

COHESITY

Release Notes

Version 7.4

April 24, 2026



© 2026 Cohesity, Inc. All rights reserved.

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

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

Published on April 24, 2026

Contents

Disclaimer	3
What's New?	3
Cohesity Feature Deprecation	11
Upgrading to 7.4	11
Considerations	19
Fixed Issues	30
Security Fixes	30
Cohesity Support	35
Documentation Feedback	36

Disclaimer

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

Cohesity provides from time to time - in release notes or in other communications to our customers - written updates about end of support for third-party software versions. Such updates are for informational purposes only, and are not a substitute for information you receive directly from third-party software publishers. Cohesity support practices align to third-party end of support, and as such Cohesity will not in any case support a version of third-party software that is no longer supported by its publisher. For further/up-to-date information, see the [Third-Party Software Support Matrix for Cohesity Data Protection](#).

What's New?

Cohesity Platform 7.4 provides new features and enhancements available for on-premises hardware, Cloud Edition, and Virtual Edition clusters. For more information, see [What's New in 7.4?](#).

For more information on upgrading from previous releases to 7.4, see [Upgrading to 7.4](#).

For more information on previous releases, see [What's New in Earlier Releases?](#)

What's New in 7.4?

The following new features and improvements are available in this release. For important information about upgrading from previous releases to 7.4, see [Upgrading to 7.4](#).

Early Access (EA) Feature

From time to time, Cohesity may add features and request for feedback on their utility and design. These features are termed as Early Access features. Early access features are limited to a closed group of testers for a limited subset of launches. Participation is by invitation only and may require signing a pre-general-availability agreement, including confidentiality provisions. These features may be unstable, change in backward-incompatible ways, and are not guaranteed to be released. There are no SLAs provided and no technical support obligations. These EA features are by default disabled and hidden and need to be enabled separately. If you wish to use these EA features you need to contact your accounts team, who will internally work within Cohesity to enable the feature on your cluster. Cohesity recommends running these features only on non-production clusters.

Data Protection

Cohesity Data Cloud for NetBackup

Data Cloud for NetBackup modernizes NetBackup on a single, scalable Cohesity Data Platform, delivering simplified operations, faster recovery, and built-in cyber resilience through unified management. This capability is available only for new deployments, which use a unified ISO for installation. For more information, see [Cohesity Data Cloud for NetBackup](#).

Databases

Agent & Connector Installation Status for MongoDB via Ops Manager

Cohesity now allows you to view the installation status of the Cohesity Agent and Connector across all MongoDB hosts managed via Ops Manager. This feature helps you quickly confirm backup readiness, detect configuration gaps, and reduce operational overhead, ensuring that all MongoDB nodes are fully prepared for reliable data protection. For more information, see [Register MongoDB Ops Manager Source](#).

Oracle RAC Alternate Restore Enhancements

Cohesity now supports selecting the target database type (Standalone or RAC) during Oracle alternate restores, including RAC-to-Standalone restores. When restoring to a RAC database, you can define the target RAC nodes and instance names using the new Target RAC Configuration section. For more information, see [Recover Oracle Database](#).

Cassandra Instant Recovery

Cohesity now introduces Cassandra Instant Recovery, a powerful new capability that reduces recovery times for Cassandra clusters. By retaining native Cassandra snapshots on the original source nodes and coordinating recovery using those snapshots, Cohesity enables a near-immediate recovery experience. For more information, see [Cassandra Instant Recovery](#).

Oracle Database Cloning Enhancements

Cohesity now supports enhanced Oracle database cloning behavior that includes guidance for managing database identifiers (DBID) during clone operations. This helps Oracle DBAs

ensure cloned databases can be correctly discovered and managed in Cohesity when required. For more information, see [Clone Oracle Database](#).

Physical Servers

FIPS-Certified OpenSSL 3 for Physical Agents (Windows and Linux x64)

Cohesity Windows and Linux 64 agents now utilize a FIPS 140-3 validated cryptographic module, strengthening cryptographic security and helping meet compliance requirements by providing an end-to-end FIPS 140-3 compliant solution.

Improved ACL Preservation and Authentication for File-Level Recovery

Cohesity now supports backup and restore of ACLs and extended attributes for Java agent-based backups, enabling higher-fidelity recovery.

File and Folder Recovery to Windows — SMB Authentication

Cohesity now supports configurable SMB authentication during file and folder recovery to Windows. You can choose an appropriate option based on your domain trust and agent configuration to ensure successful recoveries across different environments. For more information, see [Recover Physical Server Files or Folders to the Original Location](#).

NAS

REST API Support for NetApp ONTAP

Cohesity now supports NetApp REST APIs to interact with and manage ONTAP storage systems, in addition to NetApp's legacy Zephyr Application Programming Interface (ZAPI). For more information, see [Plan and Prepare for File-Runner Based NetApp ONTAP Backup](#).

Replication, SMB DFS, and CFT support for Nutanix Files Cluster Protection

Cohesity Nutanix Files protection now supports Replication, SMB DFS, and the vendor-native Change File Tracking (CFT) technology for faster incremental backups. For more information, see [Plan and Prepare for Nutanix Files Protection](#).

Support for Extended Attributes using SMB2 Protocol on NetApp ONTAP 9.15.1

Cohesity's NAS Data Protection now supports the restore of extended attributes of entities using the SMB2 protocol on NetApp. You can now preserve and restore the file content along with any associated extended attributes, which are metadata that provide additional information about the file. For more information, see [Backup NetApp ONTAP Using File-Runner](#).

Protect Nutanix Files Clusters

Cohesity now supports backup and recovery for Nutanix Files cluster deployments. You can protect multiple Nutanix Files clusters using Prism Central or one cluster at a time using Prism Element protection. For more information, see [Nutanix Files](#).

Microsoft 365

Exclusion of Teams Recordings

Cohesity now supports the exclusion of Teams recordings in Microsoft 365 Teams from the default location to help meet compliance or legal requirements. For more information, see [Recover Teams](#).

Error Report for Teams and Groups

Cohesity now supports downloading a list of all errors encountered during the backup of Teams or Groups. After the protection run completes, you can download a CSV file that reports the files that were skipped or encountered errors during the process.

Backup and restore of System Document Libraries and Templates

Cohesity now supports backup and restore of System Document Libraries and Templates for Microsoft 365 Teams and SharePoint. For more information, see [Plan and Prepare for Teams Protection](#) and [Plan and Prepare for SharePoint Protection](#).

Backup and recovery of OneNote Files

Cohesity now supports the backup and recovery of OneNote files in OneDrive, SharePoint, Teams, and Groups. For more information, see [Plan and Prepare for OneDrive Protection](#), [Plan and Prepare for SharePoint Protection](#), [Backup Microsoft 365 Teams](#), and [Plan and Prepare for Groups Protection](#).

Metadata Restore Support for SharePoint, OneDrive, Teams and Groups

Cohesity now supports metadata restore for Microsoft 365 OneDrive, SharePoint, Teams, and Groups. The restore includes four standard metadata fields: **Created By (author)**, **Modified By (editor)**, **Created Date**, and **Modified Date**. For more information, see [Plan and Prepare for OneDrive Protection](#), [Plan and Prepare for SharePoint Protection](#), [Plan and Prepare for Teams Protection](#), and [Plan and Prepare for Groups Protection](#).

