COHESITY

Protect Nutanix Acropolis Hypervisor with Cohesity Data Cloud

Bringing Simplicity, Scalability, and Resiliency to Nutanix AHV Protection

Version 3.3

April 2026

ABSTRACT

Both Cohesity and Nutanix provide hyper-converged infrastructure (HCI) that delivers simple, efficient, and cost-effective ways to keep data and applications safe. Cohesity Data Cloud is designed and engineered to provide services to organizations with an entirely new and modern data protection and recovery approach. This solution brings a snapshot-based, incremental-forever backup approach on a scalable and robust platform for your Nutanix AHV backups.

Table of Contents

Figures

Tables

Introduction to Nutanix AHV Protection

Traditional data center infrastructure has many separate silos of computing, storage, networking, and virtualization resources. Procurement, deployment, and management of these separate components is time-consuming and labor-intensive. Legacy enterprise data protection and recovery typically consists of complicated and expensive products and solutions. Today, with new types of applications driving more digital business, modernizing and protecting your data center, extending it to the cloud, and developing for the cloud have never been more critical.

The Nutanix Acropolis operating system is a hyper-converged infrastructure (HCI) platform with built-in virtualization, Acropolis Hypervisor (AHV). Nutanix AHV is native, enterprise-grade virtualization and is included with Enterprise Cloud OS.

Combining Cohesity Data Cloud with Nutanix AHV enables you to protect virtual machines and other workloads like NC2 and Nutanix Files that run on Nutanix AHV. This solution employs the Nutanix API to deliver protection using a snapshot-based, incremental-forever backup approach on a scalable, resilient platform that provides a single management pane for backup, instant and granular recovery, replication, disaster recovery, cloud archive, and cloud tiering.

Some of the advantages of this solution include:

- An API-first architecture
- A snapshot-based workflow provides granular control
- Incremental-forever backups capture, move, and store only changed data, reducing ingress and egress costs
- Distributed, parallel, and workload-optimized ingest
- Instant Mass Restore
- Global actionable search
- Global space efficiency with variable length, sliding window deduplication, and Zstd compression

Figure 1: Protect Nutanix AHV VMs with Cohesity Data Cloud

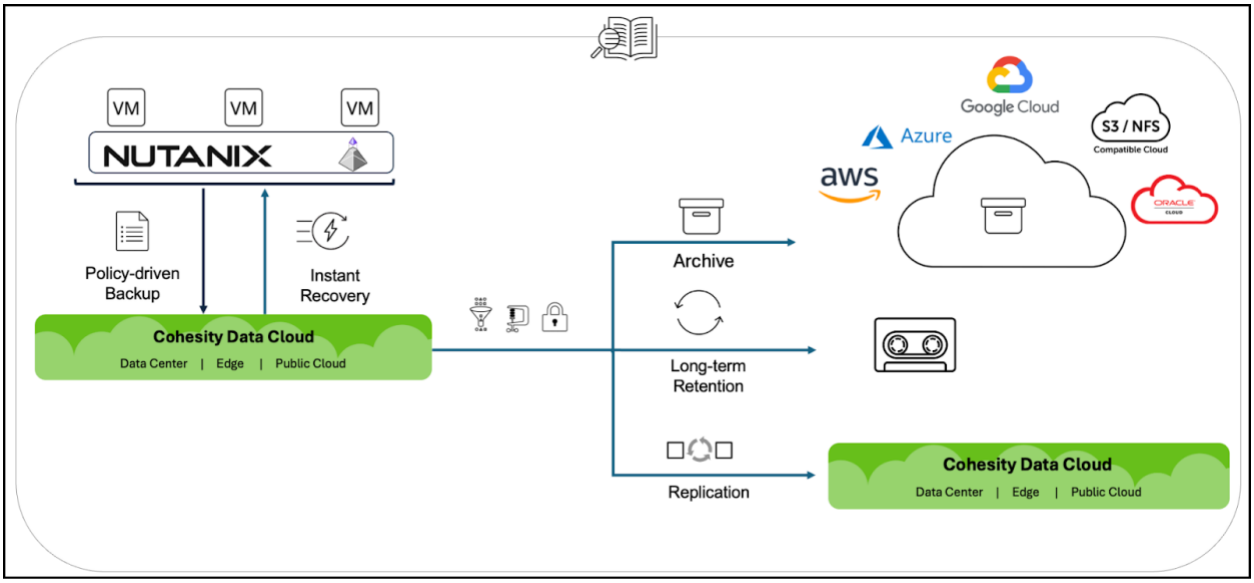

NOTE: Cohesity Data Cloud supports VMWare ESXi, Nutanix AHV, and Microsoft Hyper-V running on the Nutanix Acropolis operating system.

This guide focuses on data protection for Nutanix AHV and is written for virtualization, storage, and backup architects and administrators responsible for enterprise data protection.

Terminology

Several concepts and terms are important as you learn about Cohesity's data protection solution for Nutanix AHV.

Table 1: Terminology for Nutanix AHV with Cohesity Data Cloud

Term	Definition
Cohesity Cluster	A Cohesity cluster is an instance of Cohesity Data Cloud. Each cluster contains at least three nodes and can be scaled up as needed.
Storage Domain	A Storage Domain is a named storage location on a Cohesity cluster. It defines the deduplication, compression, fault tolerance, and encryption settings. You can create a View in a Storage Domain, which provides NFS, SMB, and S3-compatible mount paths to access the storage.
Nutanix Cluster	A Nutanix cluster is a group of three or more physical nodes working as a single entity on the Nutanix software.
Nutanix Prism	Nutanix Prism is a unified infrastructure management platform that provides actionable insights for optimizing virtualization, infrastructure management, and everyday operations.
Nutanix CVM	A Nutanix Controller Virtual Machine (CVM) is a virtual machine running Nutanix software. It runs on each node in the Nutanix cluster and serves all the I/O operations for the hypervisor and all the VMs running on the host.

Cohesity Solution for Nutanix AHV Protection

Cohesity Data Cloud integrates seamlessly with Nutanix AHV. You can register Nutanix Prism as a protection source. After registration, Cohesity Data Cloud auto-discovers Nutanix cluster nodes, VMs, and vDisks. Users can define Protection Policies for backup, replication, and archival based on business needs.

Figure 2: Nutanix AHV with Cohesity Data Cloud Solution Overview

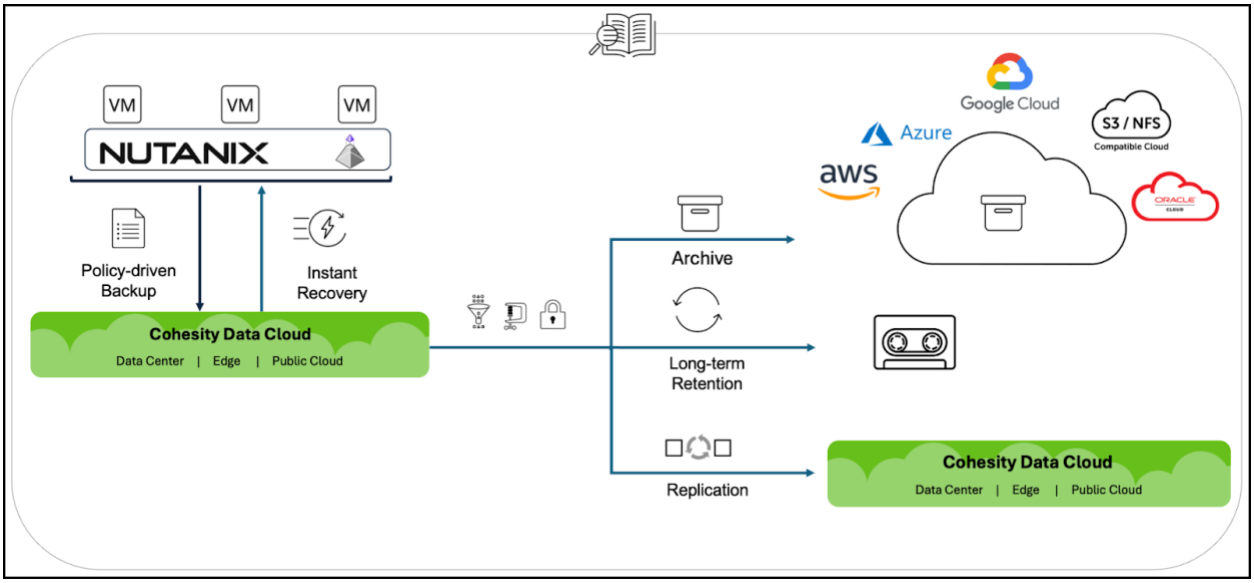

Cohesity Data Cloud Overview

Cohesity Data Cloud consolidates data and applications, including backups, files, objects, test/dev, and analytics, on a single, software-defined platform. Inspired by web-scale architecture, it is a scale-out solution based on a unique distributed file system, SpanFS™. While most organizations begin their journey to overcoming mass data fragmentation by simplifying data protection, Cohesity Data Cloud's flexible architecture allows easy expansion to many additional use cases, further increasing operational simplicity and improving TCO (Total Cost of Ownership).

Solution Features and Benefits

Cohesity Data Cloud is tightly integrated with Nutanix to provide simple, efficient, reliable, and fast data protection and recovery.

Using Cohesity Data Cloud to protect your Nutanix AHV workloads provides many significant benefits.

Table 2: Features and Benefits

Feature	Benefit
Scale-out architecture	Cohesity's scale-out architecture allows organizations to start with the right-size data platform ideal for their environment and scale out as needed, on their own terms, one node at a time. It allows for parallel backups and recoveries across all nodes and enhances performance. Both Cohesity and Nutanix utilize scale-out architecture to deliver high performance and ease of use.
API-first architecture	Every interface, component, and integration in Cohesity Data Cloud is built on the Cohesity Data Cloud REST API. As a result, it integrates seamlessly with Nutanix, making Nutanix cluster registration, backup and recovery, and archive easy to execute without needing agents or proxies on AHV or VMs.
Distributed and Parallel Ingest	With its Tier-Optimized Write Scheme (TOWS), the ingest engine ensures that data is placed onto the type of disk, SSD, or HDD that best suits the profile of the incoming data stream. In TOWS, HDDs, which prefer sequential I/O, write data out-of-place, while SSDs are used for random I/O. The ingest engine also includes adaptive data throttling to modulate backup ingest performance over the production workloads at the Hypervisor or Datastore level.
Megafile	To avoid bottlenecks when a VM is assigned to a single node, the Megafile feature splits each virtual disk in the VM into parts, 'megafile chunks,' and distributes them to all nodes in the cluster. As a result, as the number of nodes in a Cohesity cluster grows, the time to ingest decreases dramatically.
Incremental backup with Nutanix CRT	Cohesity Data Cloud integrates with Nutanix Change Region Tracking (CRT), providing faster and more efficient backups. Only the changed metadata regions that the CRT detects are backed up and written. CRT also tracks the regions that contain zeros and avoids reading those zeros, which is beneficial for full backups.

