COHESITY

Release Notes

Version 7.4.1

July 27, 2026

Published on July 27, 2026

Contents

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 7.4.1 provides new features and enhancements available for on-premises hardware, Cloud Edition, and Virtual Edition clusters. For more information, see What's New in 7.4.1?.

For more information on upgrading from previous releases to 7.4.1, see Upgrading to 7.4.1.

For more information on previous releases, see What's New in Earlier Releases?

What's New in 7.4.1?

The following new features and improvements are available in this release. For important information about upgrading from previous releases to 7.4.1, see Upgrading to 7.4.1.

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.

Controlled Availability (CA) Feature

A production-quality Cohesity product or feature made available to a limited set of customers. Contact your Cohesity account team to participate.

Data Protection

Databases

Source-Side Deduplication for MongoDB Physical Backups

Cohesity now supports source-side deduplication for MongoDB physical backups. This optimization reduces network bandwidth consumption, shortens the backup window, and optimizes storage consumption on the target, making it beneficial for large MongoDB deployments.

SAP HANA ET Log Backup Authentication Using API Keys

Cohesity now supports API key–based authentication for SAP HANA ET log backups.

With this enhancement, the API keys are securely stored on the SAP HANA host and used to authenticate ET log backup requests to the Cohesity cluster.

This feature improves security and operational reliability by eliminating frequent reauthentication due to password changes. Existing deployments remain backward compatible and continue to function using previously configured credentials if an API key is not configured.

API Support for SBT Snapshot View Allowlisting

Cohesity now provides API support to update the allowlist of SBT (System Backup to Tape) data view snapshots used for Oracle alternate restore operations. This enhancement removes the dependency on the `view_cli` command-line tool and enables programmatic management of snapshot allowlists.

File-Based Backup Support for Filestream Database

Cohesity now supports file-based backups for Microsoft SQL Server Filestream databases. Earlier, only volume-based and VDI-based backups were available for Filestream databases.

A Filestream database stores unstructured large objects in the Windows file system, controlled by SQL Server. This approach reduces database size and improves performance for large binary content compared with storing the files directly in the database tables.

PostgreSQL Database and WAL Encryption at Rest Support (EDB PostgreSQL)

Cohesity now supports backup and recovery of EnterpriseDB (EDB) PostgreSQL instances with Transparent Data Encryption (TDE) enabled, including encrypted database files and Write Ahead Log (WAL) files. Cohesity performs physical backups of the PostgreSQL data directory and preserves all encrypted files as is, without performing encryption or key management.

Edit or Remove Commit Log Archive Location for Cassandra Sources

Cohesity now allows you to edit or remove the commit log archive location for an existing Cassandra source when point-in-time recovery (PITR) is no longer required.

Physical Servers

Certificate Inventory API and Alerts for Physical Server Agents

Cohesity now supports enhancements to the existing Agents report API with additional certificate metadata for physical server agents, making it easier to audit and monitor agent certificates at scale. The 'Agent Certificate About to Expire' alert now includes additional certificate details to provide clearer context about the affected agent certificate.

Automatic Performance Tuning for Physical File Backups

Cohesity now supports automatic performance tuning for physical file-based protection, removing the need for manual configuration. At the start of each protection run, the cluster calculates and applies optimal resource allocation based on the capabilities of each host. This is applicable for Linux (x64) and AIX agents.

Scheduled Upgrades for Physical Agents Controlled Availability

Cohesity now supports scheduled upgrades for physical agents. You can configure upgrade schedules to automatically update agents to the latest version, reducing manual maintenance and ensuring agents stay current.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

NAS

REST API Support for NetApp ONTAP Controlled Availability

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).

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Automatic Pause and Resume for Isilon SMB Jobs on Authentication Errors

Cohesity now supports automatic pausing of Isilon SMB backup jobs when authentication errors occur due to password changes and resume after you update credentials on the Cohesity cluster. This prevents backup failures; any missed files are backed up in the next incremental run.

Support for Nutanix Files NFSv41 Shares

Cohesity now supports Nutanix Files NFSv4.1 shares. Cohesity also supports Distributed File System (DFS) shares over the NFSv4.1 protocol.

Microsoft 365

Recovery of SharePoint Wiki Page Controlled Availability

Cohesity now supports improved recovery coverage for SharePoint Wiki Pages as part of the SharePoint recovery workflow, including full, alternate-location, and granular (item-level) recovery.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Microsoft 365 Mailboxes Recovery to On-premises Exchange Server Controlled Availability

Cohesity now supports specifying an alternate SMTP address when recovering Exchange Online mailboxes to an on-premises Exchange Server. This provides flexibility in recovery scenarios where the target mailbox differs from the source.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Backup and Recovery for Microsoft Teams Shared Channels Controlled Availability

Cohesity now supports backup and recovery of Teams Shared Channels as part of the Teams backup and recovery workflow. This improves Teams protection coverage by ensuring that shared channel conversations and content are protected.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Backup and Recovery of Microsoft Teams Tabs Controlled Availability

Cohesity now supports backup and recovery of Teams tabs as part of the Teams backup and recovery workflow. This improves Teams protection coverage and ensures that Teams tab configurations are protected.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Granular Recovery for All Microsoft 365 Mailbox Folder Roots

Cohesity now supports granular recovery of items in Recoverable Items and Archive Recoverable Items folders. Additionally, all four Microsoft 365 mailbox folder roots, Primary Mailbox, In-Place Archive, Recoverable Items, and Archive Recoverable Items are now displayed in the granular recovery workflow, allowing you to select any root individually to browse and recover its items.

Teams Content Download

Cohesity now supports downloading Teams content, improving data accessibility and enabling you to access content during Microsoft 365 service outages or in the event of data loss or corruption.

- Download 1:1 Chats and Group Chats
- Download Channels Posts

Backup and Recovery of OneNote Files

Cohesity now supports backup and recovery of OneNote files across OneDrive, SharePoint, Teams, and Groups as part of their backup and recovery workflows, improving coverage and ensuring protection of OneNote files across different workloads.

