer-audit	bool	no	Deprecated. Send Filer Audit Logs to the Syslog Server. Default is true.
id	int	yes	The Id of a syslog server.
is-append	bool	no	Append logging parameters to the syslog server.

Parameter	Type	Required	Description
message-patterns	list of strings	no	The message with the patterns will be sent to the syslog server.
name	string	no	Name or Description of the Syslog Server.
port	int	no	Port where the Syslog Server listens.
program-names	list of strings	no	The logs of specified programes sent to the syslog server.
protocol	string	no	Protocol used by the Syslog Server. Default is tcp.
raw-message-patterns	list of strings	no	The raw message with the patterns will be sent to the syslog server.
ssh-events-log	bool	no	Deprecated. Send ssh login event logs to the syslog server. Default is false.

syslog-server update-audit-tags

To update audit tags used in cluster.

Parameter	Type	Required	Description
alert-audit-tag-name	string	no	Paramter to specify alert audit tag.
cluster-audit-tag-name	string	no	Paramter to specify cluster-audit tag.
dataprotection-events-tag-name	string	no	Paramter to specify dataprotection-events tag.
filer-audit-tag-name	string	no	Paramter to specify filer-audit tag.

system-app

To run any System app related operations.

Operation	Short Name	Description
Status		To get the status of the System apps.

trusted-cas

To run any trusted CA related operations.

Operation	Short Name	Description
list	ls	To list all trusted CAs on the cluster.
register		To register a trusted CA on the cluster.
unregister		To unregister a trusted CA on the cluster.
validate		To re-validate a trusted CA on the cluster.

trusted-cas list

To list all trusted CAs on the cluster.

Parameter	Type	Required	Description
ids	list of strings	no	Comma seperated list of trusted CA Ids to fetch.
names	list of strings	no	Comma seperated list of trusted CA Names to fetch.

trusted-cas register

To register a trusted CA on the cluster.

Parameter	Type	Required	Description
desc	string	no	Specifies Description of certificate being registered.
file-path	string	yes	Specifies filepath to PEM file with certificate to be registered.
name	string	yes	Specifies name of the certificate being registered.
only-validate	bool	no	If true, the certificates are only validated not registered.

trusted-cas unregister

To unregister a trusted CA on the cluster.

Parameter	Type	Required	Description
id	string	yes	Specifies the Id of trusted CA to be operated on.

trusted-cas validate

To re-validate a trusted CA on the cluster.

Parameter	Type	Required	Description
id	string	yes	Specifies the Id of trusted CA to be operated on.

uda-connector

To run any UDA connector related operations.

Operation	Short Name	Description
create	add	To add a new UDA connector
delete	rm	To delete a UDA connector
get		To get a list of UDA connectors in the system
update	edit	To update a UDA connector

uda-connector create

To add a new UDA connector.

Parameter	Type	Required	Description
package-path	string	yes	Path of the folder containing all the json config files for a UDA connector

uda-connector delete

To delete a UDA connector.

Parameter	Type	Required	Description
id	int	yes	Unique id of a UDA connector

uda-connector get

connectors in the system.

Parameter	Type	Required	Description
filter-index-config	bool	no	If true, only prints out the index config
id	int	no	Unique id of a UDA connector

uda-connector update

To update a UDA connector.

Parameter	Type	Required	Description
package-path	string	yes	Path of the folder containing all the json config files for a UDA connector

user-quotas

To run any user quotas related commands.

Operation	Short Name	Description
create	add	To create a user quota policy in a view.
delete	rm	To delete a user quota policy in a view.
list		To list user quota policies for a view, a user or both.
update	edit	To update a user quota policy in a view.

user-quotas create

To create a user quota policy in a view.

Parameter	Type	Required	Description
alert-limit	int	no	Threshold in GiB, over which an alert will be sent. Set -1 to remove the value.

Parameter	Type	Required	Description
alert-threshold-percentage	int	no	Threshold percentage over which an alert will be sent.
hard-limit	int	no	Maximum allowed quota in GiB. Set -1 to remove the value.
user-sid	string	no	The SID of the user.
user-unix-uid	int	no	The Unix User ID of the user.
view-name	string	yes	Name of the View.

user-quotas delete

To delete a user quota policy in a view.

Parameter	Type	Required	Description
all	bool	no	Deletes all user quota policies for this view.
user-sids	list of strings	no	Comma separated list of user SIDs.
user-unix-uids	list of ints	no	Comma separated list of UNIX User IDs.
view-name	string	yes	Name of the View.

user-quotas list

To list user quota policies for a view, a user or both.

Parameter	Type	Required	Description
cookie	string	no	Opaque string returned in a previous command output to get the next set of entries in the output.
exclude-within-alert-threshold	bool	no	Exclude the users whose usage is within the alert threshold limit.
include-usage	bool	no	Include logical usage of users in this view.
page-count	int	no	Limit the number of entries in the command output.

Parameter	Type	Required	Description
summary	bool	no	To fetch the summary of user quotas for a view or a user.
user-sid	string	no	The SID of the user.
user-unix-uid	int	no	The Unix User ID of the user.
view-name	string	no	Name of the View.

user-quotas update

To update a user quota policy in a view.

Parameter	Type	Required	Description
alert-limit	int	no	Threshold in GiB, over which an alert will be sent. Set -1 to remove the value.
alert-threshold-percentage	int	no	Threshold percentage over which an alert will be sent.
default-alert-limit	int	no	Default alert limit after which an alert is generated.
default-alert-threshold-percentage	int	no	Default threshold percentage of hard limit over which an alert is generated.
default-hard-limit	int	no	Default hard limit in bytes.
enable	bool	no	Enable user quota policies for this view.
hard-limit	int	no	Maximum allowed quota in GiB. Set -1 to remove the value.
user-sid	string	no	The SID of the user.
user-unix-uid	string	no	The Unix User ID of the user.
view-name	string	yes	Name of the View.

user

To run any User related operations.

Operation	Short Name	Description
create	add	To create a new User.
delete	rm	To delete the specified users.
delete-sessions		Deletes all active sessions for a user or all active system wide sessions
get-linux-user-mfa	lingetmfa	Get MFA config for linux support user.
linux-bash-access		To get bash access for the Linux 'support' User.
linux-sudo-access		To get sudo access for the Linux 'support' User.
linux-user	linusr	Updates password for linux user support
list	ls	To list User details.
reset-s3-secret-key		To reset the S3 account secret access key for a User.
update	edit	To update editable information of a User.
update-linux-user-mfa	linupdmfa	Updates MFA mode for linux support user.

user create

To create a new User.

Parameter	Type	Required	Description
additional-group-names	list of strings	no	Comma separated list of additional group names of the User.
domain	string	no	Domain identifies User's logging in domain.
effective-time	datetime	no	Effective time in "YYYY-MM-DDTHH:mm" format when User's credentials will be effective.
email-address	string	no	Email address of a User.
expiry-time	datetime	no	Expire time of the User in "YYYY-MM-DDTHH:mm" format

Parameter	Type	Required	Description
password	string	no	Password of the User.
primary-group-name	string	no	Primary group name of the User.
roles	string	no	Comma separated list of supported roles for users (valid roles are ADMIN, OPS, VIEW).
user-description	string	no	Description of the User.
user-name	string	yes	Username of the User.

user delete

To delete the specified users.

Parameter	Type	Required	Description
domain	string	no	Domain identifies User's logging in domain.
user-names	list of strings	yes	Comma separated list of User names of the existing users.

user delete-sessions

Deletes all active sessions for a user or all active system wide sessions

Parameter	Type	Required	Description
user-sid	string	no	The SID of the user.

user linux-sudo-access

To get sudo access for the Linux 'support' User.

Parameter	Type	Required	Description
enable	bool	yes	If set to true, sudo access will be enabled.
duration-in-hours	int	yes	Defines the number of hours for which sudo access remains active.

user linux-user

Updates password for linux user support

Parameter	Type	Required	Description
current-password	string	no	Current linux password. Needed only when updating the password.
linux-password	string	yes	New linux password
linux-username	string	yes	Linux username

user list

To list User details.