Feature	Benefit
Rapid search and recovery	Cohesity Data Cloud automatically indexes the backup data and all its associated metadata, providing wild-card global search results for near-instantaneous granular restores of individual files and data objects like VMs.
Storage efficiency	Cohesity Data Cloud leverages a unique, variable-length data deduplication technology that spans an entire cluster, resulting in significant savings across a customer's storage footprint. You can apply deduplication inline (as data is written) or post-process (after data is written) to optimize performance against backup time windows. In addition, compression of the deduplicated blocks further maximizes space efficiency.
Encryption	The Cohesity file system (SpanFS™) provides full at-rest and in-flight encryption based on the strong AES-256 standard, ensuring that the data stored on a Cohesity cluster is protected from malicious attacks.
Fast Access	Cohesity Data Cloud allows faster or instant recovery of multiple VMs. With Instant Mass Restore, you can use the Cohesity NFS views as the datastore, allowing the recovery process to happen much faster or instantly and circumventing the requirement of restoring the VMs to the original datastore before booting.
Data Resiliency	Cohesity Data Cloud enables the continuation of business operations and rapid recovery in the event of unexpected disruptions or data loss.
Data Immutability	DataLock, a time-bound, write-once, ready-many (WORM) locks on a backup snapshot that can't be modified in our file system (and extends to cloud storage by incorporating S3 object lock)
Threat Defence Architecture	Cohesity Data Cloud has built a multilayered data security architecture to combat the Ransomware 3.0 threat and help our customers achieve specific outcomes related to data protection, compliance, operational resiliency, defense against sophisticated attacks, and aggressive ransomware variants.

Feature	Benefit
Data Classification	Automated discovery and classification of sensitive data for security and compliance with out-of-the-box 200+ production-grade, AI/ML-based patterns.
Cloud Integration	Natively integrated with major public cloud providers, it provides smooth archival and tiering to the cloud.

Use Cohesity Data Cloud to Protect Nutanix AHV

Cohesity Data Cloud integration with Nutanix Prism simplifies AHV protection. Registering an AHV cluster and discovering all AHV vDisk objects is seamless. Protection policies can be assigned to AHV objects to automate data protection and recovery.

In this solution, the backup process includes the following steps:

1. Register your Nutanix cluster.
2. Use a Protection Policy.
3. Create a Protection Group.

Prerequisites

Nutanix Privileges for Cohesity

The Cohesity cluster needs to be able to perform different actions on the Nutanix source. To enable this, the user designated to connect to the source (the one who registered the source) must have sufficient privileges. Any local account with any one of the **'user admin'** or **'cluster admin'** roles work without any further configuration. For more details on the necessary privileges, please refer to "Ensuring Adequate Privileges for Cohesity".

Cohesity and AHV Cluster Firewall Port Requirements

Certain ports in the firewall must be open to allow the Cohesity cluster to transmit and receive data. The cluster can handle all traffic types on a single IP network, which is the most tested and used approach. However, the Cohesity cluster can also connect to multiple networks as required within enterprise networks.

TIP: Cohesity recommends using a single IP network.

Table 3: Required TCP Ports for backup recovery operations

Traffic Direction	Source	Destination	Port	Protocol	Usage Notes	Type of Traffic
Incoming	Data Services IP and IP of all CVMs on the Nutanix cluster and Prism Central	VIPs and VLAN IPs on Cohesity cluster	111 2049	TCP	Required for Instant Recovery	Backup and Recovery
Outgoing	Cohesity cluster	Data Services IP and IP of all CVMs on the Nutanix cluster and Prism Central.	9440	TCP	Used for REST API calls	Backup and Recovery
Outgoing	Cohesity cluster	Data Services IP and IP of all CVMs on the Nutanix cluster and Prism Central.	3205, 3260	TCP	iSCSI port for AHV backups.	Backup and Recovery

Cohesity Backup Workflow

A Cohesity Data Cloud Protection Group coordinates actions with Nutanix Prism, AHV hosts, CVM, the guest OS, and all nodes in the Cohesity cluster.

Cohesity Data Cloud backup and recovery of AHV VMs is performed using the Nutanix REST API without proxies, thereby simplifying operational management. At a high level, creating a backup involves the following steps:

1. Connect to the registered Nutanix source and authenticate with the credentials.
2. Use the snapshot API to create an application-consistent or a crash-consistent snapshot.
3. Get the virtual machine configuration details, such as VM power state, UUID, etc., and store this configuration for later recovery actions.
4. A temporary Volume Group is created with an attached vmdisk to ingest data from the Nutanix to the Cohesity cluster.
5. If a full backup is required, each disk in the snapshot is backed up. The Change Region Tracking (CRT) APIs are used for full backups, preventing zeroed regions' backups. For incremental backup, the CRT API is used to query the metadata of the regions that have changed between the current and the reference snapshot.
6. Surface the snapshots of the disks to be ingested. Read and back up the desired regions from the surfaced snapshots.
7. Use the DELETE A SNAPSHOT API call to delete snapshots and volume groups that are no longer required. One snapshot is maintained on the Nutanix cluster for reference.

Figure 3: AHV VM Incremental Backup Workflow

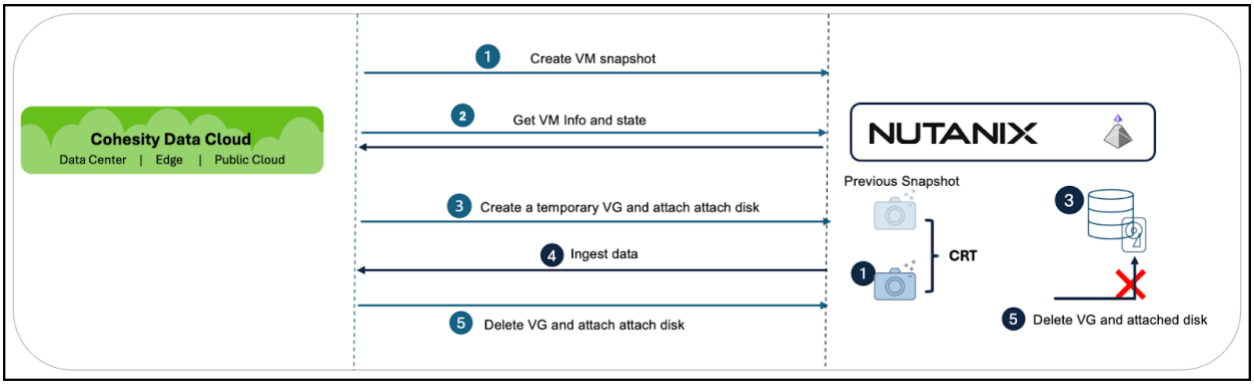

If there is no previous snapshot, Cohesity Data Cloud will do a full backup, transferring only the data that consumes storage capacity.

Register a Nutanix Cluster

In Cohesity Data Cloud, Sources are protected workloads. You can register multiple Nutanix AHV clusters as protection Sources. To register an AHV cluster, you will use a Prism hostname or IP address and the appropriate Prism user with adequate privileges to perform actions on the Source. For more information, see Nutanix Acropolis in Ensuring Adequate Privileges for Cohesity in the Cohesity documentation.

NOTE: Starting the 7.2.2_u2 release, Cohesity now supports Prism Central along with standalone Nutanix clusters.

Every Nutanix AOS cluster has two types of IP addresses: cluster Virtual IP and iSCSI Data Services IP, described in the table below.

Table 4: Nutanix AOS IP Address Types

IP Address Type	Details
Cluster Virtual IP Address	This highly available IP address allows access to management services on a Nutanix AOS cluster. It eliminates the need to connect to individual controller virtual machines and concerns about their temporary unavailability. This IP address is used to talk to the REST API services hosted by the AOS cluster.
iSCSI Data Services IP Address	This IP address is used to discover targets and ingest data using iSCSI. It is important to use this address for iSCSI traffic target discovery because it: • Load-balance storage requests to improve ingestion performance. • Enables path optimization in the cluster to prevent bottlenecks, which reduces latencies. • Eliminates the need to configure a multipathing service such as multipath I/O (MPIO).

To register a Nutanix cluster:

Create Prism Element Cluster Local User Account

1. Log in to the Nutanix Prism Element cluster and select **Settings** from the drop-down menu.
2. In the **Settings** section under **Users and Roles,** select **Local User Management** to create a new user with **Cluster Admin** or **User Admin** privileges on the Nutanix cluster.
3. Click **New User**. In the open form, enter the new user's **Username, First Name, Last Name, Email address,** and **Password.** Select the user's **Language** and the **User Admin** or **Cluster Admin** role, then click **Save**.

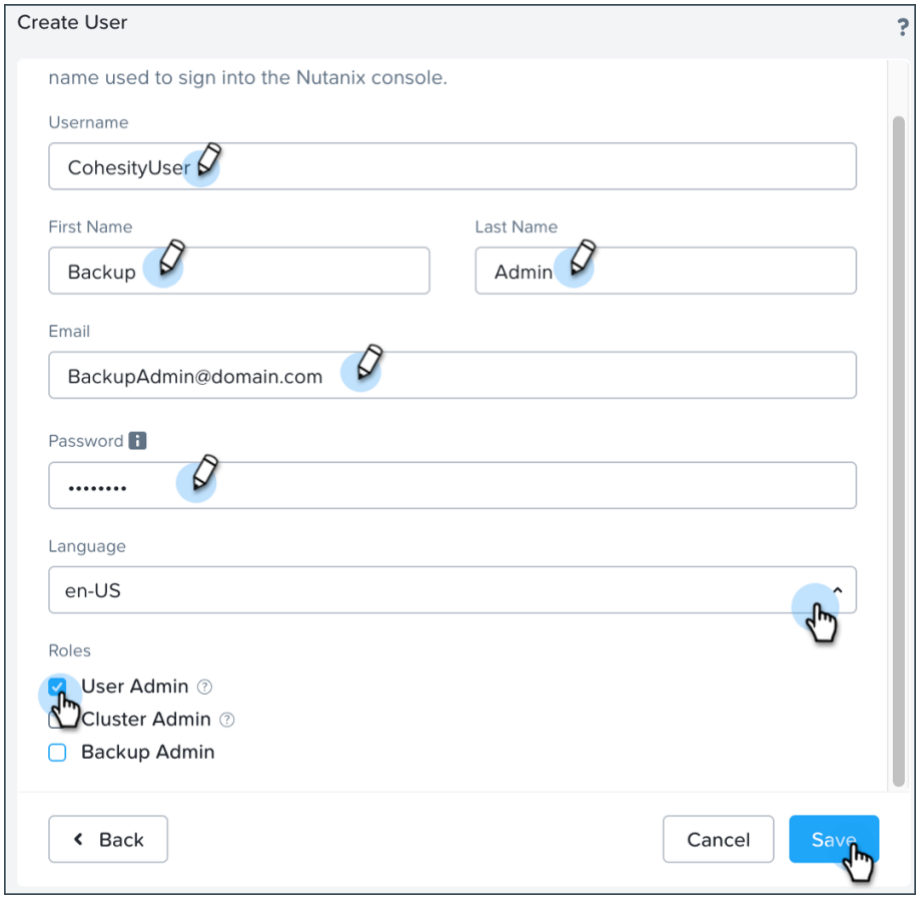

NOTE: User Administrators are automatically assigned to be cluster and backup administrators, and Cluster Admins are automatically assigned to be backup administrators.

Create Prism Central Local User Account

1. Log in to Nutanix Prism Central and select **Settings** from the drop-down menu.
2. Go to Platform Services – **Admin Center** and select **Identities.**
3. Click the "**Add Local User**" button to create a new user.
4. In the open form, enter the new user's **First Name, Last Name, Email address, Username,** and **Password.** Select the user's **Language** and **Enable User**, then click **Create**.