Backup and Recovery of Microsoft 365 Security Groups Controlled Availability

Cohesity now supports protecting users using Microsoft Entra ID security groups during the backup workflow of Mailbox and OneDrive. Selecting a security group automatically includes all members (including nested members) in the protection group, eliminating the need for manual user selection. This results in faster onboarding and an improved user experience.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Backup and Recovery of Subsites within Microsoft 365 Teams and Groups Controlled Availability

Cohesity now supports backup and recovery of subsites in Teams and Groups as part of their backup and recovery workflows. This improves Teams and Groups protection coverage by ensuring that subsites and their associated content are protected.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Backup and Recovery for SharePoint Site Pages Controlled Availability

Cohesity now supports backup and recovery of SharePoint Online site pages as part of the SharePoint backup and recovery workflow. This improves SharePoint protection coverage by ensuring that site pages are protected.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Transition to Microsoft Graph APIs for Exchange Online Mailbox Protection Early Access

Cohesity now supports Exchange Online mailbox backup and recovery using Microsoft Graph APIs, replacing Exchange Web Services (EWS) APIs. This transition ensures seamless and uninterrupted mailbox protection while providing greater flexibility and compatibility. Existing backups remain fully usable, with no need for data re-ingestion. The EWS to Graph transition affects supports only user and shared primary mailboxes in this release. Teams mailboxes, Groups mailboxes and user/shared mailboxes that protect the in-place archive continue to use EWS.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Virtualization

Support for Segmented Data Services IP Selection for iSCSI Transport in AHV Sources

You can now configure a segmented Data Services IP address for iSCSI transport for AHV sources at both the Prism Element (PE) and Prism Central (PC) levels. This enhancement provides greater flexibility in separating management and transport networks for backup and recovery operations.

Support for Existing VM Handling Options During Recovery for Nutanix AHV Controlled Availability

Cohesity now supports additional options to control how existing VMs are handled when recovering to the original location in AHV environments. These options are available for both **Copy Recovery** and **Instant Recovery** workflows.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Support for Source-Side Deduplication for VMs and PVCs in Red Hat OpenShift Virtualization Controlled Availability

You can now enable source-side deduplication for VMs and PVCs to reduce data transfer and improve backup efficiency. By eliminating duplicate data at the source, only unique data blocks are sent to the Cohesity cluster.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Red Hat OpenShift Virtualization VM File-Level Recovery (FLR)

You can now perform File-Level Recovery (FLR) for Red Hat OpenShift Virtualization VMs using an agentless approach through the QEMU Guest Agent. This enhancement enables recovery of files and folders directly within the guest, improving flexibility and simplifying recovery operations.

Instant Recovery for Red Hat OpenShift Virtualization

Cohesity now supports Instant Recovery to both original and alternate locations for Red Hat OpenShift Virtualization, enabling VMs to be immediately available in the target location after a recovery operation.

NFSv4.1 Protocol Support for VMware Recovery Operations

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

Data Protection for VMware Cloud Foundation Automation Early Access

Cohesity now supports data protection for VMware Cloud Foundation (VCF) Automation environments. This capability enables you to protect VMs managed by VCF Automation, bringing Cohesity data protection to the latest VCF platform.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Support for VM Backups on Azure NetApp Files and Pure Cloud Block Store

Cohesity now supports backup of VMs hosted on Azure NetApp Files and Pure Cloud Block Store in Azure VMWare Solution (AVS).

Kubernetes

Kubernetes Pod Labels, Annotations, and Priority Class Support

Cohesity now supports specifying priority class names for pods deployed by Cohesity in Kubernetes environments (including Velero, data mover, and temporary pods). You can also add custom labels and annotations to Cohesity-managed pods and services, enabling better integration with cluster policies and operational workflows

Kubernetes Diagnostics Collection

Cohesity now enables you to collect Kubernetes cluster diagnostics directly from the UI. This feature allows you to gather logs and resource details across cluster-scoped and namespace-scoped objects, including Velero and Datamover logs, to simplify troubleshooting and accelerate issue resolution.

Skip Persistent Volume Claim (PVC) Data During Kubernetes Backup and Recovery

You can now selectively skip Persistent Volume Claim (PVC) data during Kubernetes backup and recovery operations while preserving PVC metadata. This enhancement provides granular control at the PVC level, allowing you to choose whether to back up and restore full PVC data or only the PVC resource metadata.

Enhanced Recovery Options for Kubernetes Namespaces

Cohesity now provides enhanced recovery options for Kubernetes namespaces in both original and alternate locations. You can rename individual objects during recovery and manage conflicts when objects with the same name already exist in the target cluster. Options include overwriting existing objects or deleting and recreating them, with all actions recorded in audit logs.

Data Cloud for NetBackup

Flexible Cluster Deployment Options

Data Cloud for NetBackup now supports configuring clusters with or without media servers, including an external media-only topology. This reduces Day-1 setup complexity and simplifies deployment workflows.

Rolling Upgrade Support

Data Cloud for NetBackup supports rolling upgrades from 7.4 to 7.4.1 with minimal disruption. Backup jobs continue during upgrades with failover-based resiliency.

FIPS 140-3 Validation (Media Layer)

Data Cloud for NetBackup supports FIPS 140-3–related validation for the media layer, improving security readiness for regulated environments.

Cohesity Cloud Edition

Cohesity Cloud Edition on AWS

Support for Amazon RDS for Microsoft SQL Server Data Protection

Cohesity now supports the backup and recovery of Amazon RDS for Microsoft SQL Server at the database level by ingesting data into Cohesity Cloud Edition on AWS.

Cumulative Incremental Backup Support for Amazon RDS for Microsoft SQL Server Controlled Availability

Cohesity now supports cumulative incremental backups for Amazon RDS for Microsoft SQL Server databases. Cumulative incremental backups capture all changes since the most recent full backup and reduce backup duration and data transfer compared to full backups.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Backup and Recovery for TDE-Enabled Amazon RDS for Microsoft SQL Server Controlled Availability

Cohesity now supports the backup and recovery of Transparent Data Encryption (TDE)–enabled Amazon RDS for Microsoft SQL Server databases.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Support for Incremental Backup of Amazon RDS and Aurora for PostgreSQL

Cohesity now supports incremental (differential) backups for Amazon RDS for PostgreSQL and Amazon Aurora PostgreSQL databases. Incremental backups capture only the data changes since the last successful backup, helping reduce backup windows, optimize storage consumption, and improve backup efficiency.

Support for Materialized View Refresh After Amazon RDS and Aurora for PostgreSQL Recovery

Cohesity now supports refreshing PostgreSQL materialized views after Amazon RDS for PostgreSQL and Amazon Aurora PostgreSQL database recovery. You can choose to automatically refresh materialized views during recovery or perform the refresh manually after recovery using a downloadable refresh script from the Cohesity UI.

Register AWS Organization

Cohesity now supports registering an AWS Organization as a single source for protecting the following AWS workloads:

- Amazon Simple Storage Service Bucket
- Amazon Relational Database Service for MySQL

- Amazon Aurora MySQL
- Amazon Aurora PostgreSQL
- Amazon Relational Database Service for PostgreSQL
- Amazon Relational Database Service for Oracle
- Amazon Relational Database Service for Microsoft SQL Server
- Amazon DocumentDB
- Amazon DynamoDB
- Amazon Redshift Cluster

This enables centralized management and protection of workloads across multiple AWS accounts from a unified source registration.