Separate Recovery Task for Each Microsoft 365 Mailbox

Cohesity now creates separate recovery tasks for each Microsoft 365 mailbox. This allows you to run multiple recoveries in parallel, improving recovery performance and cluster resource utilization. For more information, see [Recover Mailbox](#).

You can also download separate PST files and view logs for individual mailboxes. For more information, see [Download the Exported Email or Mailbox Folder PST File](#).

Backup and recovery of Mailbox Recoverable Items

Cohesity now supports backup and restore of the Recoverable Items folders that includes Deletions, Purges, DiscoveryHolds, and SubstrateHolds in Microsoft 365 mailboxes. These folders preserve items that are either soft deleted or deleted from the Deleted Items folder. For more information, see [Recover Mailbox](#).

Virtualization

Recover vCD VM Placement and Sizing Policies on Alternate Recovery Location

Cohesity now allows restoration of placement and sizing policies when vCD VMs are recovered to their alternate location. For more information, see [Recover VMs from a VCD Source](#).

Support for Asynchronous Source Registration in VMware

Cohesity now supports asynchronous source registration for VMware, enabling faster environment registration while detailed environment information is retrieved in the background. For more information, see [Register or Edit a Hypervisor Source](#).

NFSv4.1 Protocol Support for VMware Recovery Operations

Cohesity now supports selecting the NFS protocol version (NFSv3 or NFSv4.1) when configuring VMware Instant Recovery, Test & Dev, Instant Volume Mount, Virtual Disk Recovery, and File-Level Recovery operations.

Support for Directly Attached Volume Groups on Nutanix AHV

Cohesity now supports backup and recovery of Directly Attached Volume Groups (VGs) for virtual machines running on Nutanix AHV. Volume Groups attached to a VM can be included during backup and optionally restored during recovery. For more information, see [Directly Attached Volume Groups on Nutanix AHV](#) and [Recover a VM with Directly Attached Volume Groups](#).

Support for Concurrent Backup Stream per Datastore for Nutanix AHV

Cohesity now supports capping concurrent backup streams per datastore for Prism Central and Prism Element, allowing you to limit the number of streams backed up simultaneously to reduce performance impact on the datastore. For more information, see [Register or Edit a Hypervisor Source](#).

Kubernetes

Sequentially Quiesce Kubernetes Workloads During Backups

Cohesity now supports sequential quiescing of Kubernetes workloads during backups. For more information, see [Quiesce Kubernetes Workloads during Backups](#).

External Targets

Object Lock Mode Selection for AWS Incremental Forever Archival

Cohesity now supports selecting an Object Lock Mode when registering an AWS target with incremental forever archival. The UI displays options for Governance (default) and Compliance modes during registration. For more information, see [Register an External Target for Archival](#).

SmartFiles

Non-Indexed Browse Support for S3 Views

Cohesity now supports non-indexed browsing for S3 Views, allowing you to browse objects within the S3 bucket hierarchy without relying on indexing. With this enhancement, download and recovery workflows are disabled for S3 Views. For more information, see [Recover to the Original View](#).

View Trust Relationships for Domains

Cohesity now allows you to view the topology of trust relationships between Active Directory (AD) domains. This functionality helps you analyze complex AD topologies with multiple domains and trusts. For more information, see [Trusted Domains](#).

Cohesity Cloud Edition and Cohesity Cloud Edition (Legacy)

Cohesity Cloud Edition Cluster and Cohesity Cloud Edition Cluster (Legacy) on AWS Secret cloud

You can now deploy Cohesity Cloud Edition Cluster and Cohesity Cloud Edition Cluster (Legacy) on AWS Secret Cloud. For more information, see [Cloud Edition Cluster](#) and [Cohesity Cloud Edition Cluster \(Legacy\)](#).

Deploy Azure Control VM Using Ubuntu Image

You can now set up the Control VM used for deploying a Cloud Edition cluster on Azure using any Ubuntu Long-Term Support (LTS) image. For more information, see [Cloud Edition cluster on Azure](#).

Backup and Recovery Support for Pinecone

Cohesity now supports the backup and recovery of Pinecone vector database indexes, which are commonly used in AI and machine learning workloads such as Retrieval-Augmented Generation (RAG), semantic search, and recommendation systems. For more information, see [Pinecone](#).

Cohesity integrates directly with Pinecone's native Backups API to protect serverless indexes without moving customer data out of the Pinecone cloud.

Cohesity Cloud Edition Cluster on Oracle Cloud Infrastructure

You can now deploy Cohesity Cloud Edition Cluster on Oracle Cloud Infrastructure (OCI). For more information, see [Cohesity Cloud Edition Cluster on Oracle Cloud Infrastructure \(OCI\)](#).

Support for Amazon DocumentDB Data Protection

Cohesity now supports the protection of Amazon DocumentDB at database-level by ingesting data into Cohesity Cloud Edition Cluster on AWS. For more information, see [Amazon DocumentDB](#).

Support for Azure SQL Database Protection

Cohesity now supports the protection of SQL databases at database-level by ingesting data into DataProtect for cloud in Azure. For more information, see [Azure SQL Database](#).

Enhanced Azure SQL Database Protection

Cohesity now allows you to set the following fields during Azure SQL database protection:

- During Backup:
 - Change Data Capture (CDC) for Backups - Cohesity requires explicit user consent to enable and use CDC data during backups. As CDC data is mandatory for SQL database protection, this option is presented in the UI to clearly communicate this intent. For more information, see [Protect Azure SQL Databases](#).
 - Batch Count - number of CDC records processed in batches to prevent consuming the entire temporary database. For more information, see [Protect Azure SQL Databases](#).
- During Recovery:
 - Elastic Pool - pre-existing elastic pool associated with the server. For more information, see [Recover Azure SQL Databases](#).

Azure SQL Managed Instance Protection Without CDC

Cohesity now supports the backup and recovery of Azure SQL Managed Instance without Change Data Capture (CDC). This allows you to perform full backups using the native backup method, improving the backup speed and efficiency. You can also specify the batch count when CDC is used in backup. For more information, see [Protect Azure SQL Managed Instance](#).

Azure VM Disk-level Recovery

You can now perform disk-level recovery of Azure VM on Cohesity Cloud Edition on Azure. For more information, see [Recover Azure Disk](#).

Support for Managed Identity Authentication for Azure Storage External Targets

Cohesity Cloud Edition cluster on Azure now supports registering Azure Storage as an external target using Managed Identity authentication for down-tiering (primary copy) and archiving (secondary copy) data. For more information, see [Cohesity Cloud Edition cluster on Azure](#).

Support for Impersonated Service Account Authentication for Google Cloud External Targets

Cohesity Cloud Edition cluster on GCP now supports registering Google Cloud as an external target using impersonated service account authentication for down-tiering (primary copy) and archiving (secondary copy) data. For more information, see [Cohesity Cloud Edition cluster on GCP](#).

Support for Cloud SQL for MySQL Server Data Protection

You can now perform database-level backup and recovery of Cloud SQL for MySQL Server using Cohesity Cloud Edition Cluster on GCP. For more information, see [Cloud SQL for MySQL](#).

Support for Bigtable Data Protection

You can now perform table-level backup and recovery of Bigtable using Cohesity Cloud Edition on GCP. For more information, see [Bigtable](#).

Unified GCP Cloud source registration

Cohesity DataProtect introduces a unified GCP Cloud source registration experience. In this release, Cohesity allows you to select multiple GCP workloads for the source registration. This enhancement simplifies the source registration by removing the need for multiple registrations. For more information, see [Register a Google Cloud Source](#).