Parameter	Type	Required	Description
domain	string	no	Domain identifies User's logging in domain.
email-addresses	list of strings	no	Comma separated list of email addresses of the users.
user-names	list of strings	no	Comma separated list of User names of the existing users.

user reset-s3-secret-key

To reset the S3 account secret access key for a User.

Parameter	Type	Required	Description
domain	string	no	Domain identifies User's logging in domain.
user-name	string	yes	Username of the User.

user update

To update editable information of a User.

Parameter	Type	Required	Description
additional-group-names	list of strings	no	Comma separated list of additional group names of the User.
current-password	string	no	Current password of the User.
domain	string	no	Domain identifies User's logging in domain.
effective-time	datetime	no	Effective time in "YYYY-MM-DDTHH:mm" format when User's credentials will be effective.
email-address	string	no	Email address of a User.
expiry-time	datetime	no	Expire time of the User in "YYYY-MM-DDTHH:mm" format
is-account-locked	bool	no	Indicate whether to lock/unlock a user account.
is-user-exempt-from-mfa	bool	no	Indicate whether to disable MFA for a user.
new-user-name	string	no	Name to which the users name should be changed.
password	string	no	Password of the User.
primary-group-name	string	no	Primary group name of the User.
remove-additional-groups	bool	no	To remove an existing additional groups of the user.
remove-primary-group	bool	no	To remove an existing primary group of the user.
roles	string	no	Comma separated list of supported roles for users (valid roles are ADMIN, OPS, VIEW).
user-description	string	no	Description of the User.
user-name	string	yes	Username of the User.

user update-linux-user-mfa

Updates MFA mode for linux support user.

Parameter	Type	Required	Description
authentication-type	string	no	Specifies Linux support user mfa authentication type: either totp, email.
current-password	string	no	Current linux password. Needed only when updating the password.
disable	bool	no	Specifies whether MFA for linux support user should be disabled.

vault

To run any Vault related operations.

Operation	Short Name	Description
create	add	To add a new Vault.
list	ls	To list the existing Vaults.
update	edit	To update an existing Vault.

vault create

To add a new Vault.

Parameter	Type	Required	Description
aws-access-key-id	string	no	AWS access key id for AWS service account.
aws-c2s-agency	string	no	AWS C2S agency for AWS C2S service account.
aws-c2s-base-url	string	no	AWS C2S base-url for AWS C2S service account.
aws-c2s-ca-cert-key-file-loc	absolute file path	no	Absolute file path for CA Cert key of AWS C2S service account.
aws-c2s-client-cert-key-file-loc	absolute file path	no	Absolute file path for client certificate key of AWS C2S service account.

Parameter	Type	Required	Description
aws-c2s-client-priv-key-file-loc	absolute file path	no	Absolute file path for client private key of AWS C2S service account.
aws-c2s-mission	string	no	AWS C2S mission for AWS C2S service account.
aws-c2s-role	string	no	AWS C2S role for AWS C2S service account.
aws-c2s-secret-key	string	no	AWS C2S secret key for AWS C2S service account.
aws-credential-blob	absolute file path	no	AWS credential blob to authenticate with credential endpoint.
aws-credential-endpoint	string	no	AWS credential REST endpoint that generates the security token.
aws-region	string	no	AWS region (e.g. us-east-1, us-west-1 etc.) for AWS service account.
aws-secret-access-key	string	no	AWS secret access key for AWS service account.
azure-client-id	string	no	Azure Client Id.
azure-storage-access-key	string	no	Storage access key for Azure service account.
azure-storage-account-name	string	no	Storage account name for Azure service account.
bucket-name	string	no	Name of the bucket or container inside the target Vault.
compression-enabled	bool	yes	Flag to enable or disable the compression.
dedup-enabled	bool	no	Flag to turn on/off post-processing deduplication. Deduplication is on by default. Inline deduplication cannot be turned on, if post-processing deduplication is off.
desc	string	no	Description of the Vault.

Parameter	Type	Required	Description
encryption-enabled	bool	yes	Flag to enable or disable the encryption.
forever-incremental-enabled	bool	no	Support forever incremental archive for archival vaults.
goog-client-email	string	no	Client email of Google service account.
goog-client-priv-key-file-loc	absolute file path	no	Absolute file path for client private key of Google service account.
goog-project-id	string	no	Project ID of Google service account.
incremental-enabled	bool	no	Support incremental archive for archival vaults.
kms-server-id	int	no	The Id of a KMS server.
name	string	yes	Name of the Vault.
nas-host	string	no	IP or host name of the NAS service host.
nas-kerberos-realm-name	string	no	NAS Kerberos realm name.
nas-mount-path	string	no	Mount path on the NAS service host.
nas-nfs-security-type	string	no	NAS NFS security type.
nas-nfs-version-number	string	no	NAS NFS version number.
nas-password	string	no	NAS service password.
nas-share-type	string	no	Mounted volume share type, e.g. CIFS, NFS. Default is NFS.
nas-username	string	no	NAS service username.
oracle-access-key-id	string	no	Access key id for Oracle service account.

Parameter	Type	Required	Description
oracle-region	string	no	Region (e.g. us-east-1) for Oracle service account.
oracle-secret-access-key	string	no	Secret access key for Oracle service account.
oracle-tenant	string	no	Tenant for Oracle service account.
qstar-host	string	no	IP or DNS name of QStar service host.
qstar-integral-volumes	list of strings	no	Comma separated names of QStar Integral Volumes.
qstar-password	string	no	QStar service password.
qstar-port	int	no	QStar WEB services API port.
qstar-share-type	string	no	Mounted QStar Integral Volume share type, e.g. CIFS, NFS. Default is CIFS.
qstar-use-https	bool	no	Use https for QStar service. Default is false.
qstar-username	string	no	QStar service username.
s3-compatible-service-url	string	no	Service URL for S3 compatible cloud.
s3-compatible-signature-version	int	no	Signature version fo S3 compatible cloud. Default is 4.
s3-compatible-use-https	bool	no	Use https for S3 compatible cloud. Default is true.

Parameter	Type	Required	Description
tier-type	string	no	Tier of the vault. For AWS, supported tier-types are Amazon-S3-standard, Amazon-S3-standard-IA, Amazon-S3-One-Zone-IA, Amazon-S3-Intelligent, Amazon-S3-Glacier, Amazon-S3-Glacier-Deep-Archive, Amazon-glacier and Amazon-S3-Glacier-IR. For Azure, supported tier types are Azure-archive, Azure-cool and Azure-hot. For GCP, supported tier types are Google-nearline, Google-coldline, Google-regional and Google-multi-regional. For Oracle, supported tier types are Oracle-standard and Oracle-archive.
usage-type	string	yes	Usage purpose of the Vault. It can be either archival or cloud-spill.
vault-type	string	no	Type of the Vault. For Archival usage, valid types are Azure, Azure-Gov-Cloud, Amazon, AWS-Gov-Cloud, Amazon-C2S, Google, Oracle, S3-compatible, NAS, and QStar-Tape. For Cloud-spill usage, valid types are NAS, S3-compatible, for AWS, vaults are Amazon, AWS-Gov-Cloud and supported tier is Amazon-S3-standard, for Azure, vaults are Azure, Azure-Gov-Cloud and supported tier is Azure-hot, for GCP, vault is Google and supported tiers are Google-regional and Google-multi-regional, for Oracle vault is Oracle and supported tier is Oracle-standard.
view-box-name	string	no	Name of the Storage Domain.

vault list

To list the existing Vaults.

Parameter	Type	Required	Description
id	int	no	ID of the Vault.
name	string	no	Name of the Vault.

vault update

To update an existing Vault.

Parameter	Type	Required	Description
aws-access-key-id	string	no	AWS access key id for AWS service account.
aws-credential-blob	absolute file path	no	AWS credential blob to authenticate with credential endpoint.
aws-credential-endpoint	string	no	AWS credential REST endpoint that generates the security token.
aws-region	string	no	AWS region (e.g. us-east-1, us-west-1 etc.) for AWS service account.
aws-secret-access-key	string	no	AWS secret access key for AWS service account.
azure-client-id	string	no	azure-client-id
azure-storage-access-key	string	no	Storage access key for Azure service account.
azure-storage-account-name	string	no	Storage account name for Azure service account.
compression-enabled	bool	no	Flag to enable or disable the compression.
desc	string	no	Description of the Vault.
encryption-enabled	bool	no	Flag to enable or disable the encryption.
forever-incremental-enabled	bool	no	Support forever incremental archive for archival vaults.
full-archive-interval-days	int	no	Interval in number of days between two Full archives.
goog-client-email	string	no	Client email of Google service account.