Amazon EBS Volume Protection Support

Cohesity now supports Amazon EBS volumes as a protected workload type on Cohesity Cloud Edition for AWS. In addition to EBS volumes protected through EC2 instance backups, you can now directly protect and recover both standalone EBS volumes and EBS volumes attached to EC2 instances.

Support for Amazon RDS for PostgreSQL Data Protection

Cohesity now supports the backup and recovery of Amazon RDS for PostgreSQL at the database level by ingesting data into Cohesity Cloud Edition on AWS.

Support for Amazon Aurora PostgreSQL Data Protection

Cohesity now supports the backup and recovery of Amazon Aurora PostgreSQL at the database level by ingesting data into Cohesity Cloud Edition on AWS.

Point-in-Time Recovery Support

You can now perform point-in-time recovery (PITR) for the following workloads backed up using snapshot-based protection in Cohesity Cloud Edition on AWS:

- Amazon Aurora PostgreSQL
- Amazon Aurora MySQL
- Amazon Relational Database Service for MySQL
- Amazon Relational Database Service for PostgreSQL
- Amazon Relational Database Service for Microsoft SQL Server

Local Secondary Index (LSI) Restore for Amazon DynamoDB Controlled Availability

Cohesity now supports the recovery of Local Secondary Indexes (LSIs) when recovering Amazon DynamoDB tables from snapshots.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Down-tier Backup Data to AWS Cold and Archive Storage

Cohesity Cloud Edition on AWS now supports backing up data directly to archive and cold cloud storage tiers as the primary copy. This reduces long-term retention costs by writing backup data to cheaper storage tiers while keeping the metadata required for incremental backups in a hot or standard tier automatically.

Support for Amazon Redshift Serverless Data Protection Controlled Availability

Cohesity now supports the protection of Amazon Redshift Serverless at database-level by ingesting into Cohesity Cloud Edition on AWS.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Cohesity Cloud Edition on Azure

Support for Azure PostgreSQL Databases Protection

Cohesity now supports the protection of Azure PostgreSQL databases at database-level by ingesting data into Cohesity Cloud Edition on Azure.

Support for Incremental Backup of Azure PostgreSQL

Cohesity now supports incremental backups for Azure PostgreSQL databases. Incremental backups capture only the data changes since the last successful backup, helping reduce backup windows, optimize storage consumption, and improve backup efficiency.

Also, Cohesity supports refreshing the PostgreSQL materialized views after database recovery. You can choose to automatically refresh materialized views during recovery or perform the refresh manually after recovery using a downloadable refresh script from the Cohesity UI.

Support for Azure Cosmos DB – MongoDB Data Protection

Cohesity now supports the protection of Cosmos DB – MongoDB databases at database level by ingesting data into Cohesity Cloud Edition on Azure.

Down-tier Backup Data to Azure Cold and Archive Storage

Cohesity Cloud Edition on Azure now supports backing up data directly to archive and cold cloud storage tiers as the primary copy. This reduces long-term retention costs by writing backup data to cheaper storage tiers while keeping the metadata required for incremental backups in a hot or standard tier automatically.

Backup and Recover Azure Blob and Data Lake Storages in Azure Government Cloud Early Access

You can now back up and recover Azure Blob Storage and Azure Data Lake Storage workloads hosted in Azure Government Cloud using Cohesity Cloud Edition on Azure.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Register Azure Source Using Managed Identity Authentication Controlled Availability

Cohesity now supports registering Azure sources in Cohesity Cloud Edition on Azure using User Managed Identity and System Managed Identity to access Azure APIs for protection services including backup and recovery

You can protect the following Azure workloads using managed identity–based authentication:

- Azure MySQL
- Azure Cosmos DB - Cassandra
- Azure Cosmos DB - MongoDB
- Azure Cosmos DB - NoSQL
- Azure Cosmos DB - Table
- Azure SQL Database
- Azure SQL Managed Instance
- Azure Table Storage
- Azure PostgreSQL
- Azure Blob Storage and Azure Data Lake Storage

This feature is not supported for Azure VM.

> **Note:** This is a Controlled Availability feature. Contact your Cohesity account team to enable the feature.

Support for Azure Cosmos DB – NoSQL Data Protection

Cohesity now supports the protection of Cosmos DB – NoSQL databases at database-level by ingesting data into Cohesity Cloud Edition on Azure.

Cohesity Cloud Edition on GCP

Support for Spanner Data Protection

You can now perform the backup and recovery of Spanner databases using DataProtect for cloud in GCP.

Support for Cloud SQL for PostgreSQL Data Protection

You can now perform database-level backup and recovery of Cloud SQL for PostgreSQL instances using DataProtect for cloud in GCP.

Support for Cloud SQL for SQL Server Data Protection

You can now perform database-level backup and recovery of Cloud SQL for SQL Server using DataProtect for cloud in GCP.

Support for Cloud SQL for AlloyDB PostgreSQL Data Protection

You can now perform the backup and recovery of Cloud SQL for AlloyDB PostgreSQL databases using DataProtect for cloud in GCP.

Support for Firestore Data Protection

You can now perform the backup and recovery of Firestore databases using DataProtect for cloud in GCP.

BigQuery Table-Level Backup and Recovery Early Access

Cohesity now supports table-level backup and recovery for Google Cloud BigQuery, allowing you to protect and recover individual tables instead of entire datasets. This enables more granular data protection, faster recovery, and reduced backup scope for large datasets.

Table-level protection supports full backups only. Incremental backup and recovery are not supported.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Virtual Edition

Virtual Edition (VE) Qualified on Red Hat OpenShift

Cohesity Virtual Edition (VE) is now qualified to run on Red Hat OpenShift. This enables customers migrating from VMware to OpenShift to continue using Cohesity VE for data protection of their OpenShift workloads and underlying virtualized infrastructure.

SmartFiles

Indexing Support for S3 Views

Cohesity now supports indexing for S3 Views, which allows search and browse across S3 View snapshots. With this enhancement, object level recovery is also supported for S3 views.

Remote Cluster Registration with Active Directory Authentication Early Access

Cohesity now allows you to register a remote cluster using either an Active Directory user or a local user on the remote cluster. To use Active Directory authentication, you can provide the domain details during remote cluster registration on the Cohesity UI. Earlier, only local user authentication was supported.

> **Note:** This is an Early Access feature. Contact your Cohesity account team to enable the feature.