Cluster Management

Data Source Connection for Multi-Tenant Network Support

Starting with this release, you can use Data Source Connection alongside Hybrid Extender to support multi-tenant environments where isolated networks may use overlapping IP address ranges and subnets. For more information, see [Data Source Connection](#).

Data Source Connectors provide enhanced lifecycle management, including in-place upgrades and controlled patching. This eliminates the need to redeploy connectors during version upgrades and reduces operational overhead while improving manageability and scalability in multi-tenant environments.

Web Server Certificates

Added the Web Server Certificates tab to view and manage the certificate used by the Cluster UI. For more information, see [Web Server Certificates](#).

Security

FortKnox Self-Managed

FortKnox Self-Managed now includes the following enhancements:

UI Support for Just-In-Time Replication to Primary Cluster

FortKnox Self-Managed now supports Just-In-Time replication from vault cluster to primary cluster using the Cohesity UI, allowing you to replicate snapshots to the remote replication cluster. For more information, see [Supported Vaulting and Recovery Configurations](#).

Improved Vaulting Status Visibility

Vaulting status for protection runs is now consistently available from both the primary cluster and the vault cluster, with enhanced filtering to quickly identify specific run states. For more information, see [Vaulting Status](#).

Indication of Paused Protection Runs

Protection runs that fall outside the configured vaulting window are now shown with a Paused status, making it easier to identify runs that were skipped due to vaulting window constraints. For more information, see [Vaulting Status](#).

Retention Validation for Vault Protection Policies

FortKnox Self-Managed now validates that the vault retention period is greater than or equal to the primary backup retention when configuring vault protection policies. For more information, see [Create Protection Policy for Vault Cluster from Primary Cluster](#).

Download Paired Clusters Report

You can now download a CSV report from the vault cluster that lists all paired clusters along with pairing details, vaulting windows, and associated protection groups. For more information, see [Pair Primary Cluster to Vault cluster](#).

Refresh API Keys

You can now refresh API keys for a vault cluster directly from the user interface, enabling communication to restore when the primary cluster is reinstalled without unpairing from the vault cluster. For more information, see [Refresh API Keys](#).

Vault Selected Objects in Edit Run

You can now edit a vault-enabled protection group run to vault only selected objects instead of reprocessing the entire group, giving you more control and reducing unnecessary vault operations. For more information, see [Edit a Protection Group Run to Vault Specific Objects](#).

Cohesity Feature Deprecation

There are no features that are deprecated in the Cohesity 7.4 release.

Upgrading to 7.4

Upgrade Paths

You can upgrade your Cohesity cluster from previous releases to 7.4. The following table provides details on supported upgrade paths.

Your Current Release	Upgrade Path to 7.4
<ul style="list-style-type: none"> • 7.3.2 • 7.3.1 • 7.3 • 7.2.2_u2 • 7.2.2_u1 • 7.2.2 • 7.2.1 • 7.2 • 7.1.2_u6 • 7.1.2_u4 • 7.1.2_u3 • 7.1.2_u2 • 7.1.2_u1 • 7.1.2 • 7.1.1 • 7.1 	7.4 directly

Note: Direct upgrades from 6.8.x to 7.4 are not supported. If your cluster is running any 6.8.x release (such as 6.8.1, 6.8.1_u1, or 6.8.2), a step upgrade is necessary. First, upgrade to 7.1.2_u6, and then proceed to 7.4.

Release Upgrade Policy

Policy	Example
<p>Cohesity will support upgrades from the latest release of the prior LTS release branch, which includes all LTS designated releases within the branch, to the most recent release of the current LTS branch.</p>	<p>6.6.0d+ (LTS designated releases: 6.6.0d_u3, 6.6.0d_u4, 6.6.0d_u5, 6.6.0d_u6) to 6.8.2 LTS designated release will be supported.</p>
<p>Cohesity will not allow upgrades by default to any release that is older in time irrespective of the release branch. Exceptions are to be managed on a case-by-case basis.</p>	<p>6.5.1f_release-20210825_596bb917 is released after 6.6.0c_release-20210822_0d731348. Therefore, an upgrade from the 6.5.1f version to the 6.6.0c version is not supported.</p> <p>This policy is also applicable to patches. If you have upgraded your Cohesity cluster to a patch released after the LTS release, upgrading to that LTS release is not supported. However, you can upgrade to any LTS version released after the patch. For example, your Cohesity clusters were upgraded to 6.6.0d_u5 in July 2022. Cohesity released 6.8.1_u1 on Nov 2022 and the 6.6.0d-p32 patch on March 2023. If you've applied 6.6.0d-p32, you cannot upgrade to 6.8.1_u1. However, you can upgrade to the upcoming 6.8.1_u2 release.</p>
<p>Cohesity will support the release N-1 upgrade without an intermediate step. (N is defined as the current release branch).</p>	<p>7.2.x to 7.4 is supported. 7.4 is the current release branch.</p>
<p>Cohesity will support the release N-2 upgrade without an intermediate step. (N is defined as the current release branch).</p>	<p>7.1.x to 7.4 is supported. 7.4 is the current release branch.</p>
<p>When a specific release is declared LTS, Cohesity will support upgrading from the open LTS releases to the new LTS release. This will include the three most recent releases on the LTS branch to the new LTS release.</p>	<p>6.8.1, 6.8.1_u1, 6.8.1_u2, 6.8.1_u3, 6.8.1_u4, 6.8.1_u5, 6.8.1_u6, 6.8.1_u7 to 6.8.2.</p>

Upgrade Considerations

Note the following about upgrading the Cohesity cluster to 7.4:

- Cohesity does not support rolling back to older versions.
- To upgrade the Cohesity cluster from a version that is no longer supported, Cohesity recommends you to upgrade to any of the supported versions mentioned in the [Upgrade Paths](#), and then perform an upgrade to the latest release version. For information on Cohesity Products that have reached the end of support, see [Cohesity Products End of Support](#).
- See to review the list of features marked for deprecation for Cohesity 7.4 and later releases.
- Before performing the upgrade, ensure that the cluster data space and metadata space utilized is less than 85%. After the cluster upgrade, the Garbage Collection algorithms take 3 to 4 days to trigger. Hence, ensure that the cluster has enough space during this period. Space constraints may lead to backup and replication failures on the Cohesity cluster.
- If you are running remote adapter jobs and the cluster is upgraded, the jobs will be disrupted during the upgrade process. The jobs will be killed and restarted multiple times during the upgrade.
- Starting 6.8.2 and 7.1.2, the Cohesity indexing service is optimized to automatically identify and delete stale directories at regular intervals, which were created for indexing. After upgrading from a version without this optimization, the cluster indexing service will remove any stale directories identified, which may result in cluster-free space increase.
- Cohesity recommends upgrading the Cohesity Agent on Physical Servers and the Cohesity installed Agent on VMs to the latest release version of the Cohesity cluster.
- Cohesity recommends upgrading the Cohesity cluster first, followed by the Cohesity Agent. Upgrading an agent before the cluster is likely to impact the existing functionality and disruptions may be observed due to agent being on a higher version than the cluster. Cohesity also recommends the agents be on the same, latest major version as the Cohesity cluster to get the latest security fixes and benefit from newer features.
- After upgrading to the latest version, if there is an IP subnet conflict, the **Enable Apps Management** toggle in **Marketplace > My Apps** is turned off. Navigate to **Settings > Summary** > click **Configure** and specify a different IP address in the **Configure Apps management network** field and then turn on the **Enable Apps Management** option.
- If you are on a Cohesity Cloud Edition cluster and using Marketplace Apps, then when you upgrade the Cohesity Cloud Edition cluster to 7.4, connectivity among the Marketplace Apps could be impacted* due to Flannel moving to etcd v3 APIs. It is

recommended to pause any Marketplace Apps before the upgrade and resume them once the upgrade is complete.