Parameter	Type	Required	Description
goog-client-priv-key-file-loc	absolute file path	no	Absolute file path for client private key of Google service account.
id	int	yes	ID of the Vault.
incremental-enabled	bool	no	Support incremental archive for archival vaults.
kms-server-id	int	no	The Id of a KMS server.
nas-host	string	no	IP or host name of the NAS service host.
nas-kerberos-realm-name	string	no	NAS Kerberos realm name.
nas-mount-path	string	no	Mount path on the NAS service host.
nas-nfs-security-type	string	no	NAS NFS security type.
nas-nfs-version-number	string	no	NAS NFS version number.
nas-password	string	no	NAS service password.
nas-share-type	string	no	Mounted volume share type, e.g. CIFS, NFS. Default is NFS.
nas-username	string	no	NAS service username.
new-name	string	no	Name of the Vault.
qstar-host	string	no	IP or DNS name of QStar service host.
qstar-integral-volumes	list of strings	no	Comma separated names of QStar Integral Volumes.
qstar-password	string	no	QStar service password.
qstar-port	int	no	QStar WEB services API port.

Parameter	Type	Required	Description
qstar-share-type	string	no	Mounted QStar Integral Volume share type, e.g. CIFS, NFS. Default is CIFS.
qstar-use-https	bool	no	Use https for QStar service. Default is false.
qstar-username	string	no	QStar service username.
s3-compatible-service-url	string	no	Service URL for S3 compatible cloud.
s3-compatible-signature-version	int	no	Signature version fo S3 compatible cloud. Default is 4.
s3-compatible-use-https	bool	no	Use https for S3 compatible cloud. Default is true.
view-box-name	string	no	Name of the Storage Domain.

view

To run any View related operations.

Operation	Short Name	Description
activate-saved-aliases		Activates the saved aliases of a cloned view.
add-subnet-to-whitelist	add-st-wl	To add a subnet to a view's subnet whitelist
allow-mount-on-windows		Allow mount path of the View to be mounted on windows machines.
cancel-failover		This command will cancel an in-progress failover on the specified view.
clone		To clone a View from an existing View.
clone-directory		To clone (or recursively copy) the contents of a directory from a source path to a destination path in the same View or a different View.
create	add	To create a new View.
create-alias		To create a View Alias.

Operation	Short Name	Description
delete	rm	To delete an existing View.
delete-alias		To delete a View Alias.
failover		This command will initiate a failover on the specified view.
get-file-lock-status		Fetches the lock details of a file in a view.
get-subnet-whitelist	ls-st-wl	To get a view's subnet whitelist
list	ls	To list View details.
list-failover		This command will list all failovers on the specified view.
list-nfs-connections	ls-nfs-con	To list the active NFS connections for views.
list-nlm-locks	ls-nl	To list the acquired NLM locks.
list-smb-connections	ls-smb-con	To list the active SMB connections for views.
list-smb-file-opens	ls-sfo	To list the active SMB file opens for views.
lock		To lock a View. A locked view cannot be modified or deleted.
lock-file		Locks a file in a view.
overwrite		To overwrite a target view with contents of source view.
remove-subnet-from-whitelist	rm-st-wl	To remove a subnet from a view's subnet whitelist
rename	ren	To rename a view.
unlock		To unlock a View.
update	edit	To update an existing View.

view activate-saved-aliases

Activates the saved aliases of a cloned view.

Parameter	Type	Required	Description
name	string	yes	Name of the View.

view add-subnet-to-whitelist

To add a subnet to a view's subnet whitelist

Parameter	Type	Required	Description
desc	string	no	Description.
name	string	yes	Name of the View.
subnet-ip	string	yes	Subnet IP for the Cluster.
subnet-mask	string	yes	Subnet mask for the Cluster. Should be either in CIDR prefix or in dotted decimal.
subnet-nfs-permission	string	no	Clients from this subnet can mount using NFS protocol. This can be kDisabled, kReadOnly, kReadWrite.
subnet-nfs-squash	string	no	Clients from this subnet can mount as root or all or none on NFS. This can be kNone, kRootSquash, kAllSquash.
subnet-s3-permission	string	no	Clients from this subnet can access using S3 protocol. This can be kDisabled, kReadOnly, kReadWrite.
subnet-smb-permission	string	no	Clients from this subnet can mount using SMB protocol. This can be kDisabled, kReadOnly, kReadWrite.

view allow-mount-on-windows

Allow mount path of the View to be mounted on windows machines.

Parameter	Type	Required	Description
enable	bool	yes	To allow/disallow the mount path of the View to be mounted on windows machines using NFS protocol.
name	string	yes	Name of the View.

view cancel-failover

This command will cancel an in-progress failover on the specified view.

Parameter	Type	Required	Description
id	int	yes	Specifies the view id.

view clone

To clone a View from an existing View.

Parameter	Type	Required	Description
datalock-expiry-usecs	int	no	Datalock duration in microseconds of the cloned view, the cloned view can only be deleted by a Data Security Officer until the datalock expires.
disable-inline-dedup	bool	no	Disable the inline dedup for the View. You can only turn off inline dedup for the View, if it is on for the containing Storage Domain.
name	string	yes	Name of the cloned View.
qos-policy	string	no	QoS Policy of the View. For example, it can be 'Backup Target High', 'Backup Target Low', 'Backup Target SSD', 'TestAndDev High', 'TestAndDev Low', 'Filer High', 'Filer Low', 'Backup Target Auto' etc.
source-view-name	string	yes	Name of the source View to be cloned.
view-description	string	no	Description of the View.

view clone-directory

To clone (or recursively copy) the contents of a directory from a source path to a destination path in the same View or a different View.

Parameter	Type	Required	Description
destination-directory-name	string	yes	The name of the new directory that is created in the specified destination directory when the contents of the source directory is cloned. The new directory must not exist.
destination-directory-path	string	yes	The directory path of the destination directory where the cloned new directory will be created. Path must include the View name e.g. /View1/Path.
source-directory-path	string	yes	The full path of the source directory in the View to clone (or recursively copy) from. Path must include the View name e.g. /View1/Path.

view create

To create a new View.

Parameter	Type	Required	Description
allow-case-insensitive	bool	no	Whether to allow case insensitive names for files and directories.
disable-inline-dedup	bool	no	Disable the inline dedup for the View. You can only turn off inline dedup for the View, if it is on for the containing Storage Domain.
enable-antivirus	bool	no	Enable antivirus scan on this view. If enabled, at least one of "enable-scan-on-access" or "enable-scan-on-close" should also be enabled.
enable-fast-durable-handle	bool	no	Whether to enable fast durable handle. If set to true, view open handle will be kept in memory, which results in a higher performance. But the handles cannot be recovered if the node or service crashes.
enable-live-indexing	bool	no	Enable live indexing feature for the view.
enable-minion	bool	no	Enable minion feature for the view.

Parameter	Type	Required	Description
enable-nfs-kerberos-authentication	bool	no	Enable kerberos authentication in NFS.
enable-nfs-kerberos-integrity	bool	no	Enable kerberos integrity in NFS.
enable-nfs-kerberos-privacy	bool	no	Enable kerberos privacy in NFS.
enable-nfs-unix-authentication	bool	no	Enable unix authentication in NFS.
enable-nfs-view-discovery	bool	no	Enable View disoverability in NFS.
enable-offline-caching	bool	no	Enable offline file caching of this view.
enable-pinning	bool	no	Enable pinning on this view. If enabled, the chunk files accessed from this view will be pinned to SSD tier.
enable-s3-folder-support	bool	no	Whether to allow s3 folder support feature on the View.
enable-scan-on-access	bool	no	Enable antivirus scan when view is accessed.
enable-scan-on-close	bool	no	Enable antivirus scan when view is closed.
enable-smb-encryption	bool	no	Enable SMB encryption on this view alias.
enable-smb-oplock	bool	no	Whether to enable SMB oplock. If enabled, a client will be able to cache data locally.
enable-smb-view-discovery	bool	no	Enable View disoverability in SMB.