Support for S3 Server Logs Predefined Template

Cohesity now allows you to designate one S3 View as the logging bucket for another S3 View using the S3 Server Logs predefined template in the Cohesity cluster. S3 server logs can be configured end-to-end through the Cohesity UI, eliminating the dependency on the AWS CLI. Using this dedicated template ensures proper defaults are applied, and logs flow automatically from the source view to the designated target bucket.

Multiple Chunks Transfer Support in Cohesity S3

Cohesity S3 now supports transferring payload in multiple chunks (Chunked Upload). This is required by default for the latest versions of S3 SDKs.

Cohesity S3 also supports the checksum algorithms, CRC32, CRC32C, CRC64NVME, SHA1, and SHA256.

Security

SNMPv3 AES-256 and SHA-2 Authentication Support

Cohesity clusters now support AES-256 encryption and SHA-2 (SHA-256, SHA-384, SHA-512) authentication protocols for SNMPv3 trap destinations and polling. This enables customers to meet FIPS 140-3 compliance requirements for SNMP monitoring without relying on deprecated algorithms (AES-128/MD5/SHA-1).

FIPS-Certified OpenSSL 3 for Physical Agents (AIX)

Cohesity AIX 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.

Certificate Inventory API and Alerts

Cohesity now supports a deduplicated, cluster-wide view of all certificates across nodes, including key metadata such as consumers, expiry status, and auto-renewal, through the Certificate Inventory API, and supports configuring certificate expiry alert thresholds using the CLI and API to enable proactive certificate monitoring and management.

FortKnox Self-Managed

Enable or Disable Physical Air-Gap

Vault administrators can now control the Physical Air-Gap setting directly from the Paired Clusters page using a toggle. When enabled, recovery or replication operations that use the interface group configured for vaulting are not permitted outside the vaulting window. When disabled, the vaulting interface group remains available outside the vaulting window, allowing recovery or replication operations from the vault cluster to proceed.

Flexible Vault Retention Configuration

You can now set the vault retention period to any value in a protection policy, including a value lower than the primary backup retention, supporting use cases that require a shorter retention for vaulted copies.

Enter Multiple Node IPs During Cluster Pairing

You can now type or paste multiple IP addresses as a comma-separated list into the VIP or Node IP Addresses field during cluster pairing, eliminating the need to enter each address individually.

Prevent Duplicate Cluster Pairing

FortKnox Self-Managed now blocks pairing a cluster that is already registered as a remote or vault-paired cluster and displays a clear error message, preventing accidental duplicate registrations.

Paired Clusters Excluded from Replication Targets

When creating a protection policy on the vault cluster, vault-paired primary clusters no longer appear in the replication target dropdown, preventing unintended selection and reducing misconfiguration risk.

Prevent Duplicate Targets in Edit Run

In the **Edit Run** dialog, replication and vault targets already associated with a run are now grayed out, preventing the same target from being added more than once to a single run.

Customer-Managed CA Integration (Venafi)

You can now integrate your Cohesity cluster with Venafi Trust Protection Platform (TPP) as a customer-managed CA. You can configure this integration using the Cohesity CLI to enable policy-based certificate issuance and automatic distribution to agents, eliminating manual certificate installation.

Cohesity Feature Deprecation

The following features are deprecated from Cohesity 7.4.1 release. You can review the following table to check whether you are currently using any features that are deprecated or

will be deprecated in the future.

Deprecated Features

Deprecated Feature Name	Description	Last Supported Cohesity Version	Alternative Solutions
Support for CentOS 6, RHEL 6 and OEL 6 as physical server platforms	The Linux physical agent will not support CentOS 6, RHEL 6 and OEL 6 platforms (including minor versions).	7.4.1	Continue to use the 7.4.1 agent or backup data using NFS shares and Cohesity's Generic NAS connector.

Upgrading to 7.4.1

Upgrade Paths

You can upgrade your Cohesity cluster from previous releases to 7.4.1. The following table provides details on supported upgrade paths.

Your Current Release	Upgrade Path to 7.4.1
• 7.4 • 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	7.4.1 directly

> **Note:** Direct upgrades from 6.8.x to 7.4.1 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 7.4.1.

Release Upgrade Policy

Policy	Example
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.	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.
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.	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. 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.
Cohesity will support the release N-1 upgrade without an intermediate step. (N is defined as the current release branch).	7.2.x to 7.4.1 is supported. 7.4.1 is the current release branch.
Cohesity will support the release N-2 upgrade without an intermediate step. (N is defined as the current release branch).	7.1.x to 7.4.1 is supported. 7.4.1 is the current release branch.
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.	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.

Upgrade Considerations

Note the following about upgrading the Cohesity cluster to 7.4.1:

- 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 7.4.1 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 7.4.1, 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 7.4.1 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.
- Starting with Cohesity version 7.2.2, the Cohesity BaseOS transitioned from RHEL 7.x to RHEL 9.x, and the time synchronization service changed from ntpd to chronyd. Due to this architectural change, Windows-based NTP sources (Windows Time Service / W32Time) are not supported for Cohesity clusters running version 7.2.2 or later. Configure Linux-based NTP sources for time synchronization before upgrading to ensure cluster stability and supportability.

Databases

- 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 7.4.1 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 7.4.1 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 7.4.1 version, if you are replicating the Mailbox data to a remote Cohesity cluster, then ensure that you upgrade the remote Cohesity cluster to the 7.4.1 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.

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**.

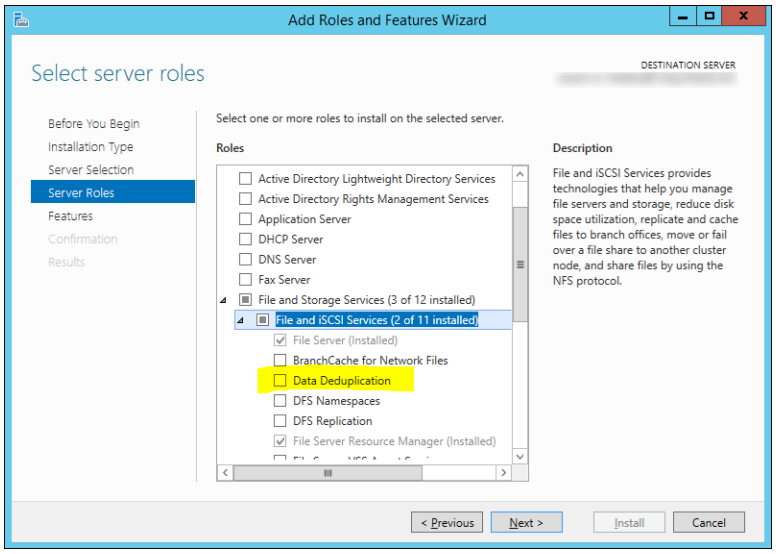

- Instant Volume Mount (IVM) restore of ReFS volumes backed up using Windows physical block-based jobs cannot be restored to an alternate Windows server running a lower version of ReFS.
- When mounting volumes on a Linux physical Server, the loop devices present on the Server are used for mounting. Therefore, 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.
- **Download** – Download is not supported for S3 workflows including indexed browsing, non-indexed browsing, search, and recovery.
- **Recover**
  - Recovery of archival S3 objects is not supported.
  - Prefix-based recovery is not allowed for S3 objects.
  - S3 object recovery is allowed only across S3 Views.
  - Recovery from S3 2.0 View to S3 1.0 View is not supported.
  - Only object-based recovery is supported; prefix-based recovery is not supported.
- **S3 Indexing**

  The following are the limitations for S3 indexed browse. Since S3 object storage is inherently non-hierarchical (flat key space), but the browse displays a folder-style hierarchy, mapping between these models can introduce ambiguous scenarios.