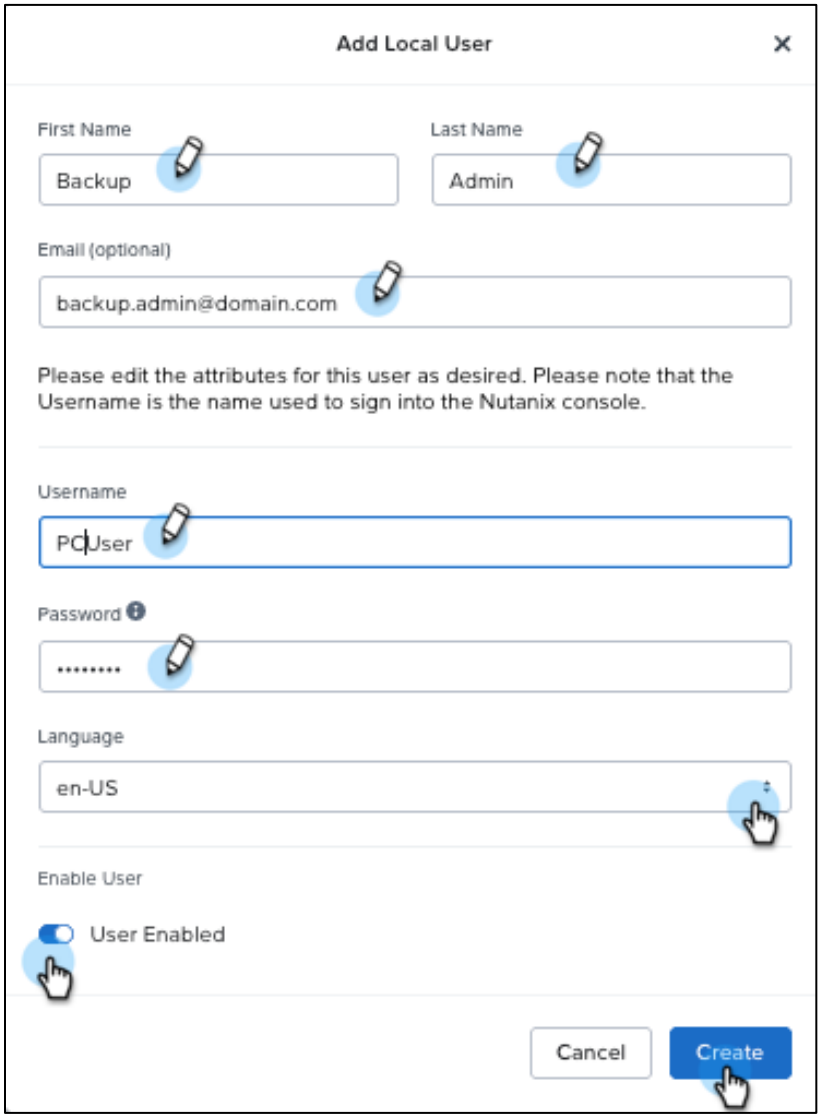

5. Map this user to the Role that has Admin privileges or adequate privileges to perform actions on the source by using Authorization Policy.