Parameter	Type	Required	Description
file-autolock-after-idle-duration	string	no	Specify the duration to lock a file that has not been accessed or modified (ie. has been idle) for a certain duration of time. For example: 2Y3M1D for 2 years, 3 months and 1 day. Do not set if it is required to disable auto lock.
file-extension-filter-blacklist	string	no	Creates a blacklist to disallow certain file extensions.
file-extension-filter-whitelist	string	no	Creates a whitelist to only allow certain file extensions.
file-lock-default-retention-duration	string	no	Specify a default retention duration for files in this view, if file lock is enabled for this view. Set to forever if the required default retention period is forever. For example: 2Y3M1D for 2 years, 3 months and 1 day.
file-lock-expiry-date	datetime	no	Specify a definite timestamp in "YYYY-MM-DDTHH:MM" format for retaining the file.
file-lock-max-retention-duration	string	no	Specify a maximum duration for which any file in this view can be retained for. For example: 2Y3M1D for 2 years, 3 months and 1 day. Set to forever if the required retention duration is forever.
file-lock-min-retention-duration	string	no	Specify a minimum retention duration after a file gets locked. The file cannot be modified or deleted during this timeframe. For example: 2Y3M1D for 2 years, 3 months and 1 day. Set forever if the required retention duration is forever.
file-lock-mode	string	no	Specify the mode of file lock. Possible values are: 'compliance' (default value) and 'enterprise'.
file-lock-protocol	string	no	Sets the file locking protocol. Possible values are: 'read-only', 'set-atime'.
filer-audit-enabled	bool	no	Whether to enable Filer Audit Logs.
name	string	yes	Name of the View.

Parameter	Type	Required	Description
pin-time-secs	int	no	Pin the chunk files accessed from this view to SSD tier for this long. By default, we pin the chunk files indefinitely.
project-domain	string	no	The Keystone project domain.
project-name	string	no	The Keystone project name.
protocol	list of strings	no	Array of protocols for the view. Valid values are NFS, NFS4, SMB, S3 and Swift and are specified in a comma separated string. Default if not mentioned during view creation will be taken as NFS, SMB and S3(Read-only)
protocol-mode	list of strings	no	Array of mode for protocols in the view. Valid values are ReadOnly (or 'r') and ReadWrite (or 'w') and are specified in a comma separated string. Number of items in the string should be equal to items in 'protocol' options or should be empty. Empty denotes that all protocols will be taken as ReadWrite mode.
qos-policy	string	no	QoS Policy of the View. For example, it can be 'Backup Target High', 'Backup Target Low', 'Backup Target SSD', 'TestAndDev High', 'TestAndDev Low', 'Filer High', 'Filer Low', 'Backup Target Auto' etc.
quota	int	no	Maximum allowed quota in GiB. Set -1 to remove the value.
quota-alert-limit	int	no	Threshold in GiB, over which an alert will be sent. Set -1 to remove the value.
quota-inherit	bool	no	Whether to inherit quota from Storage Domain. The default value of this parameter is set to true.
s3-key-mapping-type	string	no	Specifies the type of S3 key mapping config. Possible values are 'Random','Short','Long','Hierarchical' and 'ObjectId'.
security-style	string	no	Set security style in the View. This can be Native, Unified, Ntfs.
smb-access-based-enumeration	bool	no	Enable Access Based Enumeration for SMB

Parameter	Type	Required	Description
storage-domain-id	int	no	ID of a Storage Domain.
storage-domain-name	string	no	Name of the Storage Domain.
tenant-id	string	no	To specify id of the tenant owning a resource.
user-domain	string	no	The Keystone user domain.
user-name	string	no	The Keystone user name.
view-category	string	no	Category of the View. Valid values are 'BackupTarget', 'FileServices' and 'ObjectServices'. If category is not defined by user then by default 'FileServices' is selected.
view-description	string	no	Description of the View.

view create-alias

To create a View Alias.

Parameter	Type	Required	Description
alias	string	yes	Alias for the view. It allows a path inside the View to be mounted using the Alias name.
view-name	string	yes	Name of the View.
view-path	string	yes	Full path of the directory inside a view.

view delete

To delete an existing View.

Parameter	Type	Required	Description
name	string	yes	Name of the View.

view delete-alias

To delete a View Alias.

Parameter	Type	Required	Description
alias	string	yes	Alias for the view. It allows a path inside the View to be mounted using the Alias name.

view failover

This command will initiate a failover on the specified view.

Parameter	Type	Required	Description
id	int	yes	Specifies the view id.
type	string	yes	Specifies the view failover type. Type can be one of 'Planned', 'Unplanned' and 'Finalize'.

view get-file-lock-status

Fetches the lock details of a file in a view.

Parameter	Type	Required	Description
view-name	string	yes	Name of the View.
view-path	string	yes	Full path of the directory inside a view.

view get-subnet-whitelist

To get a view's subnet whitelist

Parameter	Type	Required	Description
name	string	yes	Name of the View.

view list

To list View details.

Parameter	Type	Required	Description
ids	list of ints	no	Comma separated list of View ids.

Parameter	Type	Required	Description
names	list of strings	no	Comma separated list of View names.
storage-domain-ids	list of ints	no	Comma separated list of Storage Domain IDs.
storage-domain-names	list of strings	no	Comma separated names of the Storage Domains.

view list-failover

This command will list all failovers on the specified view.

Parameter	Type	Required	Description
id	int	yes	Specifies the view id.

view list-nfs-connections

To list the active NFS connections for views.

Parameter	Type	Required	Description
names	list of strings	no	Comma separated list of View names.

view list-nlm-locks

To list the acquired NLM locks.

Parameter	Type	Required	Description
cookie	string	no	Opaque string returned in a previous command output to get the next set of entries in the output.
file-path	string	no	File path relative to the NFS mount path.
name	string	no	Name of the View.
page-count	int	no	Limit the number of entries in the command output.
verbose	bool	no	If set to true, more information is displayed.

view list-smb-connections

To list the active SMB connections for views.

Parameter	Type	Required	Description
names	list of strings	no	Comma separated list of View names.

view list-smb-file-opens

To list the active SMB file opens for views.

Parameter	Type	Required	Description
cookie	string	no	Opaque string returned in a previous command output to get the next set of entries in the output.
file-path	string	no	File path of an SMB file relative to the SMB mount path.
name	string	no	Name of the View.
page-count	int	no	Limit the number of entries in the command output.

view lock

To lock a View. A locked view cannot be modified or deleted.

Parameter	Type	Required	Description
name	string	yes	Specifies the name of the view to lock.

view lock-file

Locks a file in a view.

Parameter	Type	Required	Description
file-lock-expiry-date	datetime	no	Specify a definite timestamp in "YYYY-MM-DDTHH:MM" format for retaining the file.
view-name	string	yes	Name of the View.
view-path	string	yes	Full path of the directory inside a view.

view overwrite

To overwrite a target view with contents of source view.

Parameter	Type	Required	Description
source-view-name	string	yes	Name of the Source View to copy contents from.
target-view-name	string	yes	Name of the Target View to be overwritten.

view remove-subnet-from-whitelist

To remove a subnet from a view's subnet whitelist

Parameter	Type	Required	Description
clear-all	bool	no	Set this field to true to clear all subnet white list.
name	string	yes	Name of the View.
subnet-ip	string	no	Subnet IP for the Cluster.

view rename

To rename a view.

Parameter	Type	Required	Description
name	string	yes	Name of the View.
new-name	string	yes	New Name of the View.

view unlock

To unlock a View.

Parameter	Type	Required	Description
name	string	yes	Specifies the name of the view to unlock.

view update

To update an existing View.