  > **Note:** These limitations do not apply to non-indexed browse.

  - Ambiguity occurs between the S3 keys in the hierarchical browse View and native S3 tools in the following cases:

  - **Multiple slashes** – If a S3 key contains continuous multiple slashes, the folder View can appear different from the original key in other S3 tools.

    For example, */foo///xyz/pqr*
  - **File and folder reuse** – If the same S3 path is used both as a file and as a folder, the browse View can display only one version of the path.

    For example, */foo/pqr and /foo/pqr/abc*
  - **Leading slashes** – If one S3 key starts with a slash and another does not, they can still appear as the same path in the browse View.

    For example, *foo/pqr and /foo/pqr*
- For versioned buckets, only the latest version of an object corresponding to the snapshot is considered for indexed browsing, non-indexed browsing, search, and recovery workflows.
- Indexed browse does not support sorting, so results may appear in an order different than non-indexed browse.
- Search in indexed browse displays exact results only when the entire object name is specified.
- Search displays full object names and not prefixes.
- Search up to 10k objects is supported based on the maximum value selected. To display more results, you must refine the search query.

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 `lvol_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	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`UUID="ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc" /mnt ext4 auto 0`
		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 `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Virtual Machine	LVM Volume	The Cohesity cluster detects mount points for entries in the `/etc/fstab` file with the following formats:
		`UUID=ccd1d599-e68e-4b88-ba9b-6f75b63f1bdc /mnt ext4 auto 0`
		`/dev/mapper/VG1-root /mnt ext4 defaults 1 1`
		`/dev/VG1/root /mnt ext4 defaults 1 1`
		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 `/etc/fstab` file. For these example entries, files and directories are indexed with the `/mnt` mount path, such as `/mnt/example/test.txt`.
		If the Cohesity cluster cannot detect a mount point, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

Linux Physical	LVM Volume	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 `/mnt/example/test.txt`. If the volume is not mounted, the Cohesity cluster indexes the files and directories into a `lvol_N` directory. For example, the `/mnt/example/test.txt` file is indexed as `/lvol_1/example/test.txt`.

- 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-vmtoolsd.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 userAdd 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.

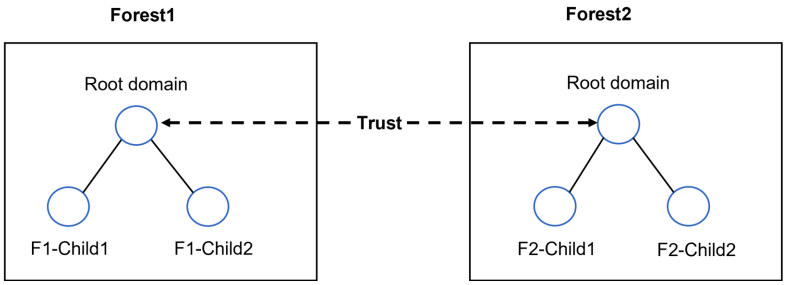

If the cluster is joined to domain F1-Child1, then users from Forest2 or any of its child domains are not authenticated/allowed-access to the cluster. Users from all child domains within Forest1 can authenticate via NTLM.

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.

Marketplace Apps

Review the following considerations:

- Dual-stack network configuration for Marketplace Apps is not supported.
- Running Apps on a Cohesity cluster with DoD mode enabled is not supported.

Fixed Issues

The **Fixed Issues** page provides a list of issues fixed in the 7.4.1 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 7.4.1 release:

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-6100	Python lzma.LZMADecompressor, bz2.BZ2Decompressor, and gzip.GzipFile use-after-free when a MemoryError occurs during decompression and the instance is reused.	High	8.1
CVE-2026-5201	gdk-pixbuf JPEG image loader heap-based buffer overflow due to improper validation of color component counts in a crafted JPEG; allows remote denial of service.	High	7.5
CVE-2026-5121	libarchive integer overflow in zisofs block pointer allocation on 32-bit systems; crafted ISO9660 image may cause heap buffer overflow and arbitrary code execution.	High	7.5
CVE-2026-4878	libcap cap_set_file() TOCTOU race condition allows a local attacker with write access to a parent directory to redirect file capability updates to an attacker-controlled file.	Medium	6.7
CVE-2026-4786	Incomplete fix for CVE-2026-4519: Python webbrowser.open() still allows command injection via URLs containing '%action', bypassing the prior mitigation for certain browser types.	High	7.1
CVE-2026-4775	libtiff putcontig8bitYCbCr44tile() signed integer overflow leading to out-of-bounds heap write via a crafted TIFF file, potentially causing denial of service or arbitrary code execution.	High	7.8
CVE-2026-46333	Linux kernel ptrace get_dumpable() logic allows processes without an MM pointer to bypass capability checks, enabling uid-0 users to access ptrace details without CAP_SYS_PTRACE.	High	7.8
CVE-2026-46300	Linux kernel net/skbuff skb_try_coalesce() loses SKBFL_SHARED_FRAG marker when transferring paged frags, allowing ESP input to decrypt in-place over page-cache-backed frags (local privilege escalation).	High	7.8
CVE-2026-4519	Python webbrowser.open() accepts URLs with leading dashes, which certain browsers interpret as command-line options, allowing command injection.	High	7.1