***Impact:** Running workloads, Protection Groups, or scans related to the Marketplace App might see network disruption during the upgrade.

- For pure PXG clusters, before upgrading from version 6.6 to 6.8.1 or above, make sure that your cluster usage is below 95%. After the upgrade, there is a known issue where the available data space may decrease. Even clusters that are using only 88% of disk space before the upgrade have experienced out-of-space errors afterward, which can lead to backup failures. To avoid disruptions, Cohesity strongly recommends reducing disk usage well below 95% before starting the upgrade process.
- Cohesity is working to qualify the Windows agent for Japanese Windows version that uses Shift-JIS encoding. Wait to upgrade agent to version 7.4 until Cohesity releases a version that supports Japanese language (Shift_JIS) based Windows OS.
- Manual disk-space verification on nodes is not required. During the upgrade, the system performs pre-upgrade checks and automatically cleans up disk space when possible. If sufficient space is still unavailable, the upgrade stops immediately and provides a clear error message indicating the specific node and path that require additional space.
- If you are leveraging storage snapshots for VMware VM backups in releases prior to 7.2 and plan to upgrade the Cohesity cluster to release 7.2 or later, ensure that SAN connectivity between Cohesity nodes and the Storage Array (Pure, HPE Alletra, Nimble, and IBM) is in place before the upgrade to prevent VM backup failures.

Databases

- Azure Cosmos DB - MongoDB is not supported on the Cohesity 7.4 version.
- The addition of the new Postgres database could cause UI slowness until the ETL process completes. The bootstrap run of the ETL process pulls the entire data set to populate the database. The initial run has a slight performance impact. In the case of upgrades, data population happens in the post-upgrade step. Subsequent upgrades will not be affected.
- After upgrading to the latest version, [to display SAP HANA log backups](#) in the Cohesity cluster, you need to modify the existing registered source and set the `et-log-backup` source registration parameter to `true`. Only the log backups triggered after enabling `et-log-backup` will be shown on the Cohesity cluster.

Note: After modifying the source configuration (with `et-log-backup=true`), a full backup is mandatory. The initial full backup must be completed before any log backups appear on the Cohesity user interface.

- If you are upgrading to version 7.4 or later, you need to [update the existing SAP HANA source configuration](#) to enable auto-discovery and entity hierarchy. Set the `--entity-hierarchy` source registration parameter to `"true."` After updating the source configuration (with `--entity-hierarchy=true`), a full backup is mandatory.
- For **SQL log backups**, if you are upgrading the cluster to version 7.2 and above, ensure that the Cohesity cluster bridge node VIPs on port 11117 are reachable from the SQL source. In case of the multitenant environment, ensure that the Hybrid Extender IPs on port 11117 are reachable from the SQL Source.

Note: The SQL log backup will fail post-upgrade if the above-mentioned port requirement is not satisfied.

- If you have upgraded your cluster to Cohesity 7.4 from an earlier version and already have an AWS account registered, you must add [new permissions](#) to protect the newly supported AWS workloads:
 - Amazon RDS (Ingest-based protection)
 - Amazon Redshift
 - Amazon DocumentDB
 - Amazon DynamoDB

Administration

- To generate a new SSH key after upgrading the Cluster, contact [Cohesity Support](#).
- Cohesity Support Engineers require a Support Channel token to remotely log into the Cluster using SSH for on-demand assistance. From your Cohesity cluster, you need to [copy the Support Channel token](#) and provide it while raising a request for on-demand assistance.
- The Secure Shell restricts access to the host commands or scripts. After you upgrade to 6.7 or later version, the secure shell might have the following impact on your existing Cohesity Data Cloud deployments:
 - Access to the bash shell using SSH will be no longer available to the support user account without authorization from Cohesity.
 - If you run custom scripts using SSH on your Cohesity cluster, the scripts may fail. In this case, Cohesity recommends the following:
 - Verify if there is an alternate method to use Cohesity CLI commands or REST API and update your scripts accordingly.
 - Verify if a corresponding Cohesity CLI command is available in the supported list of CLI commands; if so, use the supported CLI command. If

the CLI command is not available in the supported list of commands, contact [Cohesity Support](#) to enable the CLI command.

- The private binary or tools running on the Cohesity nodes might fail. Contact [Cohesity Support](#) for options to install private binaries or tools.
- Sudo access is disabled by default. For support channel access, enable the sudo access. For more information, see [Enable or Disable Linux Sudo Access](#).
- If there is a source that is registered before the upgrade and assigned to an organization, then unassigning its root entity is not allowed. You can unassign the source if it is not assigned to an organization, and it will get assigned after the upgrade.

NoSQL and Hadoop

- To continue using Cohesity NoSQL & Hadoop services on the Cohesity cluster version 7.2.2, you must upgrade the NoSQL & Hadoop service to the 7.0.0 version available on Helios.
- If you are running NoSQL and Hadoop app, Cohesity recommends the following before upgrading the Cohesity cluster:
 - Pause the protection runs by navigating to **Data Protection > Protection** . From the Action Menu (:) of the required protection run, select **Pause Future Runs**.
 - Pause the NoSQL and Hadoop app by navigating to **Marketplace > My Apps** . From the Action Menu (:) of the app instance, select **Pause**.

After upgrading the Cohesity cluster to the latest version, contact your Cohesity account team to check if the upgraded Cohesity cluster requires a new NoSQL and Hadoop app. If it requires a new version of the app, you must upgrade to the latest version of the NoSQL and Hadoop app. Once the cluster upgrade is complete, resume the app, and then the protection runs.

Microsoft 365

- If you upgrade your Cohesity cluster to 6.8.2 or later versions and currently backing up Microsoft 365, ensure that you add the required [Microsoft Graph Permissions](#) related to MS Groups to your custom application to continue using your existing Protection Groups and protect your Microsoft 365 data.
- After upgrading to the 7.4 version, if you are replicating the Mailbox data to a remote Cohesity cluster, then ensure that you upgrade the remote Cohesity cluster to the 7.4 version.

Single Node Cluster Upgrades

Single node cluster upgrades must be run when the upgrade will have the least impact. During the upgrade of a single node cluster, the node is rebooted and during the reboot, the cluster is unable to process Protection Groups, recover tasks, or any other workflow.

Virtual Edition Deployment

The following are the requirements for the Virtual Edition deployment for 6.8 and later versions:

- small (8 TB) configuration supports Virtual Machines with 12 vCPUs, 32 GB of memory, and 64 GB virtual disk to store the operating system.
- large (16 TB) configuration supports Virtual Machines with 24 vCPUs, 64 GB of memory, and 64 GB virtual disk to store the operating system.

For more information, see [Virtual Edition for VMware Setup Guide](#) and [Virtual Edition for Clustered VMware Setup Guide](#).

Replication Environments

- If the cluster replication is configured, verify that the network connectivity is functioning properly during the upgrade to ensure the cluster replication relationship is successfully upgraded to use AES-256-GCM for encryption.
- In a replication setup, when you upgrade your Cohesity cluster to Cohesity 6.6 or later and you use the default System Admin password, you will be prompted to change the password. After changing the password, you must update the new password on the replication partner cluster.
- For information about using replication between Cohesity clusters running different versions, see [Replication Compatibility](#).