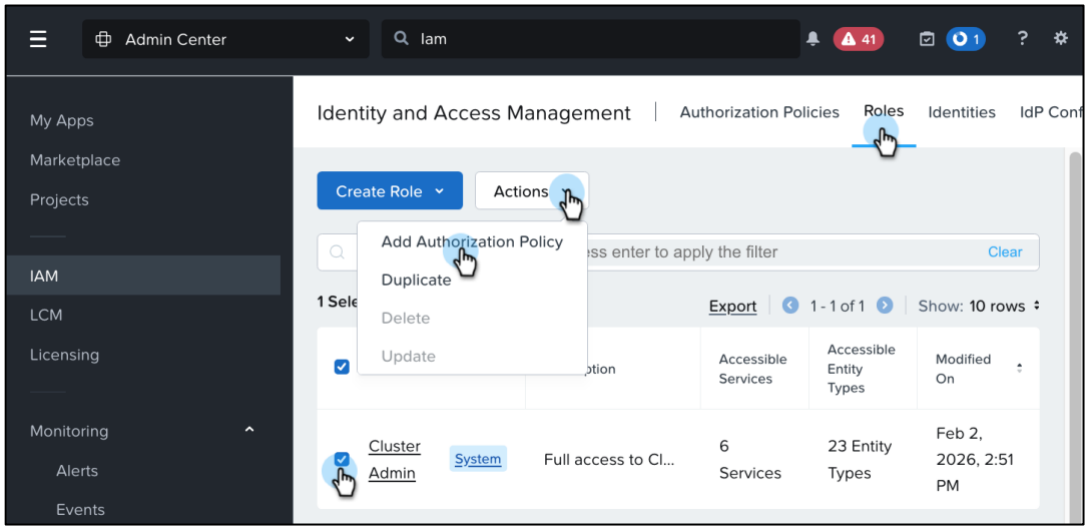

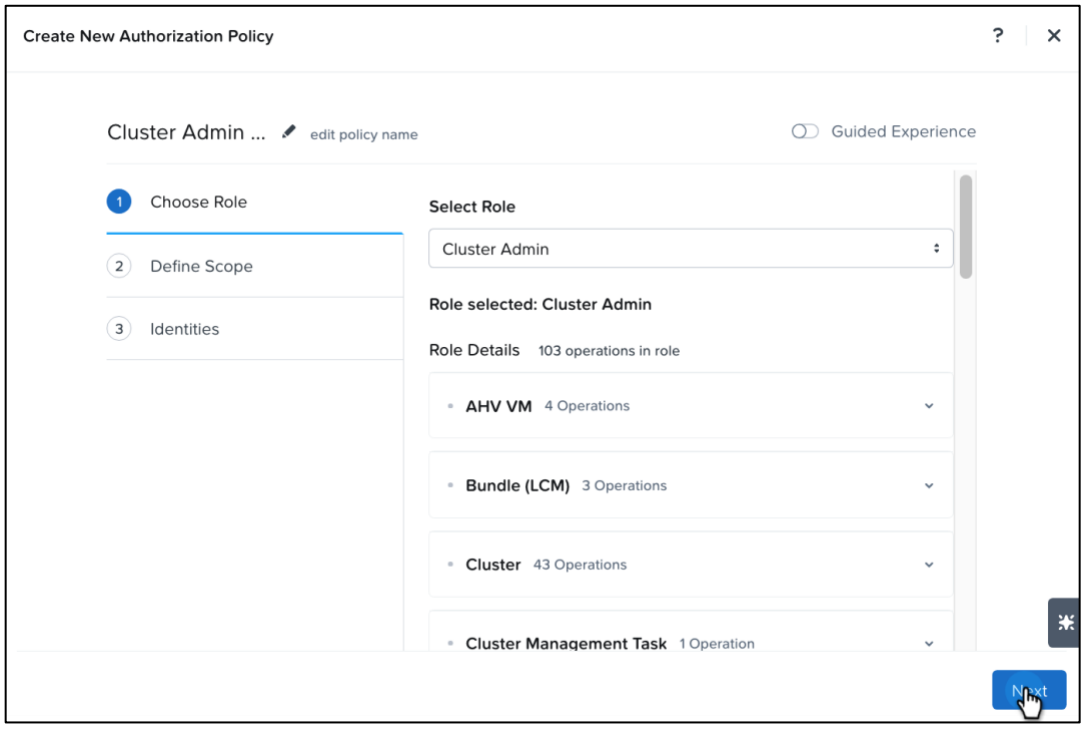

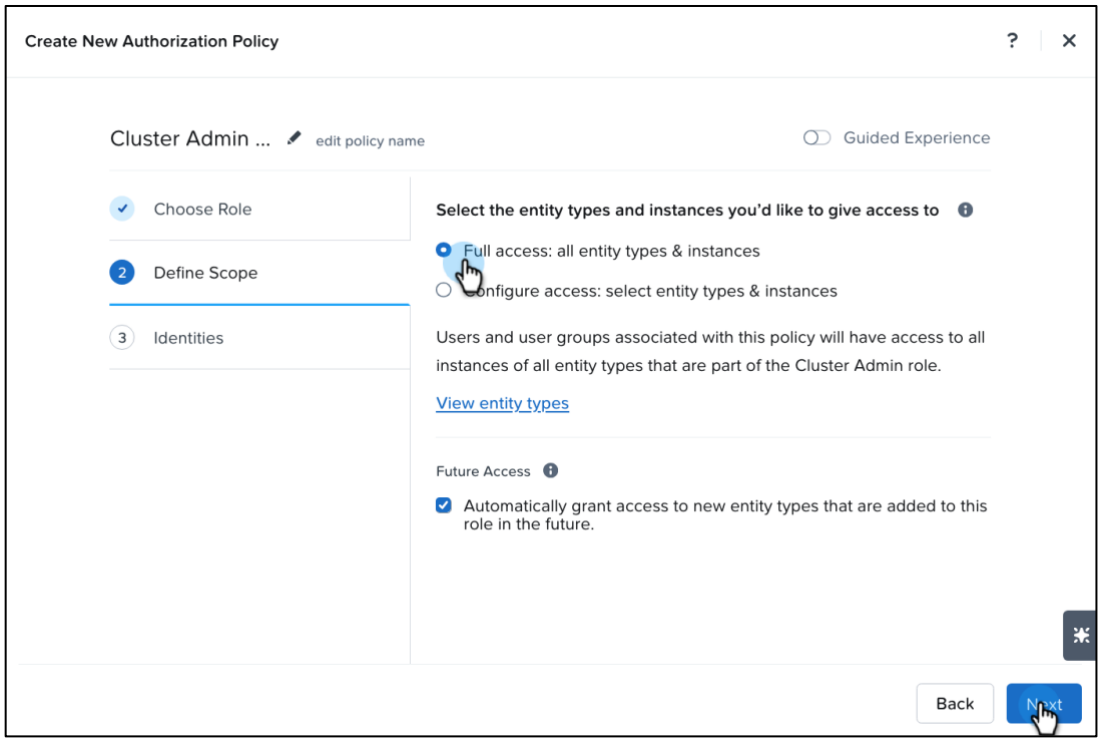

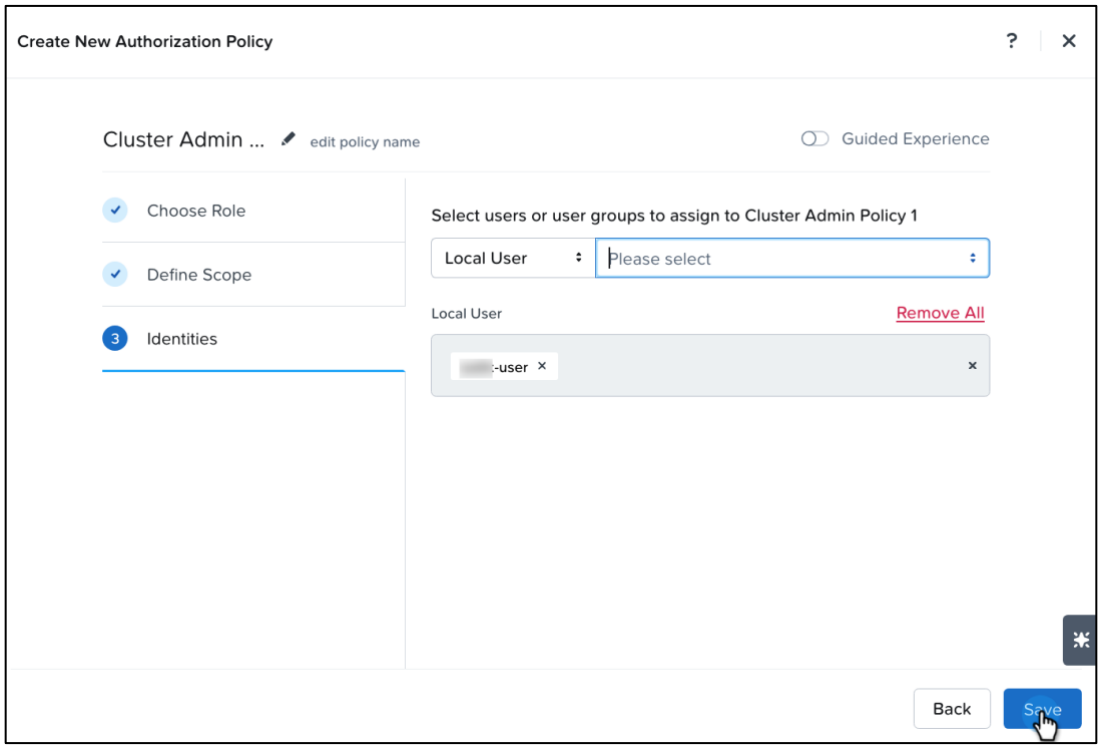

NOTE: Refer Ensure Adequate Privileges for Cohesity on the Source for details on Prism Central Granular user privileges.

6. Log in to Cohesity Data Cloud and select **Data Protection** > **Sources.** Click **Register** and select **Virtual Machines** from the drop-down.

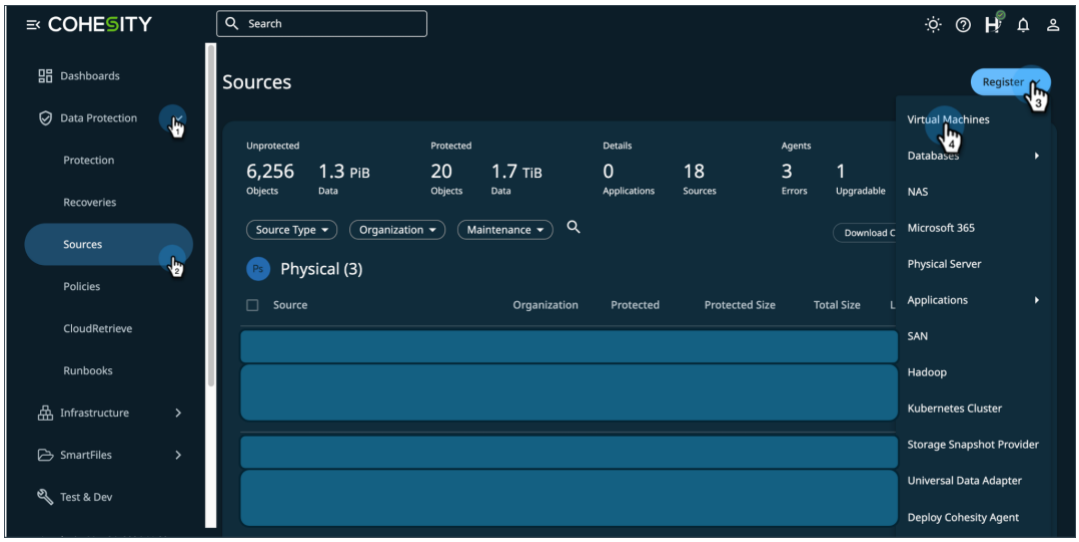

7. In the open form, under **Source Type**, choose **Acropolis: Acropolis Standalone Cluster or Acropolis: Prism Central** for AHV or **Acropolis: Other Hypervisors** for AOS on ESXI. Then, enter the **Hostname** or **IP Address** of your Nutanix Prism Central or Prism Element and the **Username** and **Password** for the user you created. (With this information, Cohesity Data Cloud can auto-discover the Nutanix Cluster objects.) Click **Register**. For more information on Source Registration refer to Register or Edit a Hypervisor Source

8. When Source Type is **Acropolis: Acropolis Standalone Cluster.**

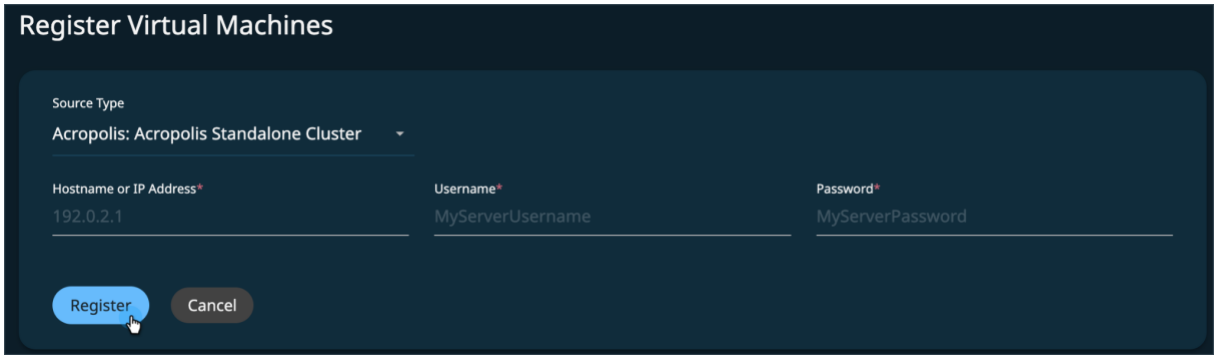

TIP: Start typing in the Select Hypervisor Source Type to narrow the choices.

9. When Source Type is **Acropolis: Prism Central.**

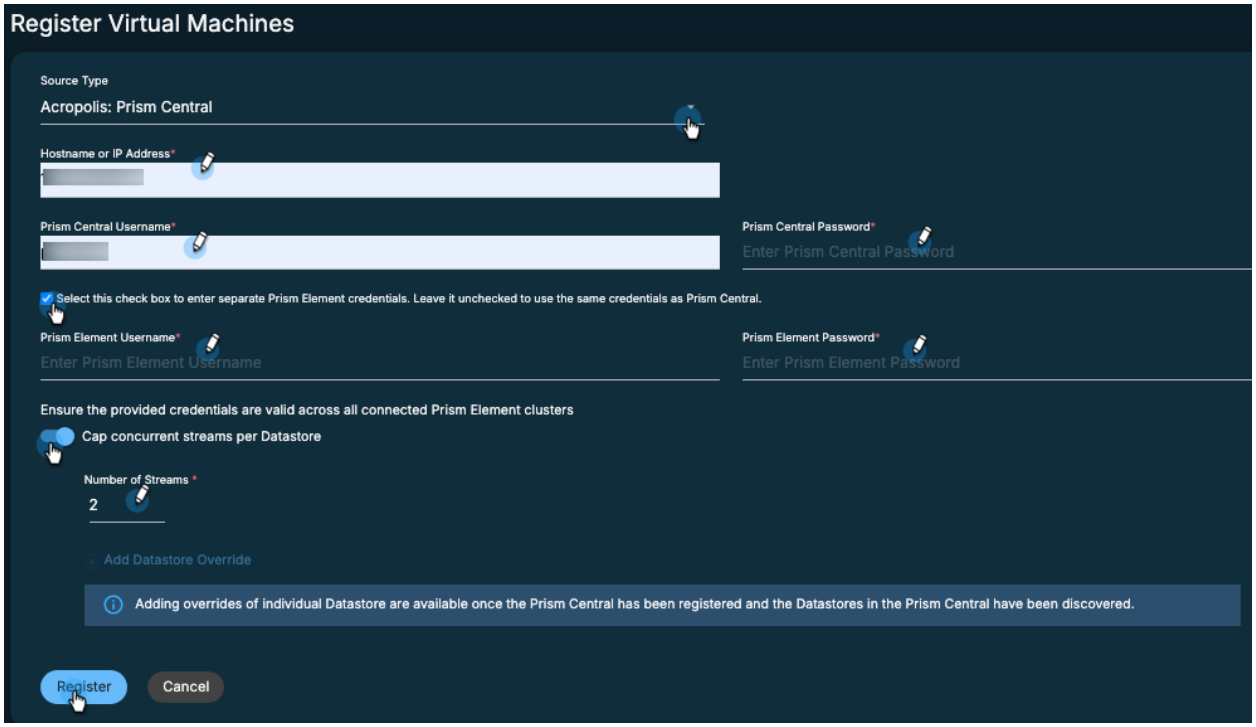

NOTE:

- If the same user account is used for both Prism Central and Prism Element credentials, you do not need to manually re-enter the Prism Element credentials.
- If any of the Prism Elements have incorrect credentials, a warning will be logged for those credentials, but the other Prism Element sources will still be registered.
- Starting release 7.4, you can configure option “Cap concurrent streams per Datastore” for fair resource utilization across workloads when registering Prism Central source.

Use a Protection Policy

A protection policy is a reusable set of settings that define how and when objects are protected, replicated, and archived. Refer to the Create and Edit a Standard Policy on Cohesity documentation for steps to configure a protection policy and its various options.

Create a Protection Group

A backup job that runs repeatedly, based on an associated policy, to back up data from a source and store it on the cluster. A Protection Group uses the schedules and settings defined in the policy to determine when and how backups are captured, archived, or replicated.

Multiple Protection Groups can use the same Protection Policy, and each Job can have only one Policy. Protection Group protects specific source objects, such as virtual servers, physical servers, Views, databases, NAS storage, and more.