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-44673	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.5
CVE-2026-4424	libarchive heap out-of-bounds read in RAR archive processing due to improper LZSS sliding window size validation after compression method transitions.	High	7.5
CVE-2026-43284	Linux kernel xfrm/ESP does not set SKBFL_SHARED_FRAG on IPv4/IPv6 UDP splice frags, allowing ESP input to decrypt in-place over externally-owned shared frag data (local privilege escalation).	High	7.8
CVE-2026-43077	Linux kernel crypto: algif_aead - Fix minimum RX size check for decryption	High	7.1
CVE-2026-4111	Flaw in RAR5 archive decompression logic of libarchive library in archive_read_data() processing path	High	7.5
CVE-2026-41035	rsync 3.0.1-3.4.1 receive_xattr() uses an untrusted length in a qsort call when --xattrs is used, leading to a use-after-free on the receiver side.	High	7.4
CVE-2026-40164	jq uses a hardcoded MurmurHash3 seed allowing attackers to precompute key collisions via a crafted ~100 KB JSON object and cause O(n²) CPU exhaustion (DoS).	High	7.5
CVE-2026-39979	jq libjq jv_parse_sized() performs an out-of-bounds read in its error-handling path via %s on a non-NUL-terminated buffer, potentially causing memory disclosure or process termination.	High	8.2
CVE-2026-35535	Sudo through 1.9.17p2 does not treat failures of setuid/setgid/setgroups during privilege drop before running the mailer as fatal, allowing local privilege escalation.	High	7.4

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-35414	OpenSSH before 10.3 mishandles authorized_keys principals option when combined with a Certificate Authority using comma characters, potentially allowing unauthorized access.	Medium	4.8
CVE-2026-35388	OpenSSH before 10.3 omits connection multiplexing confirmation for proxy-mode multiplexing sessions, allowing potential integrity bypass.	Low	2.2
CVE-2026-35387	OpenSSH before 10.3 misinterprets PubkeyAcceptedAlgorithms/HostbasedAcceptedAlgorithms entries, enabling unintended ECDSA algorithms that weaken authentication policy.	Low	3.1
CVE-2026-35386	OpenSSH before 10.3 allows command execution via shell metacharacters in a username when combined with non-default % ssh_config settings.	Low	3.6
CVE-2026-35385	OpenSSH before 10.3 scp may install downloaded files as setuid/setgid when run as root with -O (legacy protocol) without -p, allowing unintended privilege escalation.	High	7.5
CVE-2026-34986	Go JOSE (go-jose) panic in JWE decryption when alg indicates key wrapping but encrypted_key field is empty; uncaught exception causes denial of service.	High	7.5
CVE-2026-3497	OpenSSH GSSAPI patch uses sshpkt_disconnect() on error without terminating the process; attacker can send unexpected GSSAPI message during key exchange causing access to uninitialized connection variables.	High	8.2
CVE-2026-34268	Oracle Java SE / GraalVM Security component information disclosure; a difficult-to-exploit local unauthenticated attacker can read a subset of accessible data.	Low	2.9
CVE-2026-34003	X.Org X server out-of-bounds memory access in XKB key types request validation; specially crafted request leads to information disclosure or server crash.	High	7.8

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-34002	X.Org X server out-of-bounds read in XKB modifier map handling; malformed request causes the server to read beyond intended memory, potentially disclosing data.	Medium	6.1
CVE-2026-34001	X.Org X server use-after-free in XSYNC fence triggering logic (miSyncTriggerFence()); allows server crash and potential memory corruption.	High	7.8
CVE-2026-34000	X.Org X server out-of-bounds read in XKB geometry processing (CheckSetGeom/XkbAddGeomKeyAlias); allows memory disclosure or server crash.	Medium	6.1
CVE-2026-33999	X.Org X server integer underflow in XKB compatibility map handling; allows local or remote attacker to trigger buffer read overrun, causing denial of service.	High	7.8
CVE-2026-33636	libpng (1.6.36–1.6.55) out-of-bounds read and write in ARM/AArch64 Neon-optimized palette expansion path when expanding 8-bit paletted rows.	High	7.6
CVE-2026-33416	libpng use-after-free; png_set_tRNS and png_set_PLTE alias heap buffers with independent lifetimes, leaving dangling pointer allowing writes to freed memory.	High	7.5
CVE-2026-33412	Vim command injection vulnerability exists in glob() function on Unix-like systems (prior to 9.2.0202)	High	7.3
CVE-2026-33186	gRPC-Go (< 1.79.3) authorization bypass via improper validation of HTTP/2 :path pseudo-header; requests missing a leading slash bypass path-based deny rules.	Critical	9.1
CVE-2026-32283	Go standard library crypto/tls denial-of-service in TLS 1.3; multiple key update messages in a single record can deadlock the connection.	High	7.5

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-32280	Go standard library crypto/x509 denial-of-service; unbounded work during certificate chain building with large numbers of intermediate certificates can exhaust resources.	High	7.5
CVE-2026-31532	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.1
CVE-2026-31431	Linux kernel crypto: algif_aead - Revert to operating out-of-place	High	7.8
CVE-2026-31402	Linux kernel nfsd: fix heap overflow in NFSv4.0 LOCK replay cache	High	7.8
CVE-2026-3085	GStreamer rtpqdm2depay heap-based buffer overflow in X-QDM RTP payload processing; remote attackers may execute arbitrary code via crafted RTP packets.	High	8.8
CVE-2026-3083	GStreamer rtpqdm2depay out-of-bounds write in X-QDM RTP payload processing; remote attackers may execute arbitrary code via crafted RTP packets.	High	8.8
CVE-2026-3082	GStreamer JPEG Parser heap-based buffer overflow in Huffman table processing; remote attackers may execute arbitrary code via a crafted JPEG.	High	7.8
CVE-2026-2923	GStreamer DVB Subtitles out-of-bounds write in coordinate handling; remote attackers may execute arbitrary code via crafted subtitles.	High	7.8
CVE-2026-2922	GStreamer RealMedia Demuxer out-of-bounds write in video packet processing; remote attackers may execute arbitrary code via a crafted RealMedia file.	High	7.8

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-2921	GStreamer RIFF palette integer overflow in AVI file palette data handling; remote attackers may execute arbitrary code via a crafted AVI file.	High	7.8
CVE-2026-2920	GStreamer ASF Demuxer heap-based buffer overflow in stream header processing; remote attackers may execute arbitrary code via a crafted media file.	High	7.8
CVE-2026-28871	Apple Safari/iOS/macOS logic issue; visiting a maliciously crafted website may lead to a cross-site scripting attack.	Medium	4.3
CVE-2026-28859	Apple Safari/WebKit use-after-free; a malicious website may be able to process restricted web content outside the sandbox.	High	8.8
CVE-2026-28857	Apple Safari/WebKit out-of-bounds read/write or use-after-free; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2026-28421	Vim heap-buffer-overflow and segmentation fault in swap file recovery logic (prior to 9.2.0077)	Medium	5.3
CVE-2026-28417	Vim OS command injection vulnerability in netrw standard plugin (prior to 9.2.0073)	Medium	4.4
CVE-2026-27135	nghttp2 (HTTP/2 C library) missing internal state validation after session termination; a malformed frame triggers an assertion failure (denial-of-service).	High	7.5
CVE-2026-25749	Vim heap buffer overflow in tag file resolution logic when processing 'helpfile' option (prior to 9.1.2132)	High	7.3