Cohesity Cluster Patch Upgrades

- Ensure there are no cluster operations or patch updates in progress. A cluster operation is a task on a Cohesity cluster such as add or remove a node, and cluster upgrade.
- When you create a node and connect it to a Cohesity cluster, the service patch updates are done automatically but the Base OS patch is not applied. To apply Base OS patch update on the newly added node, you can refer to the link under the **Instructions** column in the [Download portal](#).

Note: Cohesity recommends that the product patch and the Base OS patch version should be the same.

Patch Upgrades in DoD Mode

If your Cohesity cluster is running on DoD mode, then you should first upgrade to 6.8.1_u2 or later and then apply a cluster patch update. For more information on DoD mode, see [Use Cohesity in DoD Mode](#).

Supported Sources for Hybrid Extender Based Organizations

From 6.6 onwards, Cohesity Platform in a multi-tenant environment displays only the sources that the Organization (tenant) can register and protect. As a prerequisite, Hybrid Extender should be enabled for Organizations (tenant).

For a list of supported sources and workflows, see [Supported Multitenancy Workflows](#).

FortKnox

After upgrading to the 7.3 version, cold vaults created in an AWS Region use Amazon S3-Glacier Deep Archive (GDA) and require a minimum retention period of 180 days. In earlier releases (prior to 7.3), cold vaults were created on Amazon S3-Glacier Flexible Retrieval (FR). Recovery starts up to 12 hours after initiation, due to the data hydration process required for data access from cold vaults.

Considerations

Review these considerations before you install the software for the first time or upgrade from a previous version.

Data Protection

Instant Volume Mount

Review the following considerations:

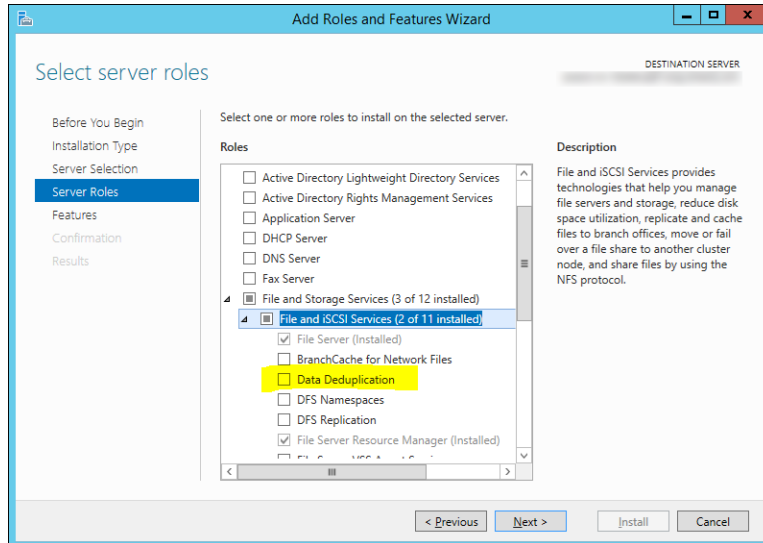
- When recovering a file or instantly mounting a volume from a Windows VM or Physical Server Backup Source that has Windows deduplication installed and enabled for one or more volumes, you must choose a target machine that also has Windows deduplication installed (it does not have to be enabled for any volume). (However, this rule does not apply to Nutanix AHV VMs. If AHV VMs are enabled with Windows deduplication, the only supported recovery option is full VM recovery.)

If the target does not have Windows deduplication installed:

- File level recovery will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."
- Instant volume mounting will fail with the error message: "Windows Data Deduplication role is not installed on the target machine. Retry recovery after installing the Windows Data Deduplication role on the target machine."

To determine if Windows deduplication is installed on the Source or target machine, follow the steps given below:

1. Open **Server Manager**.
2. Select **Roles and Features > File and Storage Services > File and iSCSI Services**.
3. Select the **Data Deduplication** check box, if necessary.
4. Click **Next** until the **Install** button is enabled and then click **Install**.



- Instant Volume Mount (IVM) restore of ReFS volumes backed up using Windows physical block-based jobs cannot be restored to an alternate Windows server running a lower version of ReFS.
- When mounting volumes on a Linux physical Server, the loop devices present on the Server are used for mounting. Therefore, the number of volumes that can be mounted depends on free loop device availability. By default, the number of available loop devices is 8, but this number can be customized. If the number of configured loop devices is the default of 8, up to eight volumes can be mounted. In this example, if an attempt to mount more than 8 volumes occurs, the mounting of all the volumes after the 8th volume fails and errors are reported.
- Tearing down a cloned database or instant volume mount deletes the mounted volumes. Any new or modified data on these volumes will be deleted along with the volumes, so ensure you back up any important data before teardown.
- Review the following considerations when performing instant volume mount to Hyper-V VMs:
 - Only Windows VMs are supported.
 - Dynamic Disks (LDM and LVM) are not supported.
 - The Bring Disks Online option requires the following:

- VM must be part of Active Directory, the VM and the Hyper-V host must be in the same AD.
- Users must execute "winrm quickconfig" to enable winrm on the target VM and remote powershell must be enabled from Hyper-V host to VM.
- Instant volume mount and file level recovery from Gen 1 to Gen 2 type VMs is not supported.
- If SCVMM is unregistered from the Cohesity cluster, ensure you tear down all instant volume mounts. Not tearing them down can prevent the VM from being backed up when the source is registered with a different Cohesity cluster.
- Instant volume mounting Hyper-V 2012 R2 VMs without a SCSI controller is not supported. This is because Hyper-V disallows dynamically adding a SCSI controller, which is required to add the virtual disks.
- On 2012 R2 VMs, if an instant volume mount disk is attached to a VM during a Protection Group run, that snapshot cannot be application-consistent. If this occurs, the event viewer may contain a VSS-catastrophic error or similar message.
- Instant Volume Mount for NetApp stub file is not supported.
- You cannot instantly mount a volume from a VM to a physical server, and vice-versa.

File and Object Services

NFS

Review the following considerations:

- NFS mount names and names of files contained in the mount support ASCII and UTF-8 character codes only.
- When mounting a View, the `-o atime` option for the `mount` command improves the performance marginally. For performance reasons even if you specify the `atime` option, the Cohesity cluster does not record the access time. The `-o noatime` option is always in effect and the Cohesity cluster only records the access time when files are created or modified.
- When data is deleted from a view, it may take up to a day for the disk space to become available again and visible from utilities such as `df`.
- To register an Oracle RAC or a RAC node as physical server, "host" command must be executed on each of the nodes of that RAC.
- Cohesity recommends using a Linux client with kernel version 4.x or higher.
- NFSv4.1 considerations:
 - If you use a single client machine to mount an NFS4.1 view with different node IPs, all mount requests will go to a single node on the Cohesity node and might

result in inefficient workload balancing.

- **Workaround:** If you want to mount a single NFSv4.1 View using different node IPs, Cohesity recommends to use multiple clients for better performance. However, you can use each of these NFS clients to mount Views from different Cohesity clusters.
- LOCKT operations are not supported.