To create a Protection Group:

1. Log in to Data Cloud and select **Data Protection** > **Protect**. In the Protect dropdown, select **Virtual Machines.**

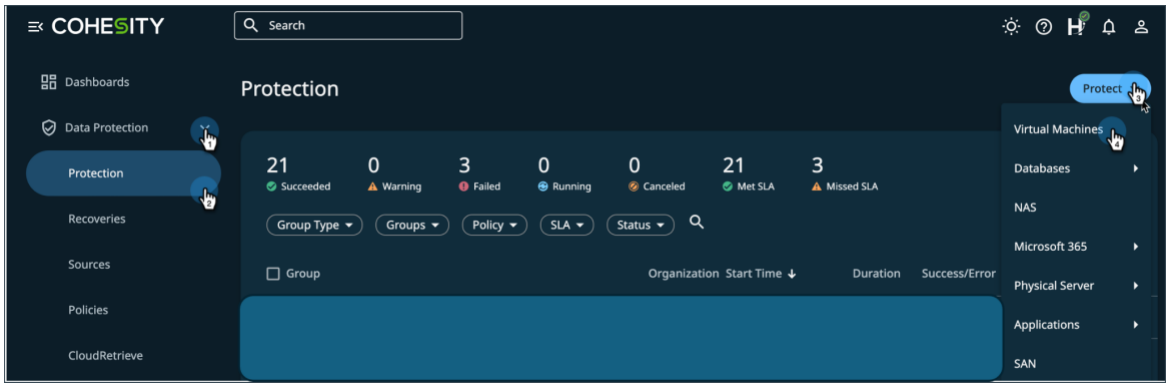

2. In the form that opens, click **Add Objects.**

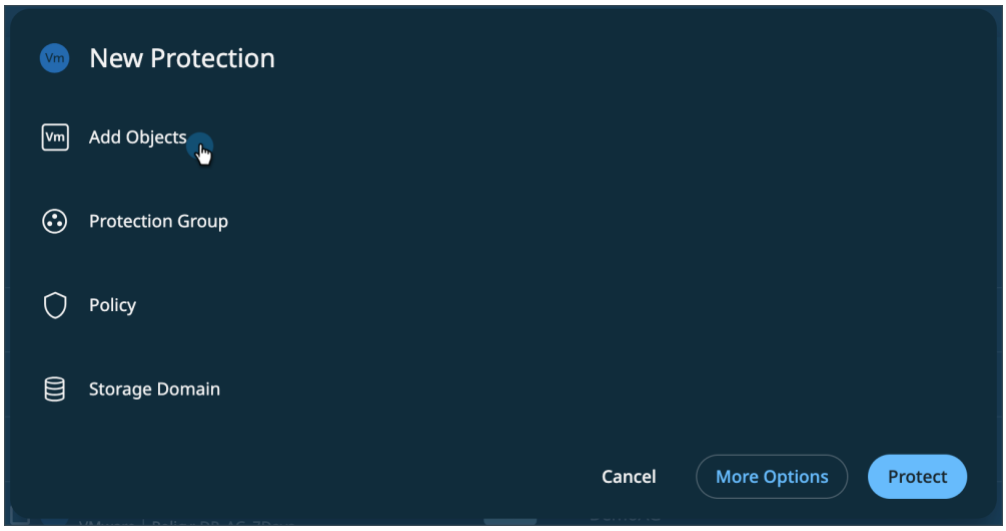

3. Under **Registered Source,** select the desired source (Prism Element or the Prism Central as applicable) and the VM objects to protect, and click **Continue**.

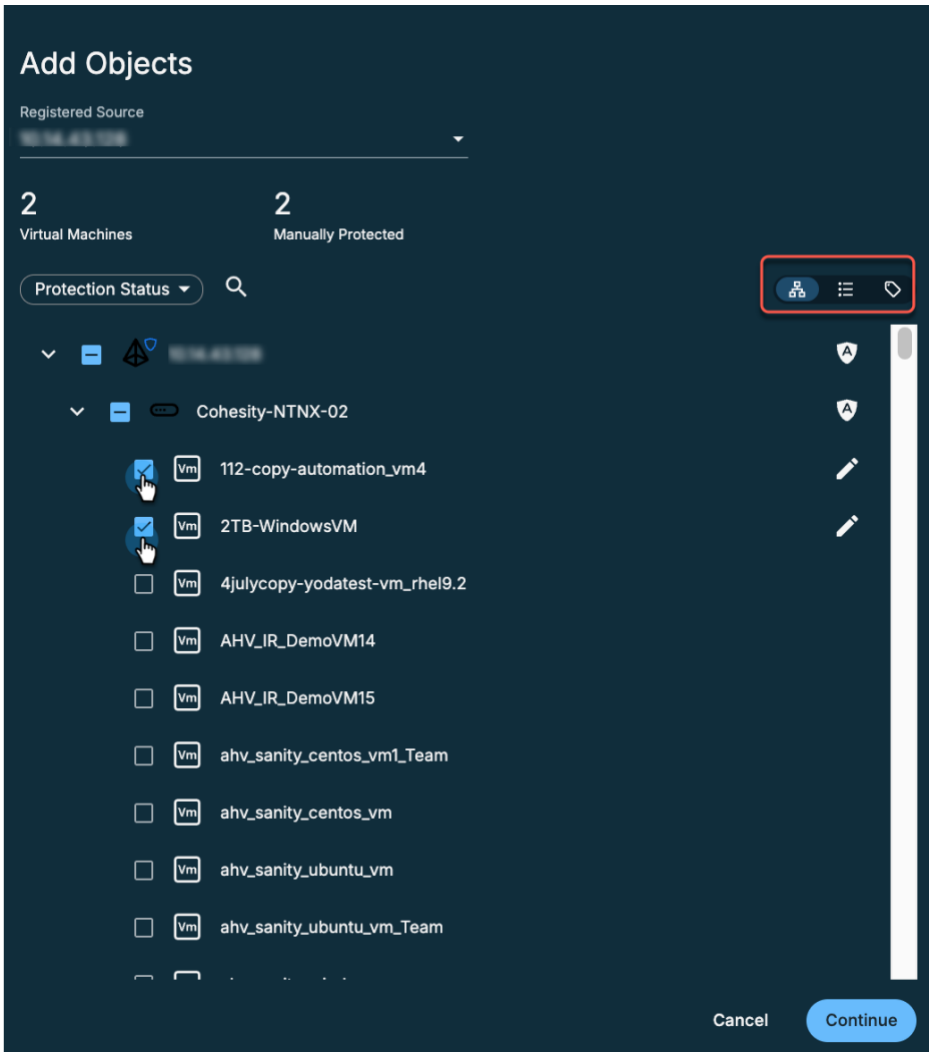

Starting Cohesity release 7.3, VM object selection or exclusion by VM tags (VM Categories – Value) is supported with the following options:

- Auto Protect This Tag
- Select All Child Objects
- Exclude This Tag

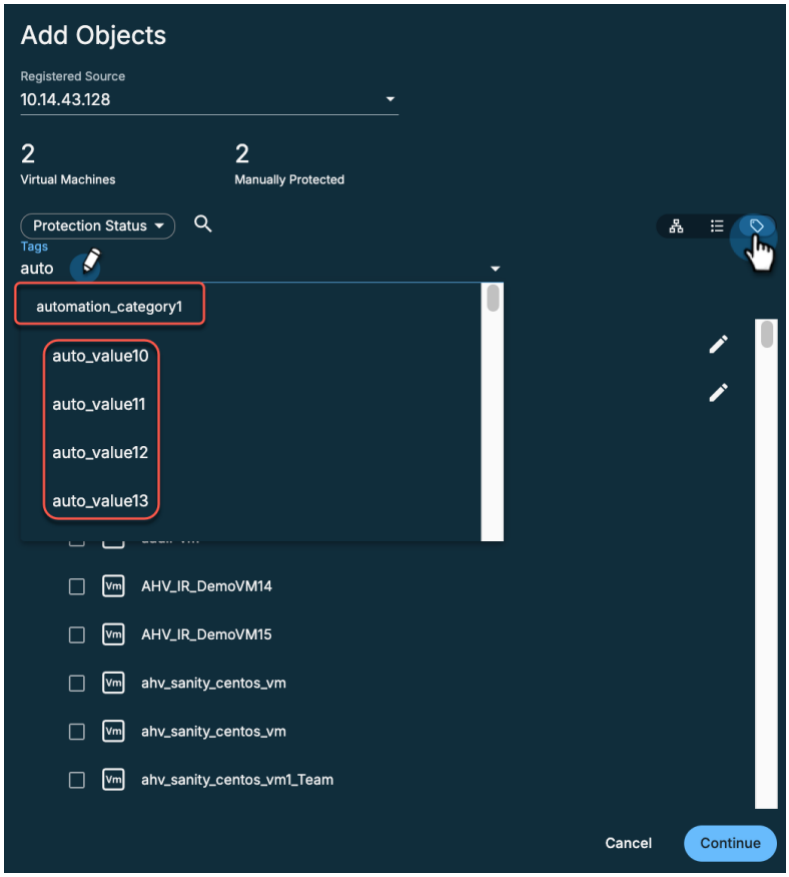

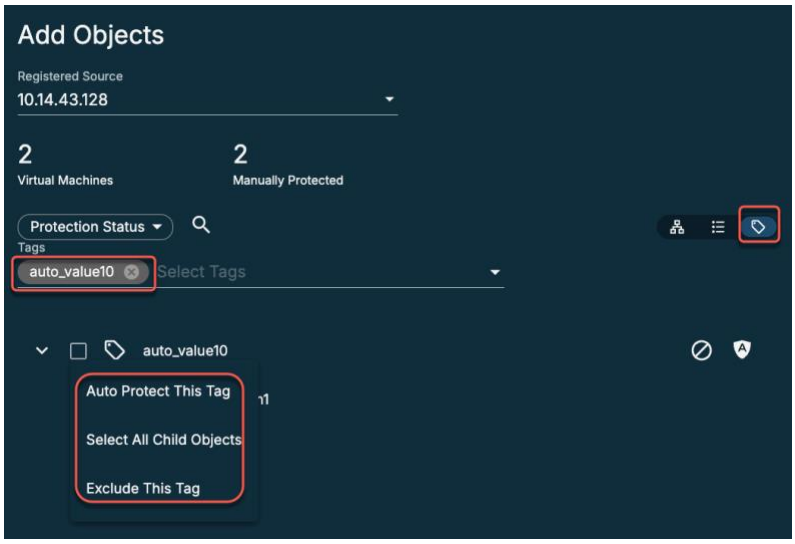

TIP: With Auto-Protect enabled on a protected object, Cohesity Data Cloud dynamically protects that object and any future child objects whenever the Protection Group runs. Click the **Auto-Protect** shield icon on the right to apply it to the object.

4. Select an appropriate name for the Protection Group. Then select the **Protection Policy** you need. You can choose a predefined Policy (**Bronze**, **Silver**, or **Gold**), a custom Policy, or click **New Policy** to create a new one. In the same form, select the target **Storage Domain** to use.