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-25679	Go standard library net/url (url.Parse) insufficient validation of IPv6 host literals; accepts invalid URLs, potentially causing denial-of-service or security bypass.	High	7.5
CVE-2026-25646	libpng vulnerability in PNG raster image file processing	High	7
CVE-2026-23868	giflib double-free vulnerability due to shallow copy in GifMakeSavedImage and incorrect error handling; may cause denial-of-service.	High	7
CVE-2026-23865	FreeType integer overflow in tt_var_load_item_variation_store; parsing HVAR/VVAR/MVAR tables in OpenType variable fonts may lead to an out-of-bounds read.	Medium	5.3
CVE-2026-23490	pyasn1 (Python ASN.1 library) prior to 0.6.2 memory exhaustion via malformed RELATIVE-OID with excessive continuation octets; allows remote denial-of-service.	High	7.5
CVE-2026-23401	Linux kernel KVM x86/mmu: drop/zap existing present SPTE even when creating an MMIO SPTE	High	8.1
CVE-2026-23243	Linux kernel RDMA/umad: Reject negative data_len in ib_umad_write	High	7.3
CVE-2026-23231	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.5
CVE-2026-23209	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-23204	Linux kernel net/sched: cls_u32: use skb_header_pointer_careful() to prevent out-of-bounds read	High	7.3
CVE-2026-23193	Linux kernel scsi target iscsi: Fix use-after-free in iscsit_dec_session_usage_count()	High	7.1
CVE-2026-23171	Linux kernel bonding: fix use-after-free due to enslave fail after slave array update	High	7
CVE-2026-23144	Linux kernel mm/damon/sysfs: cleanup attrs subdirs on context dir setup failure	High	7.3
CVE-2026-23111	Linux kernel netfilter nf_tables: fix inverted genmask check in nft_map_catchall_activate()	High	7.8
CVE-2026-23097	Linux kernel migrate: correct lock ordering for hugetlb file folios	High	7.3
CVE-2026-23066	Linux kernel rxrpc: Fix recvmsg() unconditional requeue of socket	High	7.4
CVE-2026-23001	Linux kernel macvlan: fix possible use-after-free in macvlan_forward_source() via missing RCU protection	High	7.8
CVE-2026-22801	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	Medium	6.6

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-22695	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	Medium	6.1
CVE-2026-22021	Oracle Java SE / GraalVM JSSE component resource exhaustion via HTTPS; an unauthenticated network attacker may cause a partial denial of service.	Medium	5.3
CVE-2026-22018	Oracle Java SE / GraalVM Libraries component resource allocation flaw; an unauthenticated network attacker may cause a partial denial of service.	Low	3.7
CVE-2026-22016	Oracle Java SE / GraalVM JAXP component information disclosure; an unauthenticated network attacker may gain complete unauthorized read access to all accessible data.	High	7.5
CVE-2026-22013	Oracle Java SE / GraalVM JGSS component protection mechanism failure; a network attacker with user interaction may gain unauthorized read access to data.	Medium	5.3
CVE-2026-22007	Oracle Java SE / GraalVM Security component information disclosure; a local unauthenticated attacker may gain unauthorized read access to a subset of accessible data.	Low	2.9
CVE-2026-20691	Apple Safari/WebKit authorization issue; a maliciously crafted webpage may be able to fingerprint the user.	Medium	4.3
CVE-2026-20676	Apple Safari web extensions state management flaw; a website may be able to track users through Safari web extensions.	Medium	4.3
CVE-2026-20665	Apple Safari/WebKit state management flaw; processing maliciously crafted web content may prevent Content Security Policy from being enforced.	Medium	5.4

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2026-20664	Apple Safari/WebKit out-of-bounds write; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2026-20652	Apple Safari/WebKit memory handling flaw; a remote attacker may be able to cause a denial-of-service.	High	7.5
CVE-2026-20644	Apple Safari/WebKit use-after-free/out-of-bounds write; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2026-20643	Apple Safari/WebKit Navigation API cross-origin issue with improper input validation; processing maliciously crafted web content may bypass Same Origin Policy.	Medium	5.4
CVE-2026-20636	Apple Safari/WebKit memory handling flaw; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2026-20635	Apple Safari/WebKit memory handling flaw; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2026-20608	Apple WebKit (Safari/iOS/macOS/visionOS) state management flaw; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2026-1299	Python email BytesGenerator class did not properly quote newlines in headers when serializing email messages	High	7.1
CVE-2026-0994	Denial-of-service in Python google.protobuf.json_format.ParseDict() via max_recursion_depth limit bypass	High	7.5

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-71085	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.5
CVE-2025-68811	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.1
CVE-2025-68800	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.3
CVE-2025-68741	Linux kernel scsi: qla2xxx: Fix improper freeing of purex item	High	7.3
CVE-2025-68349	Linux kernel NFSv4/pNFS: null layout crash when NFS_INO_ LAYOUTCOMMIT not cleared in pnfs_mark_layout_stateid_invalid	High	7.5
CVE-2025-68121	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.4
CVE-2025-68114	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.3
CVE-2025-67873	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	Medium	6.1
CVE-2025-62231	X.Org X server / Xwayland XkbSetCompatMap() improper bounds checking causes an unsigned short integer overflow, leading to memory corruption or a crash via specially crafted input.	High	7.3

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-62230	X.Org X server / Xwayland XKB extension frees data structures without detaching related resources, causing a use-after-free on client disconnect that can lead to memory corruption or crash.	High	7.3
CVE-2025-62229	X.Org X server / Xwayland Present extension improper error handling leaves dangling pointers, causing a use-after-free that can lead to memory corruption, crash, or arbitrary code execution.	High	7.3
CVE-2025-61729	Go HostnameError.Error() unbounded host list in error string construction	High	7.5
CVE-2025-61728	Go standard library archive/zip uses a super-linear file name indexing algorithm, enabling denial of service via a maliciously constructed ZIP archive.	High	7.5
CVE-2025-61726	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.5
CVE-2025-61662	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.8
CVE-2025-47906	Go os/exec LookPath may return unexpected executable paths when PATH contains executables	Medium	6.5
CVE-2025-46299	Apple WebKit (Safari/iOS/macOS/tvOS/visionOS/watchOS) memory initialization flaw; processing maliciously crafted web content may disclose internal states of the app.	Medium	6.5
CVE-2025-43541	Apple WebKit type confusion in Safari/iOS/macOS/visionOS; processing maliciously crafted web content may lead to an unexpected Safari crash.	High	8.8