Parameter	Type	Required	Description
clear-auto-lock	bool	no	Clears the Auto Lock duration from the file lock config of the view.
clear-expiry-timestamp	bool	no	Clears the expiry timestamp from the file lock config of the view.
clear-max-retention-time	bool	no	Clears the max retention duration from the file lock config of the view.
clear-min-retention-time	bool	no	Clears the min retention duration from the file lock config of the view.
disable-inline-dedup	bool	no	Disable the inline dedup for the View. You can only turn off inline dedup for the View, if it is on for the containing Storage Domain.
enable-antivirus	bool	no	Enable antivirus scan on this view. If enabled, at least one of "enable-scan-on-access" or "enable-scan-on-close" should also be enabled.
enable-fast-durable-handle	bool	no	Whether to enable fast durable handle. If set to true, view open handle will be kept in memory, which results in a higher performance. But the handles cannot be recovered if the node or service crashes.
enable-minion	bool	no	Enable minion feature for the view.
enable-nfs-kerberos-authentication	bool	no	Enable kerberos authentication in NFS.
enable-nfs-kerberos-integrity	bool	no	Enable kerberos integrity in NFS.
enable-nfs-kerberos-privacy	bool	no	Enable kerberos privacy in NFS.
enable-nfs-unix-authentication	bool	no	Enable unix authentication in NFS.

Parameter	Type	Required	Description
enable-nfs-view-discovery	bool	no	Enable View disoverability in NFS.
enable-offline-caching	bool	no	Enable offline file caching of this view.
enable-pinning	bool	no	Enable pinning on this view. If enabled, the chunk files accessed from this view will be pinned to SSD tier.
enable-scan-on-access	bool	no	Enable antivirus scan when view is accessed.
enable-scan-on-close	bool	no	Enable antivirus scan when view is closed.
enable-smb-oplock	bool	no	Whether to enable SMB oplock. If enabled, a client will be able to cache data locally.
enable-smb-view-discovery	bool	no	Enable View disoverability in SMB.
file-autolock-after-idle-duration	string	no	Specify the duration to lock a file that has not been accessed or modified (ie. has been idle) for a certain duration of time. For example: 2Y3M1D for 2 years, 3 months and 1 day. Do not set if it is required to disable auto lock.
file-extension-filter-blacklist	string	no	Creates a blacklist to disallow certain file extensions.
file-extension-filter-whitelist	string	no	Creates a whitelist to only allow certain file extensions.
file-lock-default-retention-duration	string	no	Specify a default retention duration for files in this view, if file lock is enabled for this view. Set to forever if the required default retention period is forever. For example: 2Y3M1D for 2 years, 3 months and 1 day.
file-lock-expiry-date	datetime	no	Specify a definite timestamp in "YYYY-MM-DDTHH:MM" format for retaining the file.

Parameter	Type	Required	Description
file-lock-max-retention-duration	string	no	Specify a maximum duration for which any file in this view can be retained for. For example: 2Y3M1D for 2 years, 3 months and 1 day. Set to forever if the required retention duration is forever.
file-lock-min-retention-duration	string	no	Specify a minimum retention duration after a file gets locked. The file cannot be modified or deleted during this timeframe. For example: 2Y3M1D for 2 years, 3 months and 1 day. Set forever if the required retention duration is forever.
file-lock-mode	string	no	Specify the mode of file lock. Possible values are: 'compliance' (default value) and 'enterprise'.
file-lock-protocol	string	no	Sets the file locking protocol. Possible values are: 'read-only', 'set-atime'.
filer-audit-enabled	bool	no	Whether to enable Filer Audit Logs.
name	string	yes	Name of the View.
pin-time-secs	int	no	Pin the chunk files accessed from this view to SSD tier for this long. By default, we pin the chunk files indefinitely.
project-domain	string	no	The Keystone project domain.
project-name	string	no	The Keystone project name.
protocol	list of strings	no	Array of protocols for the view. Valid values are NFS, NFS4, SMB, S3 and Swift and are specified in a comma separated string. Default if not mentioned during view creation will be taken as NFS, SMB and S3(Read-only)
protocol-mode	list of strings	no	Array of mode for protocols in the view. Valid values are ReadOnly (or 'r') and ReadWrite (or 'w') and are specified in a comma separated string. Number of items in the string should be equal to items in 'protocol' options or should be empty. Empty denotes that all protocols will be taken as ReadWrite mode.

Parameter	Type	Required	Description
qos-policy	string	no	QoS Policy of the View. For example, it can be 'Backup Target High', 'Backup Target Low', 'Backup Target SSD', 'TestAndDev High', 'TestAndDev Low', 'Filer High', 'Filer Low', 'Backup Target Auto' etc.
quota	int	no	Maximum allowed quota in GiB. Set -1 to remove the value.
quota-alert-limit	int	no	Threshold in GiB, over which an alert will be sent. Set -1 to remove the value.
quota-inherit	bool	no	Whether to inherit quota from Storage Domain. The default value of this parameter is set to true.
security-style	string	no	Set security style in the View. This can be Native, Unified, Ntfs.
smb-access-based-enumeration	bool	no	Enable Access Based Enumeration for SMB
tenant-id	string	no	To specify id of the tenant owning a resource.
use-file-extension-filter	bool	no	Enables or Disables the existing filter on file extensions for the view.
view-category	string	no	Category of the View. Valid values are 'BackupTarget', 'FileServices' and 'ObjectServices'. If category is not defined by user then by default 'FileServices' is selected.
view-description	string	no	Description of the View.

vlan

To run any Vlan related operations.

Operation	Short Name	Description
add-dns-zone		To add a dns zone name and vips mappingfor a VLAN interface in the Cluster.
create	add	To add a VLAN for the Cluster.

Operation	Short Name	Description
delete	rm	To remove a VLAN from the Cluster.
remove-dns-zones		To remove dns zone for a VLAN interfacein the Cluster.
update	edit	To edit a VLAN for the Cluster.
update-dns-zone		To update dns zone for a VLAN interfacein the Cluster.

vlan add-dns-zone

To add a dns zone name and vips mappingfor a VLAN interface in the Cluster.

Parameter	Type	Required	Description
dns-zone-name	string	yes	To remove internal dns entries specified by given dns host name.
dns-zone-vips	list of strings	yes	To update internal dns vips.
vlan-interface-group	string	yes	Name of a network interface group.

vlan create

To add a VLAN for the Cluster.

Parameter	Type	Required	Description
apps-ips	list of strings	no	Comma separated list of VLAN IP addresses for Apps.
ecmp-enabled	bool	no	When set, ECMP will be enabled in the vlan.
family	int	no	Param to specify IP address family, 1=IpV4(default) or 2=IpV6.
hostname	string	no	Hostname for the VLAN IPs.
id	int	yes	ID of the VLAN.
interface-group	string	no	Name of a network interface group.

Parameter	Type	Required	Description
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-ip	string	no	Subnet IP of the VLAN.
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.
subnet-mask-ip4	string	no	Subnet mask for IPv4 type subnet of the Cluster.
vlan-ips	list of strings	no	Comma separated list of VLAN IP addresses.
vlan-name	string	no	VLAN Name for vlan-id created

For example:

```
iris_cli vlan create id=100 interface-group=intf_group1 vlan-ips-
s=150.0.0.1,150.0.0.2,150.0.0.3 subnet-gateway=150.0.0.100 subnet-mask-bits-
s=24 subnet-ip=150.0.0.0
```

vlan delete

To remove a VLAN from the Cluster.

Parameter	Type	Required	Description
vlan-interface-group	string	no	Name of a network interface group.

For example:

```
iris_cli vlan delete vlan-interface-group=intf_group1.100
```

vlan remove-dns-zones

To remove dns zone for a VLAN interfacein the Cluster.

Parameter	Type	Required	Description
dns-zone-names	list of strings	yes	Name of a network interface group.
vlan-interface-group	string	yes	Name of a network interface group.

vlan update

To edit a VLAN for the Cluster.