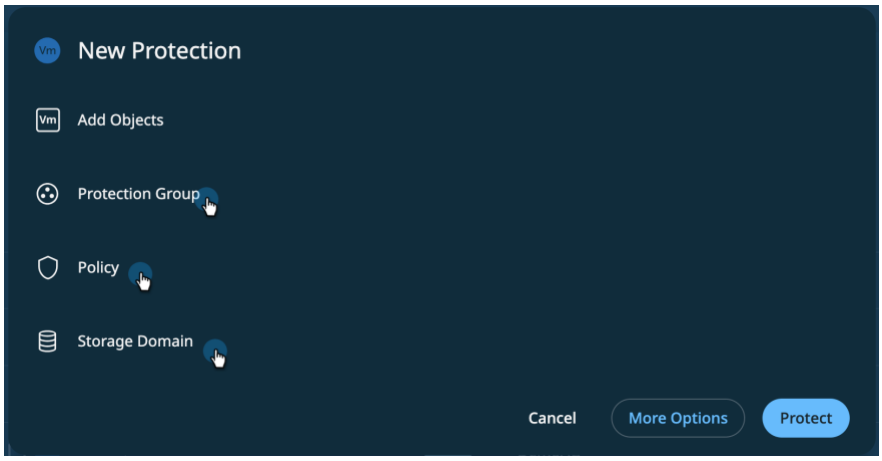

5. If you need to change any of the **Additional Settings**, click **More Options.** When you're done, click **Protect.**

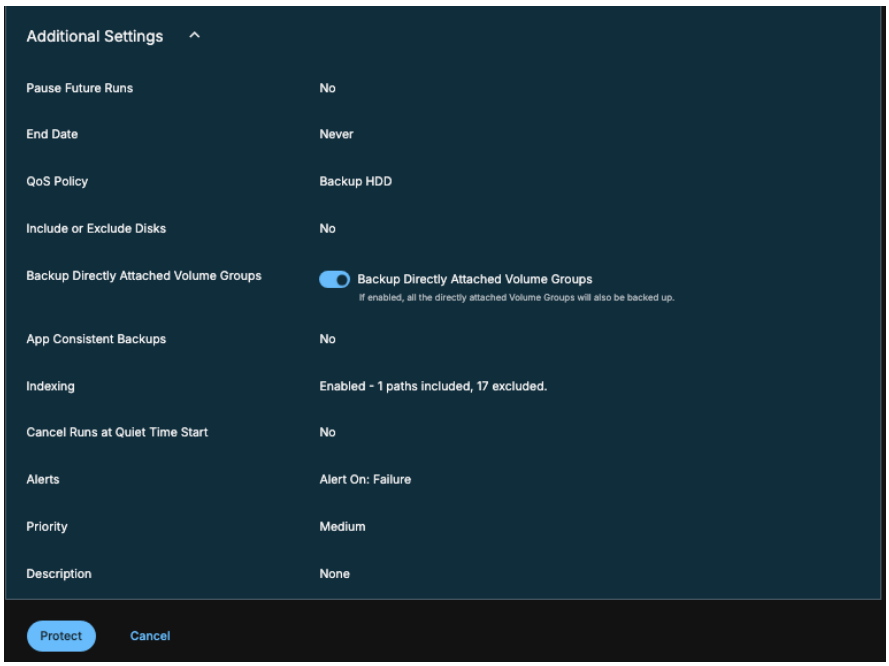

NOTE: Starting release 7.4, Cohesity allows the backup of Volume Groups that are directly attached to the VM.

Protection Groups combine operational flexibility with the business requirements defined in a Protection Policy. For advanced VM Protection Group settings, see Add or Edit a Protection Group for Virtual Servers in the Cohesity documentation.

Disk Inclusion/Exclusion

Cohesity supports the inclusion/exclusion of Nutanix AHV VM disks for backup. This feature allows users to include or exclude specific disks, improving backup time and preventing unnecessary disk backups.

To enable Disk Inclusion/Exclusion:

1. Click the Edit icon next to the protected object in the Add Object page.

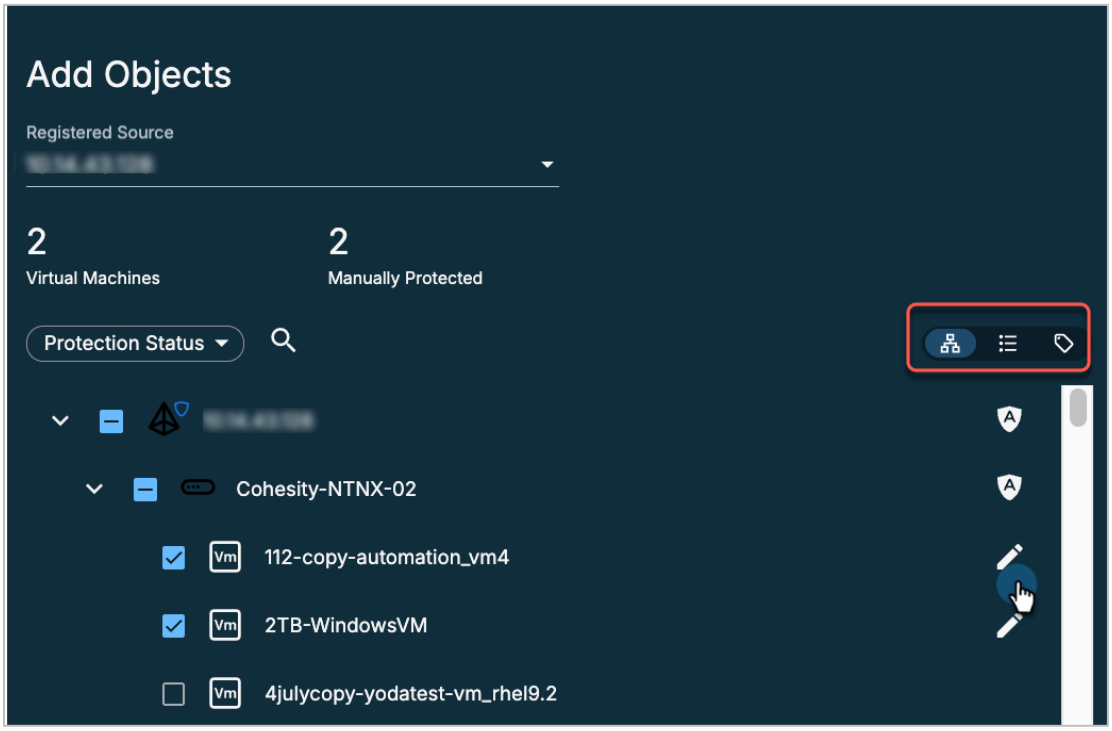

2. Toggle the **Enable Disk Inclusion/Exclusion** button to include/exclude disks.

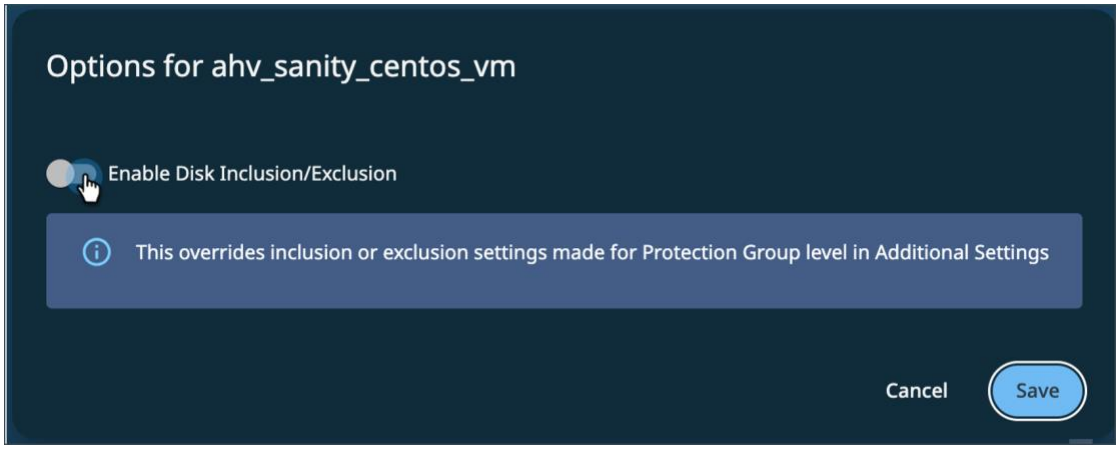

3. Click **Include** and select the disk from the drop-down menu to include the disk/disks in the backup. Click **Exclude** and select the disk from the drop-down menu to exclude the disk/disks from the backup.

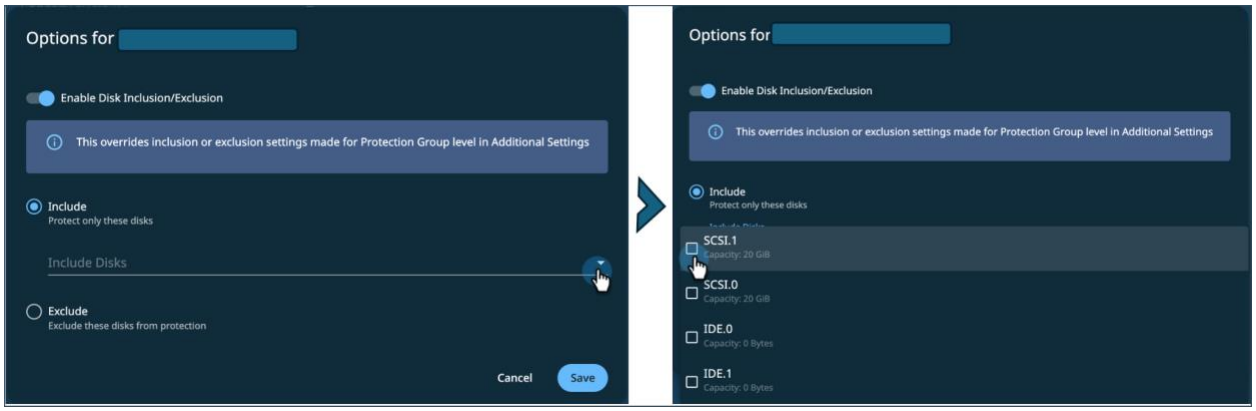

4. Click **Save**.

> **NOTE**: Starting release 7.3, Cohesity supports the backup and recovery of vTPM metadata for encrypted VMs. Cohesity backups up the vTPM metadata with the VM and does not require any additional configuration. The administrator should backup the vTPM user data manually.

Application Consistent Snapshots

Cohesity supports application-consistent backup for Nutanix Acropolis. By enabling application-consistent backups for a Protection Group, the guest operating systems of all the VMs in the Protection Group will be quiesced before snapshots of these VMs are created. Nutanix Guest Tools (NGT) must be installed on the AHV VM for application-consistent snapshots.

For App-Consistent Backups:

1. Toggle **App-Consistent Backups** in **Additional Settings** for a Protection Group if you want the guest operating systems of all the AHV VMs in the Protection Group to be quiesced before snapshots of these VMs are created.
2. If the App-Consistent backups toggle is on, the "**Take a Crash Consistent backups if unable to perform an App Consistent backup"** toggle is available. Toggle it on if you want the Cohesity cluster to capture a crash-consistent snapshot if it fails to capture an app-consistent snapshot.

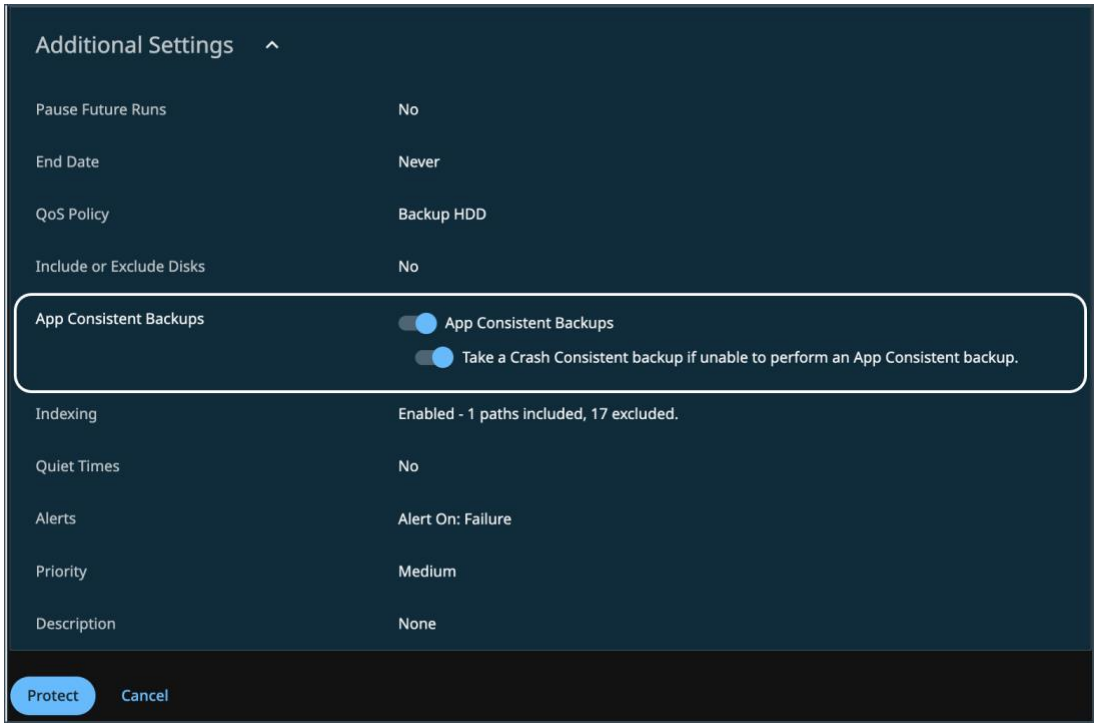

NOTE: Nutanix does not support app consistent backup of Volume Groups, hence, all the Volume Group Snapshots are crash consistent.