SMB

Review the following considerations:

- Keeping with the industry standard of change notification for SMB shares, recursive change notifications are not sent due to their effect on process load and network traffic.
- Filenames that contain UTF-16 character codes ranging from U+D800 to U+DFFF are not allowed in Cohesity SMB shares.
- For Linux clients that are members of AD, using "client max protocol = SMB2" in the [global] section of /etc/samba/smb.conf is not supported. Use "client max protocol = SMB3".
- Cohesity SMB shares do not support alternate data streams.
- You can add Cohesity SMB shares as a [Microsoft Distributed File System \(DFS\)](#) target, but note that SmartFiles does not support any additional features or functionalities provided by [Microsoft DFS](#).
- Windows behavior prevents Cohesity SMB shares from being automatically discoverable. Use the `net view` command to probe the cluster explicitly using `\\<Cluster-machine-account-name>` or `\\<Cluster-vip-FQDN>` or `\\<Cluster-VIP>`.

SMB Multichannel

Review the following consideration:

The option to advertise multiple IP addresses on the cluster is not supported.

S3

Review the following considerations:

- You must use one of the following accounts to create an S3 View:
 - A local Cohesity user.
 - An Active Directory user that was explicitly added to the Cohesity cluster and assigned a role. This user does not rely on an AD group for access to the Cohesity cluster.

Important:

You cannot create an S3 View using one of the following accounts:

- An AD user that has Cohesity cluster access through an Active Directory group only
- An SSO user
- A Helios user

- To create a SmartFiles S3 View in a multi-tenant environment, log in to the Cohesity cluster as an Organization user. If you create the S3 View while impersonating an organization, the Service Provider administrator becomes the owner of the S3 View.
- Access Control Lists (ACLs) can be set on a bucket using the AWS CLI.
- You cannot use NFS to mount newly created S3 Views. However, if there are existing S3 Views that were configured to use NFS, you can mount such S3 Views using NFS.
- The maximum number of versions allowed per S3 object is 500,000.
- Cohesity recommends excluding any unsupported header(s) from your requests. By doing so, you can prevent any potential unintended consequences that may arise from using unsupported headers.
- By default, only non-indexed browse is supported for S3 Views.
- Download and Recovery workflows are not supported for S3 Views.
- Browsing archived S3 snapshots is not supported.

Indexing and File Recovery

Review the following considerations:

- The Indexing Helper Service is not supported on a Cohesity cluster that is running on DoD mode. When DoD mode is not enabled, both the proxy and the host machines are available and there is improved resiliency for mounting of volumes. This improved resiliency is lost when the entire dependency is on the host node to perform the volume mounts.
- The Cohesity cluster attempts to index all files and folders to a drive on both Windows and Linux systems. If the Cohesity cluster is unable to find mount point information about files or directories, it indexes and displays these files and directories in the `lv01_N` directory, where `N` is a unique number such as `1`.

On Windows systems, if the Cohesity cluster finds the mount point information about files and directories, it indexes and displays these files and directories with a drive letter such as `C:`.

Linux LVM indexing supports the following LVM types only: Linear, Striped, Mirrored, Mirrored + Striped, Thin. On Linux systems, how files and directories are indexed and displayed is dependent on the conditions specified in the following table.

Server Type	Volume Type	
Linux Virtual Machine	Simple Volume	<p>The Cohesity cluster detects mount points for entries in the <code>/etc/fstab</code> file with the following formats:</p> <pre>UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0</pre> <pre>UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0</pre> <p>If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point that was specified in the <code>/etc/fstab</code> file. For these example entries, files and directories are indexed with the <code>/mnt</code> mount path, such as <code>/mnt/example/test.txt</code>.</p> <p>If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a <code>lvvol_N</code> directory. For example, the <code>/mnt/example/test.txt</code> file is indexed as <code>/lvvol_1/example/test.txt</code>.</p>
Linux Virtual Machine	LVM Volume	<p>The Cohesity cluster detects mount points for entries in the <code>/etc/fstab</code> file with the following formats:</p> <pre>UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0</pre> <pre>/dev/mapper/VG1-root /mnt ext4 defaults 1 1</pre> <pre>/dev/VG1/root /mnt ext4 defaults 1 1</pre> <p>If the Cohesity cluster can detect a mount point, it indexes and displays files and directories in the volume with the mount point specified in the <code>/etc/fstab</code> file. For these example entries, files and directories are indexed with the <code>/mnt</code> mount path, such as <code>/mnt/example/test.txt</code>.</p> <p>If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a <code>lvvol_N</code> directory. For example, the <code>/mnt/example/test.txt</code> file is indexed as <code>/lvvol_1/example/test.txt</code>.</p>
Linux Physical	LVM Volume	<p>The Cohesity agent can only return mount data when the volume is mounted on the Linux physical Server. If the volume is mounted, the Cohesity cluster indexes and displays files and directories in the volume with the mount point such as <code>/mnt/example/test.txt</code>.</p> <p>If the volume is not mounted, the Cohesity cluster indexes the files and directories into a <code>lvvol_N</code> directory. For example, the <code>/mnt/example/test.txt</code> file is indexed as <code>/lvvol_1/example/test.txt</code>.</p>

- Cohesity supports recovering files/folders from NTFS (Windows VMs) to Windows VMs, and from Linux VMs to Linux VMs only.
- **Error:** When recovering files or folders, the virtual disks are part of the target VM. These virtual disks are attached as SCSI disks that can be any of the supported adapter types: LSI Logic Parallel, LSI Logic SAS or VMware Paravirtual. During this step, you may encounter the following error: "Disk adapter with required slots - <n> is not available. Try creating a new adapter". Here, <n> is the number of virtual disks that are being attached. This can occur if the VM's disk adapter does not have the required number of slots (one SCSI adapter can support 15 virtual disks).

Solution: Attempt the operation *after* creating a new SCSI adapter. Additionally, the number of virtual disks where files and folders are being recovered from is limited to 15 at a time. Remove some files (or folders) and retry the recovery.

- For RHEL7, if Open VM Tools is installed instead of VMware Tools, TMPDIR may not point to /tmp. When recovering to location "/tmp/<SOME_DIRECTORY>", files may be recovered to a different location.

Example: If the recovery location is '/tmp/DIR1', files are recovered to a different location, such as '/tmp/systemd-private-c74aea179e9a43c789a19306d880274f-vmttoolsd.service-9GhOBD/tmp/DIR1'

- When unzipping a zip file that was created by downloading files and folders from an archived Snapshot, if the file or folder name has encoded characters, unzip the zip file using the corresponding encoding. For example if a file name in the zip file has a UTF-8 character, unzip the file using the following command:

```
unzip -O UTF-8 Download-Files_Sep_20_2018_3-17pm_3090.zip
```

- For Linux VMs, Cohesity supports file recovery from LVM volumes. One LVM volume can consume more than one loop back device, so Linux VMs may support fewer than 8 volumes when configured with the default number of loop devices.
- When recovering a Linux file, the Cohesity Linux Agent runs the following commands in sudo:
 - mount
 - umount
 - findmnt
 - timeout
 - blkid
 - lsof
 - ls
 - rsync
 - losetup
 - dmsetup