Parameter	Type	Required	Description
add-apps-ips	list of strings	no	Comma separated list of IPs to add to the VLAN for Athena Apps.
add-ips	list of strings	no	Comma separated list of IPs to add to the VLAN.
family	int	no	Param to specify IP address family, 1=IpV4(default) or 2=IpV6.
hostname	string	no	Hostname for the VLAN IPs.
pool-name	string	no	To group the follower IPs to a pool.
remove-hostname	bool	no	Whether to remove the configured Hostname for the VLAN IPs.
remove-ips	list of strings	no	Comma separated list of VLAN IP addresses to remove from the VLAN.
subnet-gateway	string	no	Subnet gateway for the Cluster.
subnet-ip	string	yes	Subnet IP of the VLAN.
subnet-mask-bits	int	no	Subnet mask bits for the Cluster.
subnet-mask-ip4	string	no	Subnet mask for IPv4 type subnet of the Cluster.
vlan-interface-group	string	yes	Name of a network interface group.
vlan-name	string	no	VLAN Name for vlan-id created

vlan update-dns-zone

To update dns zone for a VLAN interfacein the Cluster.

Parameter	Type	Required	Description
add-resolved-vips	list of strings	no	To add resolved vips to the zone.

Parameter	Type	Required	Description
dns-zone-name	string	yes	To remove internal dns entries specified by given dns host name.
dns-zone-vips	list of strings	yes	To update internal dns vips.
remove-resolved-vips	list of strings	no	To remove resolved vips from the zone.
vlan-interface-group	string	yes	Name of a network interface group.

Cluster Health Check

hc_cli is a Cohesity support-automation tool which checks for commonly seen issues on a Cohesity cluster. The output of the verbose command includes links to a knowledge base article which describes how to resolve the error. After using hc_cli to perform the initial checks, you need to open a case with Cohesity Support to review your cluster's health.

Run hc_cli

To run hc_cli, run the `hc_cli` run command.

To view results, run the `hc_cli show --failed` command.

To view the built-in help, run the `hc_cli` command without any arguments or append `-h` at the end of the command:

```
$ hc_cli
usage: hc_cli.py [-h]
{debug,ls-tests,show,run,version,generate-doc,status} ...
hc_cli.py: error: too few arguments
```

See the following table for a list of hc_cli commands and description.

Operation	Description

debug	• Runs utilities which are more resource-intensive, collect system state, or are informational.
	• Available utilities are:
	  • performance_traces.py
	  • For the Backup Manager service:
	  • get_magneto_vm_protection_flags.py
	  • get_job_hierarchy.py
	  • For the Cluster Manager service:
	  • top_cpu.py
	  • services_memory_check.py
	  • network_errors.py
	  • For the Indexing service:
	  • yoda_persistent_queue.py

ls-tests	• Lists available lists.

show	• Prints the most recent results.
	• Available flags include:
	  • **-h/--help**: Prints help output.
	  • **-C/--current-node**: Prints output from only the current node.
	  • **-F/--failed**: Prints only failed checks.
	  • **-v/--verbose**: Prints verbose output.

Operation	Description

run	• Initiates a run.
	• Available parameters are:
	  • **-h/--help**: Prints help output.
	  • **--test-ids**: Runs only the specified tests.
	  For example, `hc_cli run --test-ids=10003,10010` runs only tests 100003 and 10010.
	  • **-u/--username & -p/--password**: Defines a non-standard UI username and password. If not defined, credentials default to username admin and password admin.
	  Example: `hc_cli -u mario -p mushroomkingdom`.
	  • **-v/--verbose**: Prints verbose output. If not specified, prints only pass or fail status.
	  • **--type**: Runs all checks of a certain type. Options are:
	  • **hardware**: hardware-level checks.
	  • **software**: Cohesity software-specific checks.
	  • **network**: network checks and configuration validation checks.
	  • **-C/--current-node**: Runs hc_cli only on the current node.
	  • **--max-timeout**: Configures a timeout value. Generally recommended only for larger clusters where hc_cli might run for an extended period. A value of 99999 covers larger clusters.

version	• Displays the build time of the installed hc_cli package.
generate-doc	• Prints the built-in documentation.

status	• Prints run status.
	• Status values are:
	  • **Runs Ready** if no run is in progress.
	  • **Running** if a run is in progress.

clean	• Removes hc_cli and support-toolbox from the ~/bin directory.

Available Checks

To fetch the available checks from the cluster, run the `generate-doc` command.

Sample available checks:

Test ID	Test Name	Type	Category
10001	SAS Controller / smartctl Error Check	kHardware	kEssential, kSmartFiles, kPreUpgrade, kPostUpgrade
10002	Iostat Disk Utilization Percentage Check	kHardware	kEssential, kSmartFiles
10003	Custom Binary Check	kSoftware	kPreUpgrade
10004	Partition Size Utilization Check	kSoftware	kEssential
10006	Helios Connectivity Check	kSoftware	kEssential

Support Toolbox

The support toolbox and its cli interface, hc_cli, are a set of tools built to enable self service of a cluster by identifying issues and providing documentation with the goal of resolution, or if escalation is needed, a more targeted support case.

The support toolbox and its constituent components are all designed to assist in identifying known issues, validating cluster health, and to certify a cluster is prepared for an upgrade.

As Cohesity keeps this tool in constant development and updates are constantly being built, refer to the following Knowledge Base article to get the latest data on the toolbox, execution, and how to use new features:

- Support Toolbox overview
- Updating the support toolbox to the latest version
- Running hc_cli health checks
- Preupgrade checks (recommended to review prior to an upgrade)
- List of checks with links to their KBs
- Showing verbose output post check

Using the Secure Shell

The Secure Shell restricts SSH access to the host operating system shell, further tightening access to the Cohesity cluster while still providing users with the ability to execute CLI commands, run diagnostic tools, and access logs. In previous releases, support users could log on to Cohesity nodes using SSH and access the host operating system shell. Access to the host operating system is now restricted through SSH.

The Secure Shell can be accessed using only the support user account. To enable the support user account, see Manage Support User.

Using the Secure Shell, you can perform the following:

- Run iris_cli commands
- View the logs (read-only)
- Basic connectivity checks

> **Note:**
> The basic connectivity checks can be performed using the following utilities:
>
> - ping and curl will be available from 6.8.1 and later versions.
> - dig/host, nping, tcptraceroute, and elinks will be available from 6.8.1_u7 and later versions.

Considerations

After upgrading to the 6.8 or later releases, the secure shell might have the following impact on your existing Cohesity deployments:

- If you run custom scripts using ssh on your Cohesity cluster, the scripts might fail. In this case, Cohesity recommends the following:
  - Verify if there is an alternate method to use Cohesity CLI commands or REST API and update your scripts accordingly.
  - Verify if a corresponding Cohesity CLI command is available in the supported list of CLI commands; if so, use the supported CLI command. For more information, see Using the Secure Shell. If the CLI command is not available in the supported list of commands, contact Cohesity Support to enable the CLI command.
- The private binary or tools running on the Cohesity nodes might fail. Contact Cohesity Support for options to install private binaries or tools.
- Sudo access is disabled by default. For support channel access, you need to enable sudo access for the support user to the host shell using the following command:

```
iris_cli user linux-sudo-access enable=true duration-in-hours=xxx
```

> **Note:** You can enter a duration within the range of 1 to 720 hours.

Secure Shell Prompt

To identify whether you are running the secure shell, once you log into the shell using the support user account, the shell prompt displays the term "***restricted***," this suggests that

you have logged into the Secure Shell of a Cohesity node.

For example:

```
$ ssh support@10.x.x.16
support@10.1.4.16's password:
Last login: Fri Dec 17 03:32:32 UTC 2021 from 10.2.209.99 on pts/1
Last login: Fri Dec 17 03:46:25 2021 from 10.2.209.99
***************** Welcome to Cohesity *****************

WARNING: Unauthorized access to this system is forbidden.
By accessing this system, you agree that your actions
may be monitored if unauthorized usage is suspected.

*********************************************************
[support@restricted-haswell2-bqkp91500082-node-2 ~]\> iris_cli
 Username: admin
 'admin' Password:
```

Enable Host Shell Access

You can access the host shell on port 2222 using the ***support*** user account.

To enable host shell access:

1. Log on to the node's secured shell using the ***support*** user. Use the following command:

   ```
   $ ssh support@<node_IP>
   ```

   Where `node_IP` is the IP address of the node in your Cohesity cluster.