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-43536	Apple WebKit use-after-free in Safari/iOS/macOS; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2025-43535	Apple WebKit (Safari/iOS/macOS/visionOS) memory handling flaw; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2025-43531	Apple WebKit race condition in Safari/iOS/macOS/tvOS/visionOS/watchOS; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2025-43529	Apple WebKit use-after-free allowing arbitrary code execution via maliciously crafted web content; actively exploited in targeted attacks against iOS users (CISA KEV).	High	8.8
CVE-2025-43511	Apple WebKit use-after-free in Safari/iOS/macOS/visionOS/watchOS; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2025-43501	Apple WebKit (Safari/iOS/macOS/visionOS) buffer overflow; processing maliciously crafted web content may lead to an unexpected process crash.	High	8.8
CVE-2025-43457	Apple WebKit use-after-free in Safari/iOS/macOS/visionOS/watchOS; processing maliciously crafted web content may lead to an unexpected Safari crash.	High	8.8
CVE-2025-43214	Apple WebKit (Safari/iOS/macOS/tvOS/visionOS/watchOS) memory handling flaw; processing maliciously crafted web content may cause an unexpected Safari crash.	High	8.8
CVE-2025-43213	Apple WebKit (Safari/iOS/macOS/tvOS/visionOS/watchOS) memory handling flaw; processing maliciously crafted web content may cause an unexpected Safari crash.	High	8.8

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-40322	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.1
CVE-2025-40304	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.3
CVE-2025-40294	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.3
CVE-2025-40271	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7
CVE-2025-40269	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.3
CVE-2025-40240	Linux kernel sctp: avoid NULL dereference when chunk data buffer is missing	High	7.5
CVE-2025-40096	Linux kernel drm/sched: Fix potential double free in drm_sched_job_add_resv_dependencies	High	7
CVE-2025-39933	Linux kernel smb: client: recv_done verifies data_offset, data_length and remaining_data_length	High	7.1
CVE-2025-39760	Linux kernel usb: core: config: Prevent out-of-bounds read in SS endpoint companion parsing	High	7.1

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-38730	Linux kernel io_uring/net: commit partial buffers on retry to prevent data loss	High	7.1
CVE-2025-38703	Linux kernel drm/xe: Make dma-fences compliant with safe access rules	High	7
CVE-2025-38568	Linux kernel net/sched: mqprio: fix stack out-of-bounds write in tc entry parsing	Medium	4.7
CVE-2025-38415	Linux kernel Squashfs: check return result of sb_min_blocksize to prevent integer overflow	High	7.1
CVE-2025-38349	Linux kernel eventpoll: don't decrement ep refcount while still holding ep mutex	High	7
CVE-2025-38248	Linux kernel bridge mcast: Fix use-after-free during router port configuration	High	7.3
CVE-2025-38206	Linux kernel exfat: fix double free in delayed_free	High	7.3
CVE-2025-38180	Linux kernel net: atm: fix /proc/net/atm/lec handling to prevent memory issues	High	7.1
CVE-2025-38154	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	Medium	6.7

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-38141	Linux kernel dm: fix dm_blk_report_zones to prevent incorrect zone reporting	High	7
CVE-2025-38129	Linux kernel page_pool: Fix use-after-free in page_pool_recycle_in_ring	High	7.3
CVE-2025-38106	Linux kernel io_uring: fix use-after-free of sq->thread in __io_uring_show_fdinfo()	High	7.1
CVE-2025-38051	Linux kernel SMB client: race condition use-after-free in cifs_fill_dirent during concurrent readdir	High	7
CVE-2025-38024	Linux kernel RDMA/rxe: Fix slab-use-after-free Read in rxe_queue_cleanup	High	7.3
CVE-2025-38022	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.1
CVE-2025-37882	Linux kernel usb: xhci: Fix isochronous Ring Underrun/Overrun event handling	Medium	6
CVE-2025-37861	Linux kernel scsi mpi3mr: Synchronous access between reset and tm thread for reply queue	Medium	6.3
CVE-2025-22056	Linux kernel netfilter: nft_tunnel: fix geneve_opt type confusion addition	High	7

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-21791	Linux kernel vrf: use RCU protection in l3mdev_l3_out() to prevent race condition	High	7.1
CVE-2025-21786	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.8
CVE-2025-21647	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.1
CVE-2025-15367	Python poplib module command injection via newlines when passed user-controlled commands	High	7.1
CVE-2025-15366	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.1
CVE-2025-14831	GnuTLS denial-of-service via excessive CPU and memory consumption with specially crafted malicious certificates	Medium	5.3
CVE-2025-14512	GLib GIO escape_byte_string() integer overflow causes heap buffer overflow and denial of service when processing malicious file or remote filesystem attribute values.	Medium	6.5
CVE-2025-14174	Google Chrome ANGLE component (Mac) out-of-bounds memory access via a crafted HTML page; actively exploited in the wild (CISA KEV).	High	8.8
CVE-2025-14087	GLib (GVariant parser) heap corruption via buffer-underflow when processing maliciously crafted input strings, potentially leading to denial of service or code execution.	Medium	5.6

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2025-12801	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	Medium	6.5
CVE-2025-12084	Python xml.dom.minidom quadratic complexity in nested element building using appendChild() and _clear_id_cache()	High	7.5
CVE-2025-10158	rsync malicious receiver client can trigger an out-of-bounds read of a heap buffer via a negative array index, requiring at least read access to the remote module.	Medium	4.3
CVE-2024-56645	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.1
CVE-2024-56603	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.3
CVE-2024-54456	Linux kernel NFS: potential buffer overflow in nfs_sysfs_link_rpc_client() via unchecked strcat on fixed-size buffer	High	7.1
CVE-2024-53229	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7
CVE-2024-53216	Linux kernel nfsd: release svc_expkey/svc_export with rcu_work to prevent use-after-free	High	7.1
CVE-2024-47727	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.8

CVE Name	Details	Threat Severity	CVSS Base Score
CVE-2024-12086	rsync server-side flaw allows enumeration of arbitrary client file contents by sending specially crafted checksum values, enabling byte-by-byte data reconstruction (information disclosure).	Medium	6.1
CVE-2023-52356	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	High	7.5
CVE-2023-40403	Vulnerability requiring remediation in 7.4.1 (NVD record pending full analysis)	Medium	6.5
CVE-2026-45447	OpenSSL PKCS#7/S/MIME signature verification (PKCS7_verify) use-after-free when the SignedData digestAlgorithms field is an empty ASN.1 SET; potentially enabling remote code execution.	High	8.8

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