- lvs
 - vgs
 - lvcreate
 - lvremove
 - lvchange
- For Linux Logical Volume Manager (LVM), if all the disks for a volume group are not found by the Cohesity cluster, the Cohesity cluster will not process that volume group. As a result of that, no volumes of this volume group will be recognized or indexed by the Cohesity cluster.
 - Indexing, file recovery and browsing files and folders on VMs are not supported for drives with disk-level encryption (such as BitLocker). On physical Servers, however, these workflows are supported.
 - Encrypted VMs are not indexed.
 - If a Windows VM includes volumes created from a storage pool (Microsoft Storage spaces), VMDK recovery, IVM, and FLR are not supported.
 - Cohesity does not support indexing of Microsoft Storage Spaces.
 - File level recovery for VMware ESXi environments does not support RAID-5 volumes on dynamic disks. Simple, striped, spanned and mirrored volumes on dynamic disks are supported.
 - A VMware Tools service restart during a Recovery operation may disrupt Recovery. If the VMware Tools service restarts during a Recovery operation, the following error message is returned: The guest operations agent could not be contacted. After multiple retries to contact the guest operations agent, an error message stating that it started the copy but it could not get the status is returned. Go to the recovery location to verify whether the operation succeeded.
 - Recovering files to a VM where vMotion is in process is not supported.
 - File recovery is not supported for ReFS volumes in these environments: physical, VMware, Hyper-V and AHV.
 - Encrypted folders that have been renamed or deleted cannot be recovered.
 - Recovering files/folders with names longer than 200 characters may return an error. This is due to Windows behavior when handling files/folders with long names.
 - After making system configuration changes to a Windows 8 or Windows 2012 System VM, such as renaming an existing drive letter or adding a new disk, these changes may not immediately take effect due to a Windows registry refresh issue. To force the drive letters to be updated on the VM, reboot the system in the VM. This issue affects how files are indexed by the Cohesity cluster and displayed while browsing the contents of the VM.

- Considerations when recovering to physical servers that run:
 - Windows 2012 or later - None

If the OS does not support your recovery, you must recover to an alternate physical server running Windows 2012 Server or later, or use downloads.
- File-based recovery to Windows VMs does not support hardlinks and alternate data streams.
- Downloading files and folders from tape archive locations is not supported.
- Recovering files and folders from VMs to physical servers and from physical servers to VMs is not supported.
- The downloadable zip file can contain regular files and folders only; symlinks are not supported. When unzipping the downloaded files/folders, use a zip utility that supports the ZIP64 format.
- Recovering files to Linux VMs is not supported in the following cases:
 - When run as a non-root user that does not have sudo access
 - If `ALL=(ALL) NOPASSWD:ALL` is not set for the recover user in the `/etc/sudoers` file
 - If `requiretty` is not disabled for the recover user in the `/etc/sudoers` file

Recovering to Linux VMs requires `requiretty` to be disabled for the recover user in the `/etc/sudoers` file, otherwise recovery will fail. To disable `requiretty` for a recover user Add the following line in the `/etc/sudoers` file, where `<USERNAME>` is the name of the recover user with sudo access: `Defaults: <USERNAME> !requiretty`
 - The recovery directory path length is greater than 4096 characters.
 - There is not enough space in `/tmp` for Cohesity to push `linux_agent`.

Replication and Archival

Review the following considerations:

- Backups that are taken on the Full (No CBT) schedule are not currently archived by the Cohesity cluster. Other full backups (first Protection Group run, failed CBT) can be archived because they are not initiated by the Full backup schedule.
- In production environments, Cohesity recommends not replicating from one single node Cohesity cluster Virtual Edition to another single node Cohesity cluster Virtual Edition. Cohesity recommends replicating from Cohesity cluster Virtual Editions to Cohesity clusters running directly on hardware.
- If you have a Protection Group that is capturing and replicating Snapshots multiple times a day, Cohesity recommends configuring the replication schedule to copy

Snapshots daily instead of replicating Snapshots after each protection run. If the replication schedule is too frequent, the replication may lag behind the capturing of Snapshots resulting in a backlog of replication tasks.


- If Snapshots of a VM are replicated to a remote Cluster and the VM is renamed in the vCenter Server, the Cohesity UI on the remote Cluster displays the original VM name in the protection run Details page. However, you can search for new VM name while recovering or cloning and the search results displays the new VM name. Replication is not affected by this issue.

Access Management

Active Directory

Review the following considerations:

- Due to Windows client authentication cache behavior, after you add or remove a Cohesity cluster from an Active Directory domain, clients must log out and log in again to access the Cohesity cluster.
- The Cohesity cluster is added as one or more computer entities with no back-end RPC management API implementation.
- Users from trusted domains with trust type External cannot access Cohesity SMB shares.
- Active Directory lookup to external (non-transitive) trust via LDAP referral setup in AD is not supported.
- Active Directory lookup to a non-Windows-AD trust (Kerberos v5 Realms) is not supported.
- Consider the following trusted domains and forests.



If the cluster is joined to domain F1-Child1, then users from Forest2 or any of its child domains are not authenticated/allowed-access to the cluster. Users from all child domains within Forest1 can authenticate via NTLM.

If the cluster is joined to domain Forest1, then users from all child domains of Forest1 and users from the Forest2 domain only can access the cluster via NTLM. Users from child domains of Forest2 cannot access the cluster via NTLM.

Multitenancy

Review the following considerations:

Organizations (Tenants)

- If a VMware vCloud Director (vCD) source sub-object is assigned to a tenant, the recovery of VMs and vApps to an alternate location will fail in 6.2 release. When an entire vCD is registered within a tenant, then recovery to both original location and alternate location is supported.
- Enabling multitenancy for a cluster cannot not be undone. You cannot revert the cluster to a single tenancy state.
- If a single-tenant cluster is configured with remote access to a multitenant-enabled cluster, the Organizations page will not be available when accessing the multitenant cluster. The workaround is to enable multitenancy on the single tenancy cluster (it is not necessary to add any organizations.)

Hybrid Extender VM

Review the following considerations:

- Hybrid extender supports source registration and backup only for Windows and Linux physical sources. AIX, HPUX, Solaris physical sources are not supported with hybrid extender.
- Currently, Cohesity does not support the auto-upgrade of the Hybrid Extender. Therefore, you must upgrade the Hybrid Extender after upgrading the Cohesity cluster from one major release to another major release. For example, if you are upgrading the Cohesity cluster from 6.5.1 to 6.6, use the Hybrid Extender version provided with 6.6.
- When you're upgrading to maintenance releases such as 6.5.1e, you need not upgrade the Hybrid Extender. However, Cohesity recommends that the version of Cohesity cluster and the Hybrid Extender to be same.
- If a tenant deploys multiple Hybrid Extender VMs, SMB and NFS sessions do not failover to the next available Hybrid Extender VM. Cohesity depends on the hypervisor that is hosting the Hybrid Extender VM to ensure high availability. If the hypervisor does not support high availability, I/O requests fail.
- Hybrid Extender does not support the following features:
 - S3
 - SMB Multichannel
 - Keystone
 - Kerberos client for NFS
 - SSO
 - NFS authentication

Virtualization

If you are leveraging storage snapshots for VMware VM backups in releases prior to 7.2 and plan to upgrade the Cohesity cluster to release 7.2 or later, ensure that SAN connectivity between Cohesity nodes and the Storage Array (Pure, HPE Alletra, Nimble, and IBM) is in place before the upgrade to prevent VM backup failures.

Fixed Issues

The **Fixed Issues** page provides a list of issues fixed in the 7.4 release and its associated patch and update releases. Each fixed issue contains an issue ID and a brief description.

On the [Fixed Issues](#) page, select one of the following options to view the fixed issues:

- **Filter By Version**—Select a version to filter the fixed issues by a specific version.
- **Search By Issue ID**—Enter an issue ID to search for a specific fixed issue.
Example: ENG-225665 or 225665.