2. Log on to the `iris_cli` and generate a support token using the following command:

   ```
   $ iris_cli support host-access-token
   ```

   > **Note:** By default, the support token is only valid for an hour. You can extend the support token validity for up to 12 hours using the following command:
   > `iris_cli support host-access-token validity-in-hours=12`

3. Obtain the Cohesity cluster information using the following command:

   ```
   $ iris_cli cluster info
   ```

4. Share the following details with Cohesity Support to generate an access code:
   - Host Access Token generated in step 2
   - Cluster ID obtained in step 3

- Cohesity Software Version obtained in step 3
- The number of hours you require the host shell access

5. After you receive the access code from Cohesity Support, enable host shell access using the following command in `iris_cli`:

```
$  iris_cli support host-access ssh-enable=true access-code=<access_
code>
```

Where `access_code` is the access code you received from Cohesity Support.

> **Important:** Ensure that you enable host shell access before the support access token you generated in **Step 2** expires.

6. Access the host shell of the node on port 2222 using the ***support*** user. Use the following command:

```
$ ssh support@<node_IP> -p 2222
```

Cohesity REST API

Cohesity provides a REST API that supports programmatically interfacing with a Cohesity cluster, for example:

- Send a `GET /public/protectionJobs` REST API request that returns all the Protection Groups found on a Cohesity cluster.
- Send a `GET /public/protectionSources/virtualMachines` REST API request that returns all the VMs found in a protected Source (vCenter Server) that are registered on a Cohesity cluster.

Complete REST API documentation is available on the Cohesity Developer Portal.

Cohesity provides ready-to-use integrations with the automation and orchestration tools of your choice to streamline your operations. Under the hood, all integrations are powered by the Cohesity REST API. For more information, see API Integrations.

View and Run the API Directly on a Cluster

Cohesity also provides a REST API interface directly on the Cohesity cluster which provides documentation and the ability to run operation requests directly against the Cohesity cluster and view the response.

1. Log in to the Cohesity Dashboard from a browser by entering a Node IP Address or VIP IP address of a Cohesity cluster into the browser address bar.
2. Click the ? icon in the upper right corner and select **REST API**.

A new browser tab opens and the Cohesity REST API operations are displayed as nested under group tags such as **AccessTokens** and **Active Directory**.

> **Note:** By default, **AccessTokens** are used. **Sessions** will be used if session management is enabled.

3. To view the documentation on the response:
   a. Click **Model** next to **Example Value**.

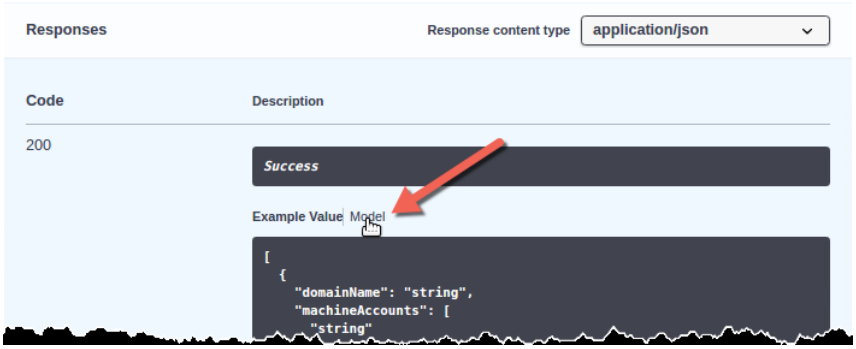

   b. Keep clicking to expand the documentation hierarchy.

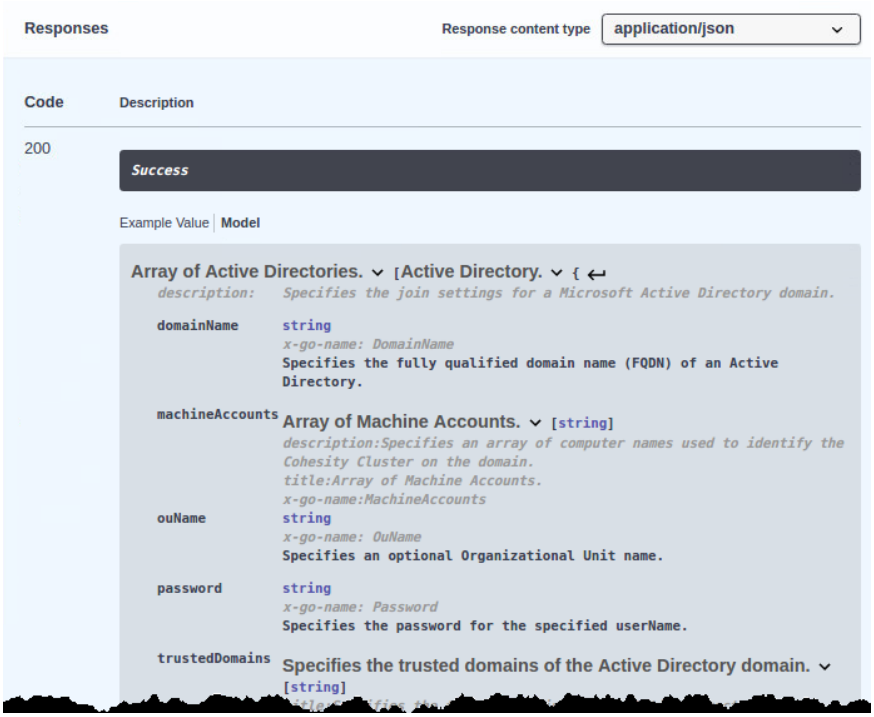

4. To run an operation:
   a. Click on the operation to expand it.
   b. Click **Try it out**.

c. Enter any required parameters.

d. Click **Execute**.

The operation is run against the Cohesity cluster you are logged in to and the results are displayed under the Response body.

Specifying Access Credentials

When you run Cohesity REST API operations from the browser, the Cohesity cluster uses the Login credentials you specified to log in to the Cohesity Dashboard.

You can run the Cohesity REST API operations programmatically in two ways:

1. Access tokens
2. Sessions

Access Tokens

> **Note:** By default, AccessTokens are used.

From your REST client program, you must make a `POST /public/accessTokens` request to supply the valid Login credentials before calling other Cohesity REST API operations. This POST request returns an access token and type in the response that is generated by the Cohesity cluster. Subsequent requests to other Cohesity REST API operations must specify the returned access token and type by setting `Authorization` in the http header in the following format:

```
Authorization: <token_type> <access_token>
```

The generated token is valid for 24 hours. If a request is made with an expired token, the `Token expired` error message is returned. Add code to your REST client to check for this error and request another access token before reissuing the request.

Sessions

> **Note:** Sessions will be used if session management is enabled.

From your REST client program, you must make a `POST /public/v2/users/sessions` request to supply the valid Login credentials before calling other Cohesity REST API operations. This POST request returns the session id in the response that is generated by the Cohesity cluster.

If session management is enabled, the subsequent requests to other Cohesity REST API operations must specify the returned session id by setting the session-id header in the http header in the following format:

```
session-id: "oxxx0yyyyy81Rzzzz//aaaaaO7qbbbbbAcccccc="
```

The session id is valid till the Inactivity Timeout (1 hour) or Absolute Timeout (24 hours).

If a request is made with an expired session id, the expired error message is returned. Add code to your REST client to check for this error and request another session id before reissuing the request.

Generate an API Key

To generate an API key:

1. Navigate to **Settings** > **Access Management** and click **API Keys**.

   > **Note:** By default, the API Key is disabled. Contact Cohesity Support to enable the feature.

2. Click **Add API Key**.
3. Select the user for which the API key should be generated.
4. Enter a **Name** for the API key.
5. Click **Add** to generate the API key.
6. Copy the API key to use with the application.

Chapter 36: Supplemental Information

This chapter includes the following topics:

- Open Source in Cohesity Products 5147
- Explore Our Videos 5148
- Helios Documentation 5149
- Cohesity Hardware Data Sheets and Compliance Guide 5149
- Hardware Replacement 5149
- Technical Guides 5149
- Cohesity Support 5149
- Glossary 5150
- Documentation Archive 5154
- Documentation Feedback 5155

Open Source in Cohesity Products

The tables below include information regarding open source software (OSS) distributed in or with the version of the Product (as such term is defined in the Cohesity EULA).