Support for Nutanix Cloud Cluster (NC2)

Cohesity Data Cloud integrates seamlessly with NC2 as a Nutanix source for AWS and Azure bare metal. This way, Nutanix users can extend their data protection and security capabilities to their data estate with minimal effort.

The AOS version supported by Cohesity applies to the NC2 cluster as well. Refer to the Supported Software page on Cohesity documentation for the list of supported AOS versions.

Registering an NC2 source is like registering a Nutanix AOS cluster. To register an NC2 cluster, refer to the Register or Edit a Hypervisor Source.

The instant recovery of Nutanix Cloud Clusters (NC2) AHV on Azure is supported.

Recover Nutanix AHV VMs, Files, and Folders

You can recover from Cohesity AHV backups at different granularity levels to both the original and alternate locations. You can recover entire VMs or search, locate, and recover individual files and folders.

The options to recover AHV VMs to different locations are:

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

You can also recover and download specific files and folders from snapshots created by a Cohesity Protection Group. The recovery task extracts the files stored in snapshots and creates new instances in their original location or a new location, depending on your options during recovery.

The options to recover or download individual files and folders are:

- Recover Files or Folders: Recover files or folders to their original or new locations. Cohesity Agent must be installed on the VMs for file and folder-level recovery.
- Download a File or Folder: Download a file or folder to your local machine from an existing snapshot.

Recovery Methods

Recovery from Cohesity AHV backups can be performed in two ways:

1. **Copy Recover**y: In this recovery method, the VMs will be available in the target location only after all the data is copied to the target storage from the source location (Cohesity cluster or cloud). Once the restore is complete, the VM will be available to use.
2. **Instant Recovery**: In this recovery method, the VMs will be instantly available after recovery in the target location. The data will be moved to the target storage later. The VM can be used in production during Instant Recovery while the data is copied to the VM from the Cohesity cluster. In such scenarios, the VM performance may be low as the Storage migration is still in progress.

Recover AHV VMs

You can recover your protected AHV VMs from a Cohesity cluster or a currently registered Cloud Archive. The recovery process is the same as recovering from the Cohesity cluster.

This restore task extracts the VM virtual disk files stored in snapshots. It creates new instances of the VMs in their original locations or a new location, depending on the options you choose during recovery:

Cohesity Data Cloud progresses through several processes to restore AHV virtual machines:

1. Create a new Volume Group on a Nutanix cluster and allocate the disks of the desired size.
2. Use iSCSI to log in to the new volume group from Cohesity Data Cloud to copy the data to the newly allocated disks.
3. Use the VM configuration recorded during backup to instantiate a new VM using the POST /VM endpoint.
4. Add the clone of the disks prepared in Step 2 to the restored virtual machine.
5. Delete the disks and volume group created in Step 1.
6. Power on the virtual machine (if configured).

Figure 4: AHV VM Restore Workflow

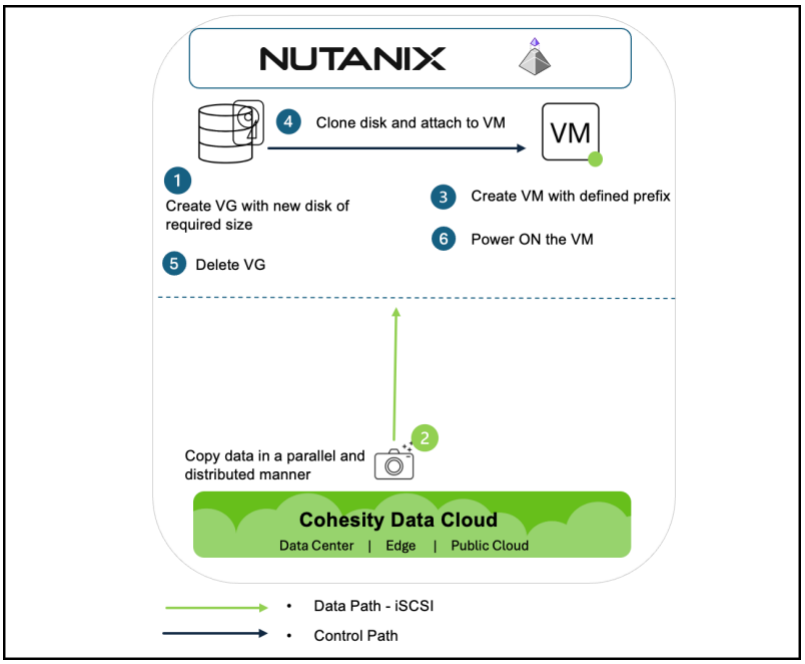

To restore your AHV VM:

1. Log in to Cohesity Data Cloud and select **Data Protection** > **Recoveries** > **Recover** > **Virtual Machines** > **VMs**.

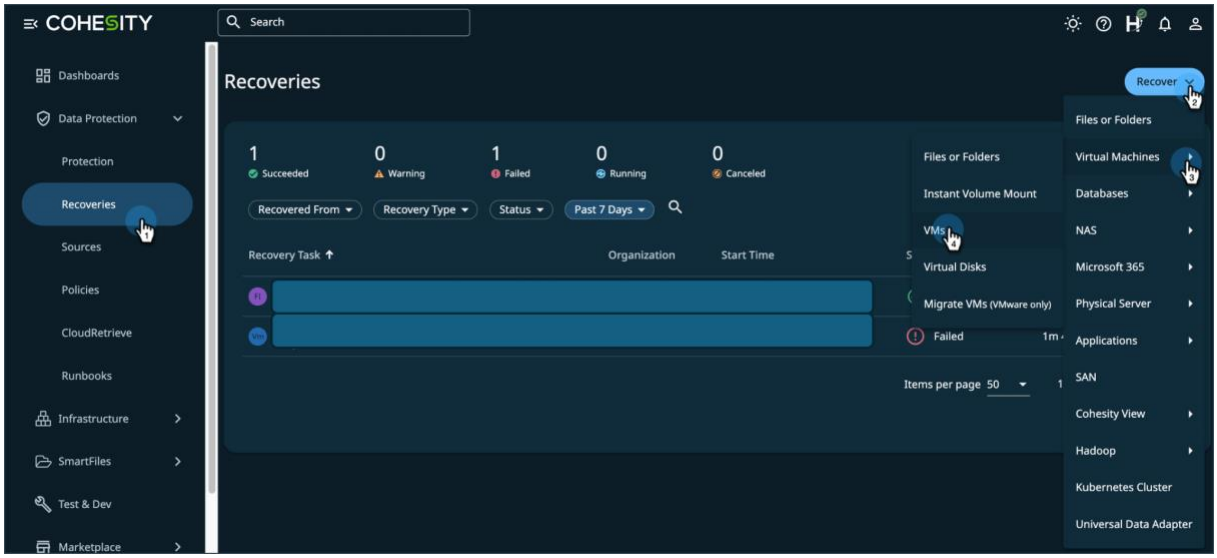

2. To recover the VMs, filter using the Source, Protection Group, Storage Domains or Date. Enter part or all the VM or Protection Group name on the screen. You can use the * wildcard to list all the VMs and Protection Groups. Select the VMs you need from the list or select an entire Protection Group to recover all the VMs, and then click **Next: Recover Options**.

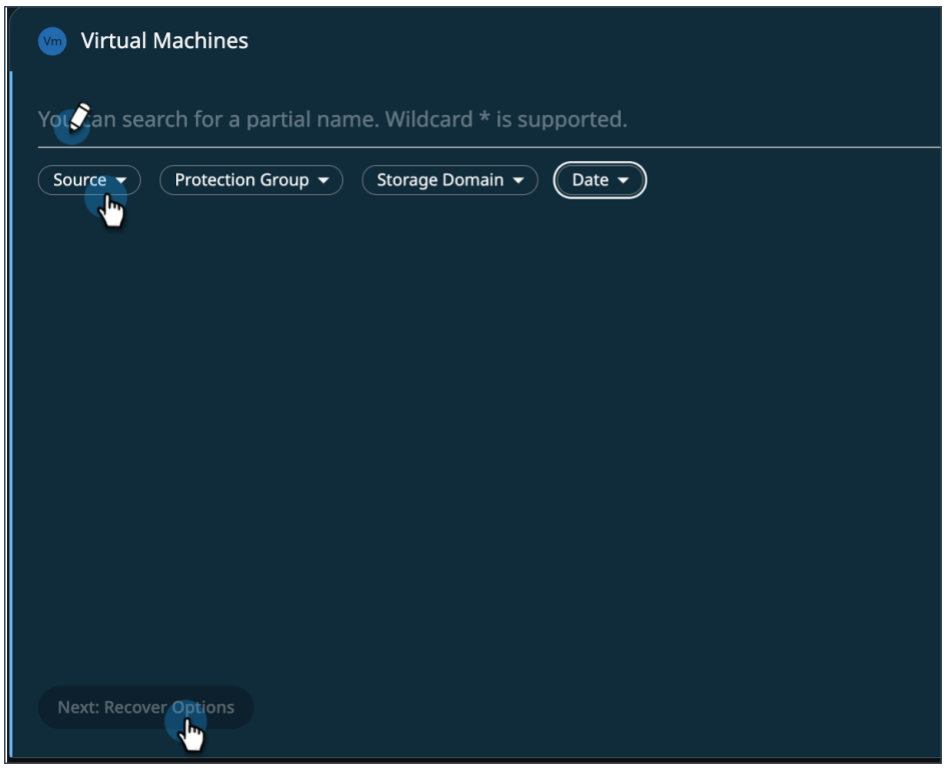

3. Select the location to recover the virtual machine. The Virtual Machine can be recovered to the **Original Location** or **New Location**. In the **Recovery Method,** select **Copy Recover** or **Instant Recovery**. Refer to the Recovery Methods section for details on the recovery methods.

   NOTE:

   - You can recover VMs from one Prism Central to another, as well as recover VMs within the same Prism Central instance, but on a different Prism Element (cluster).
   - You cannot perform mass instant recovery across different clusters. All VMs must be recovered to the same cluster for mass instant recovery or done individually.