Security Fixes

Cohesity CVE patch releases utilize the Base OS patch within the software bundle to hold the CVE and related security fixes. BaseOS patch may contain critical CVE fixes, kernel updates, driver updates, and optionally bug fixes for other user-mode packages. Customers can review the fixes and determine if they want to skip a base OS patch and apply just software patches. All patches are cumulative if a patch is skipped and applied using a later patch release.

The following table lists the Common Vulnerabilities and Exposures (CVEs) fixed in the 7.4 release:

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-68285	Use-after-free condition in Linux kernel libceph map handling (CVSS not yet published by NIST or CNA)	N/A	N/A
CVE-2025-39981	Linux kernel Bluetooth MGMT use-after-free due to race in mgmt_pending handling	Critical	9.8
CVE-2025-68615	Buffer overflow in net-snmp snmptrapd via crafted SNMP trap packets	Critical	9.8

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-11561	SSSD AD integration flaw allows privilege impersonation	High	8.8
CVE-2025-14523	libsoup duplicate Host header handling enables request smuggling	High	8.2
CVE-2025-32988	GnuTLS double-free vulnerability in SAN otherName export leading to DoS or memory corruption	High	8.2
CVE-2025-32990	GnuTLS certtool template parsing heap buffer overflow leading to DoS	High	8.2
CVE-2025-5987	libssh ChaCha20 OpenSSL error handling flaw may cause undefined behavior	High	8.1
CVE-2025-11083	GNU Binutils heap-based buffer overflow in elf_swap_shdr	High	7.8
CVE-2025-23142	Linux kernel SCTP use-after-free in sendmsg	High	7.8
CVE-2025-37849	Linux kernel KVM arm64 vGIC cleanup flaw leading to use-after-free	High	7.8
CVE-2025-37891	Linux kernel ALSA UMP SysEx buffer overflow	High	7.8
CVE-2025-38724	Linux kernel nfsd race condition leading to use-after-free	High	7.8
CVE-2025-39864	Linux kernel cfg80211 use-after-free in wifi subsystem	High	7.8
CVE-2025-39955	Linux kernel TCP Fast Open socket state mishandling leading to retransmission errors	High	7.8
CVE-2025-39966	Linux kernel iommufd race condition causing use-after-free during aborted object creation	High	7.8
CVE-2025-40186	Linux kernel TCP subsystem race condition causing memory corruption	High	7.8
CVE-2025-40300	Linux kernel memory management flaw leading to potential use-after-free	High	7.8

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-6020	linux-pam pam_namespace local privilege escalation via symlink and race attacks	High	7.8
CVE-2025-8941	Privilege escalation via symlink and race condition flaws in linux-pam pam_namespace module	High	7.8
CVE-2025-39881	Use-after-free vulnerability in the Linux kernel kernfs subsystem during polling when an open file is released, which can lead to kernel crashes.	High	7.8
CVE-2025-59375	libexpat memory amplification leads to resource exhaustion and denial of service	High	7.5
CVE-2025-8194	Infinite loop in Python tarfile module when parsing archives with negative offsets	High	7.5
CVE-2025-9086	Out-of-bounds read in curl cookie path comparison allowing secure cookie overwrite	High	7.5
CVE-2025-39806	Linux kernel HID multitouch slab out-of-bounds access	High	7.1
CVE-2025-64720	Out-of-bounds read in libpng png_image_read_composite when processing palette images with optimized alpha handling	High	7.1
CVE-2025-65018	Heap buffer overflow in libpng png_image_finish_read when processing 16-bit interlaced PNG images	High	7.1
CVE-2025-66293	Out-of-bounds read in libpng simplified API when processing palette PNG images with transparency and gamma correction	High	7.1
CVE-2025-39883	Linux kernel memory-failure handling bug where unpoison_memory() accesses an uninitialized page, triggering VM_BUG_ON_PAGE and causing a kernel panic.	High	7.1
CVE-2025-39971	Linux kernel i40e driver index validation flaw causing potential memory corruption	High	7

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-39983	Linux kernel Bluetooth HCI event processing use-after-free vulnerability	High	7
CVE-2025-40047	Linux kernel io_uring waitid wait queue race condition	High	7
CVE-2025-40154	Linux kernel ASoC Intel bytcr_rt5640 driver invalid input mapping causing OOB access	High	7
CVE-2025-40176	Linux kernel TLS async decryption failure leading to use-after-free risk	High	7
CVE-2025-40778	Linux kernel driver validation error leading to memory safety issues	High	7
CVE-2025-40780	Linux kernel device handling flaw resulting in possible memory corruption	High	7
CVE-2025-68287	Race condition in Linux kernel dwc3 USB request handling leading to use-after-free	High	7
CVE-2024-37891	urllib3 Proxy-Authorization header not stripped on cross-origin redirects	Medium	6.5
CVE-2025-45582	Windows component information disclosure allowing spoofing over a network	Medium	6.5
CVE-2025-47906	Go os/exec LookPath may return unexpected executable paths when PATH contains executables	Medium	6.5
CVE-2025-48964	iputils ping denial of service via integer overflow from crafted ICMP Echo Reply	Medium	6.5
CVE-2025-6395	GnuTLS NULL pointer dereference in cipher suite handling causing denial of service	Medium	6.5
CVE-2024-36350	AMD processor transient execution vulnerability leading to potential privileged information leakage	Medium	5.6

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2024-36357	AMD processor L1D cache transient execution vulnerability allowing data leakage	Medium	5.6
CVE-2024-56690	Linux kernel pcrypt error handling can cause unnecessary panic	Medium	5.5
CVE-2025-13499	Wireshark Kafka dissector crash causing denial of service	Medium	5.5
CVE-2025-21795	Linux kernel NFSD shutdown callback hang	Medium	5.5
CVE-2025-39843	Linux kernel SLUB allocator lock recursion leading to deadlock	Medium	5.5
CVE-2025-39918	Linux kernel mt76 WiFi driver linked-list corruption due to improper cleanup	Medium	5.5
CVE-2025-39925	Linux kernel J1939 CAN protocol reference count handling flaw causing device unregister hang	Medium	5.5
CVE-2025-9714	Uncontrolled recursion in libxml2 XPath evaluation leading to stack overflow	Medium	5.5
CVE-2025-32989	Heap-buffer-overread in GnuTLS SCT extension parsing causing information disclosure	Medium	5.3
CVE-2025-68161	Missing TLS hostname verification in Apache Log4j Socket Appender	Medium	4.8
CVE-2025-39697	Linux kernel NFS race causing potential denial of service	Medium	4.7
CVE-2025-53905	Vim tar.vim path traversal allows arbitrary file overwrite via crafted tar archive	Medium	4.1
CVE-2025-53906	Vim zip.vim path traversal allows arbitrary file overwrite via crafted zip archive	Medium	4.1
CVE-2025-61984	OpenSSH username control characters may enable code execution with ProxyCommand	Low	3.6
CVE-2025-61985	OpenSSH ssh:// URI null byte injection may enable code execution with ProxyCommand	Low	3.6

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to [Cohesity Support](#), to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the [Cohesity Support Portal](#) to create a new case.
- Click the (?) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the [Cohesity Knowledge Base](#).
- Log in to the [Cohesity Support Portal](#) to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.

4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. [Click here](#) to send us your feedback!

Ensure that you provide the following details in your email:

- Document name
- Topic name
- Page number