To request a similar list with respect to other versions, please send a request to: oss@cohesity.com.

Cohesity Product: Software

Cohesity Product	Version	Link
Cohesity Software	7.4	OSS List

Cohesity Product	Version	Link
Cohesity Software	7.3.2	OSS List
Cohesity Software	7.3.1	OSS List
Cohesity Software	7.3	OSS List
Cohesity Software	7.2.2_u3	OSS List
Cohesity Software	7.2.2_u2	OSS List
Cohesity Software	7.2.2_u1	OSS List
Cohesity Software	7.2.2	OSS List
Cohesity Software	7.2.1	OSS List
Cohesity Software	7.2	OSS List
Cohesity Software	7.1.2_u4	OSS List
Cohesity Software	7.1.2	OSS List
Cohesity Software	7.1.1	OSS List
Cohesity Software	7.1	OSS List
Cohesity Software	7.0.1	OSS List
Cohesity Software	7.0	OSS List
Cohesity Software	6.8.2	OSS List
Cohesity Software	6.8.1_u2	OSS List
Cohesity Software	6.8.1	OSS List
Cohesity Software	6.6	OSS List

Explore Our Videos

You can learn about Cohesity products and solutions by watching videos on the Cohesity How-to Videos page.

Helios Documentation

Helios is Cohesity's SaaS-based management platform that provides a single dashboard and global management of all your Cohesity Clusters. For more information, see Helios Documentation.

Cohesity Hardware Data Sheets and Compliance Guide

- Cohesity Model C2xx5 (C2305, C2505, C2515, and C2605)
- Cohesity Model C3000
- Cohesity Model C4000
- Cohesity Model C5000
- Cohesity Model C6000
- Cohesity Model CX8000
- Safety and Compliance Guide
- Cohesity Media Sanitization Brief

Hardware Replacement

Hardware replacement procedures (PDFs) are available at: Hardware Universe.

Technical Guides

Additional Solution Guides, White Papers, Deployment Guides, Reference Architectures and Best Practices are available at Technical Guides.

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to Cohesity Support, to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the Cohesity Support Portal to create a new case.
- Click the (**?**) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the Cohesity Knowledge Base.
- Log in to the Cohesity Support Portal to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

   > **Note:** Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Glossary

A

Agent

A lightweight software that is installed on Operating systems like Windows, Linux, AIX, Solaris, HP-UX etc running on a Physical Server or a Virtual machine enabling communication between OS and Cohesity to protect data.

B

Backup
: A recoverable snapshot of a registered source, such as a physical or virtual server, database or storage array, created by a Protection Group, used to protect against data loss or corruption.

C

Clone
: The ability to clone objects from snapshots created by a Protection Group. The cloned objects are copied to a new location, rehydrated and restored to a running state.

Cloud Tier
: The process of moving rarely used and inactive data to the cloud when HDD usage on the Cohesity cluster exceeds a certain threshold.

CloudArchive
: The process of copying snapshots created by Protection Groups from the Cohesity cluster to a registered external target.

CloudArchive Direct
: The process of archiving data directly to an external target without retaining a copy of the data on the cluster. Only the metadata and indexes, required for quick search and recovery, are stored on the cluster.

CloudRetrieve
: The process of retrieving an archived Protection Group and its run details from an external target to a different cluster. Used for geo-redundancy and disaster recovery.

Cluster
: An instance of Cohesity Data Cloud (Self-Managed).

D

Data In
: Data sent from a protected source to Cohesity.

Data Written
: Data to be written post-deduplication.

Deduplication
: The process of eliminating redundant copies of data to reduce storage usage on the cluster.

  The Inline Deduplication process deduplicates data as it is written to the cluster and the Post-process deduplication deduplicates data after it is written to the cluster.

E

External Target
: Any storage location outside the Cohesity cluster where data is stored. They can be for cloudarchive, cloudarchive direct, or cloud tier.

F

Firewall Attachment
: An active Cohesity rule that includes a profile and multiple ipsets.

Firewall Profile
: A group of TCP or UDP ports.

Full Archive
: A full copy of a Protection Group that is archived.

I

Incremental Archive
: An archive that records just the changed data since the most recent archive.

Inline Deduplication
: The process of eliminating redundant copies of data as the data is written to a Cohesity view.

Instant Mount
: A fast and reliable way to access data via the instant volume mounting feature. Some benefits include granular recovery of Microsoft Exchange, SQL and SharePoint data.

L

Logical Data
: Data residing in a protected source.

N

Node
: A physical or virtual appliance that runs the Cohesity software. One or more nodes comprise a Cohesity cluster.

P

Cohesity Data Cloud (Self-Managed)
: The Cohesity software consolidates secondary data and applications, including backups, files, objects, test/dev, and analytics, on a single, software-defined platform.

Protection Policy
: Business mapping on reusable settings that define how often to protect (RTO), how long to retain, how many copies to create like, archive, replication, tape, and where to store the data.

Protection Group
: A backup job that runs repeatedly, based on an associated policy, to back up data from a source and store it on the cluster. A Protection Group uses the schedules and

settings defined in the policy to determine when and how backups are captured, archived or replicated, the operational management of data to be protected etc.

R

Registration
: The process of registering a source with the Cohesity cluster so the source can be protected.

Remote Adapter Protection Group
: A backup job that runs scripts or executables on a remote Linux machine, typically for sources that are not supported by Cohesity.

Replication
: The process of automatically copying Protection Group snapshots from one Cohesity cluster to another Cohesity cluster for data resiliency purposes.

S

Source
: An object in the IT infrastructure of your company that you want to protect, such as a VM, physical server, Pure Storage volume, MS SQL Server, and NAS files.

Storage Consumed
: Data written with all copies required for honoring the resiliency setting.

Storage Domain
: A named storage location on a cluster. A storage domain defines deduplication, compression, erasure coding, fault tolerance and more. Views are created in storage domains. Storage domains were formerly known as View Boxes.

T

Test&Dev
: The process of cloning a VM, database, or view from a backup or existing file share for test and development purposes.

V

View
: A named storage location in the storage domain in a Cohesity cluster, which has NFS, SMB, and S3 mount paths.

Documentation Archive

Cohesity maintains and updates documentation for the versions listed in the following set of tables.

Actively Updated

The following documentation sets are actively updated for the time period from general availability (GA) to end of support (EoS). Active documentation updates include changes to the release notes, procedures, support matrix, fixed issues, and so forth.

Version	Release Notes	Release Date
7.4.1 documentation	7.4.1 release notes	First released June 2026
7.4 documentation	7.4 release notes	First released March 2026
7.3.2 documentation	7.3.2 release notes	First released February 2026
7.3.1 documentation	7.3.1 release notes	First released December 2025
7.3 documentation	7.3 release notes	First released October 2025
7.2.2 documentation	7.2.2 release notes	First released March 2025
7.2.1 documentation	7.2.1 release notes	First released November 2024
7.2 documentation	7.2 release notes	First released July 2024
7.1.2 documentation	7.1.2 release notes	First released March 2024 *This is a designated long-term support release version.*
6.8.2 documentation	6.8.2 release notes	First released March 2024 *This is a designated long-term support release version.*

As-is

The following documentation sets remain available, but Cohesity no longer provides updates. For more information on EoS of a release version, see Cohesity Products End of Support and End of Life.

Version	Release Notes	Release Date
7.1.1 documentation	7.1.1 release notes	First released December 2023
7.1 documentation	7.1 release notes	First released September 2023
7.0.1 documentation	7.0.1 release notes	First released June 2023
7.0 documentation	7.0 release notes	First released January 2023
6.8.1 documentation	6.8.1 release notes	First released December 2020
6.8 documentation	6.8 release notes	First released May 2022
6.6 documentation	6.6 release notes	First released December 2020
6.5.1 documentation	6.5.1 release notes	First released August 2020

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. Click here to send us your feedback!

Ensure that you provide the following details in your email:

- Document name
- Topic name
- Page number