   This page provides additional recovery options mentioned in the table below:

   Table 5: Recover Options

Recovery Options	Details
Create Empty Disks for Excluded Disks	Enable this option if you want Cohesity to create blank disks for every disk you excluded from the backup. The metadata of the excluded disks will be used to create the blank disks. These blank disks are available post-recovery and only have the metadata of the original disk, but no data. You can configure and run databases and applications on these disks, saving you from the effort of creating new disks. This feature is especially helpful if you have many VMs with excluded disks.
Network	By default, the VMs to be recovered do not have a virtual Network Interface Card (vNIC) attached. Enable the Attach option to attach a virtual Network Interface Card (vNIC) to each VM to be recovered. **Recover VM to Original Location:** With the Attach option enabled, each VM will connect to its original network and the following option is displayed: • **Start Connected**— Enable this option to connect to the original network when the VMs reboots. If this option is disabled, the VMs are not connected to any network on reboot. **Recover VM to New Location:** With the Attach option enabled, you need to select a network to attach the vNIC to a new network. Additionally the following options are displayed: • **Start Connected**— Enable this option to connect to the original network when the VMs reboots. If this option is disabled, the VMs are not connected to any network on reboot.

Recovery Options	Details
	• **Preserve MAC Address—** Enable this option to preserve the MAC address when recovering to an alternate location.
Recovery Directly Attached Volume Groups	Choose whether to restore the Volume Groups that were attached to the VM at the time of protection. Only available with Copy recovery option.
Rename	Add prefix and/or suffix strings to the names of the new VMs created by this task.
Power State	The Power Stare is set to Power On by default. Disable **Power On** if you want the recovered VMs to remain powered off after they are created.
Continue on Error	Enable **Continue recovery even if errors occur when recovering VMs** if you want the recovery task to continue even if errors occur when recovering VMs. For example, if one of the VMs cannot be created, the Cohesity cluster will still attempt to create the other VMs. If this option is disabled, the recovery task is discontinued if VM creation fails. However, if VM creation succeeds but VM storage migration fails, the task continues. By default, this option is disabled.
Cluster Interface	By default, the **Auto Select** option is enabled, and the recovery task automatically selects the correct VLAN. Select a configured interface group from the **Interface Group *** drop-down if you disable this option. This option is only available with Instant Recovery
Task Name	Change the default name of the recovery task.

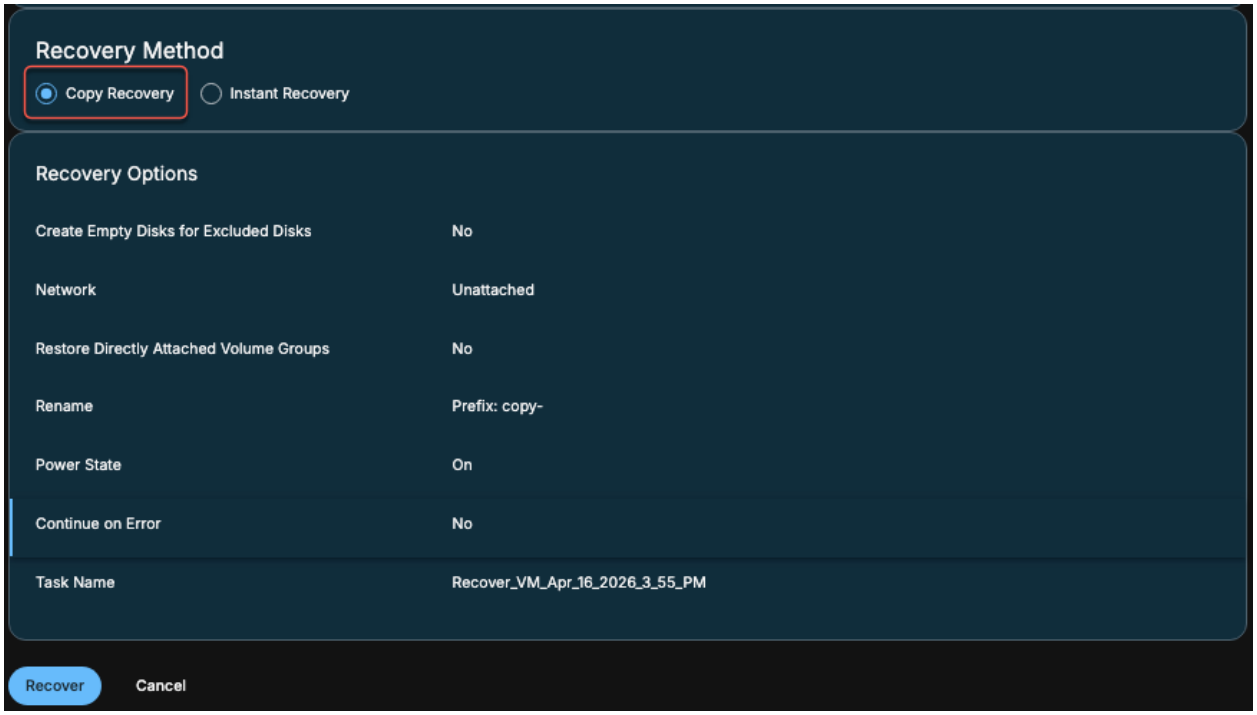

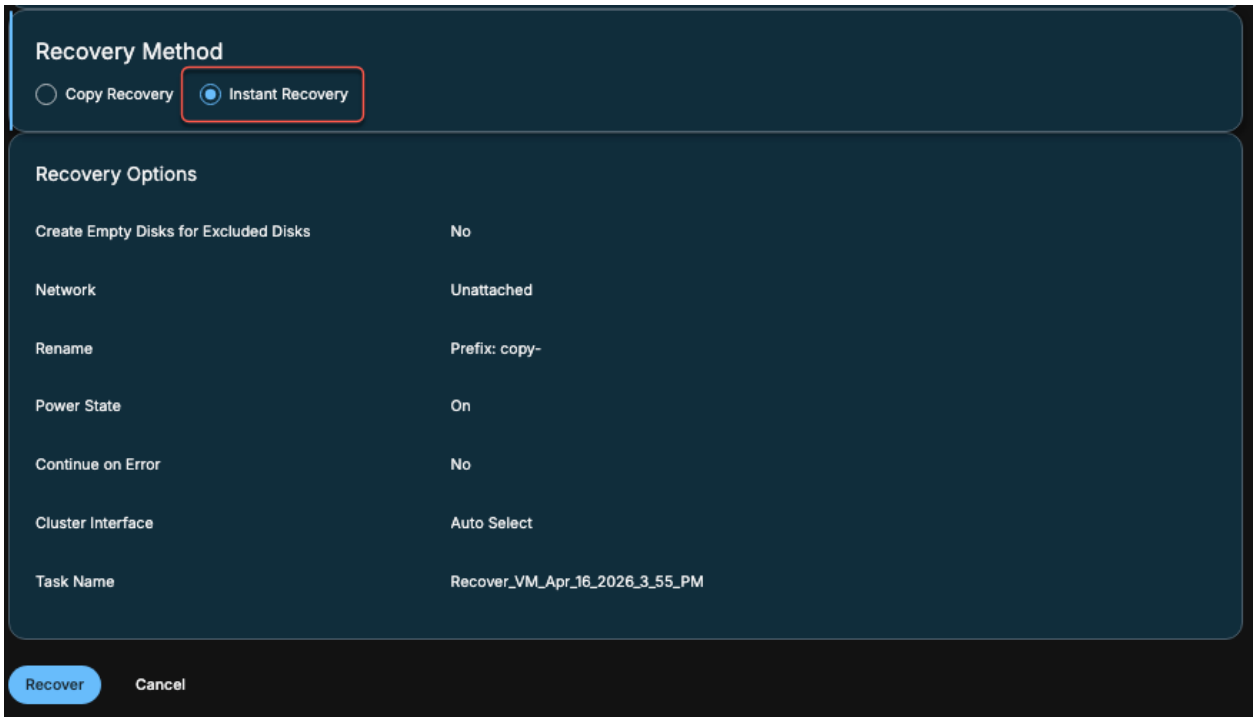

Click on each option to expand the **Recovery Options**.

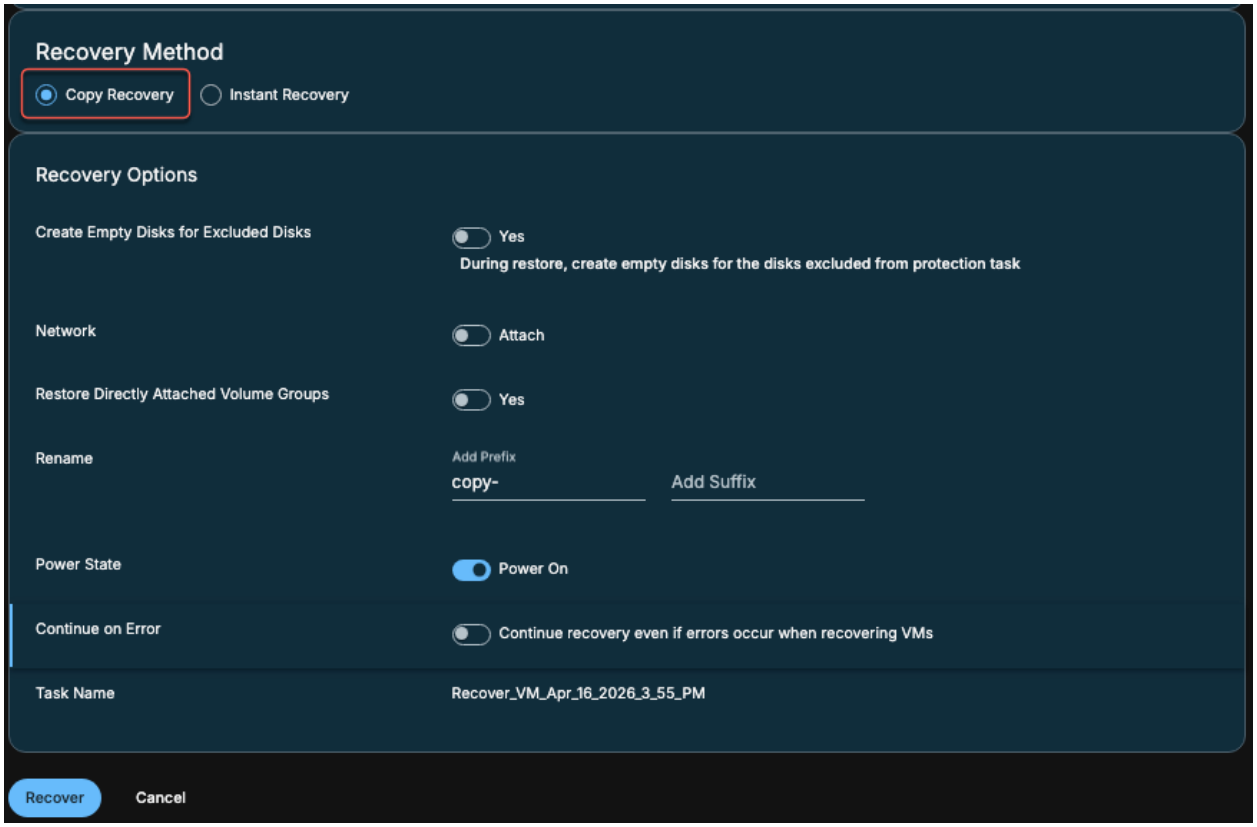

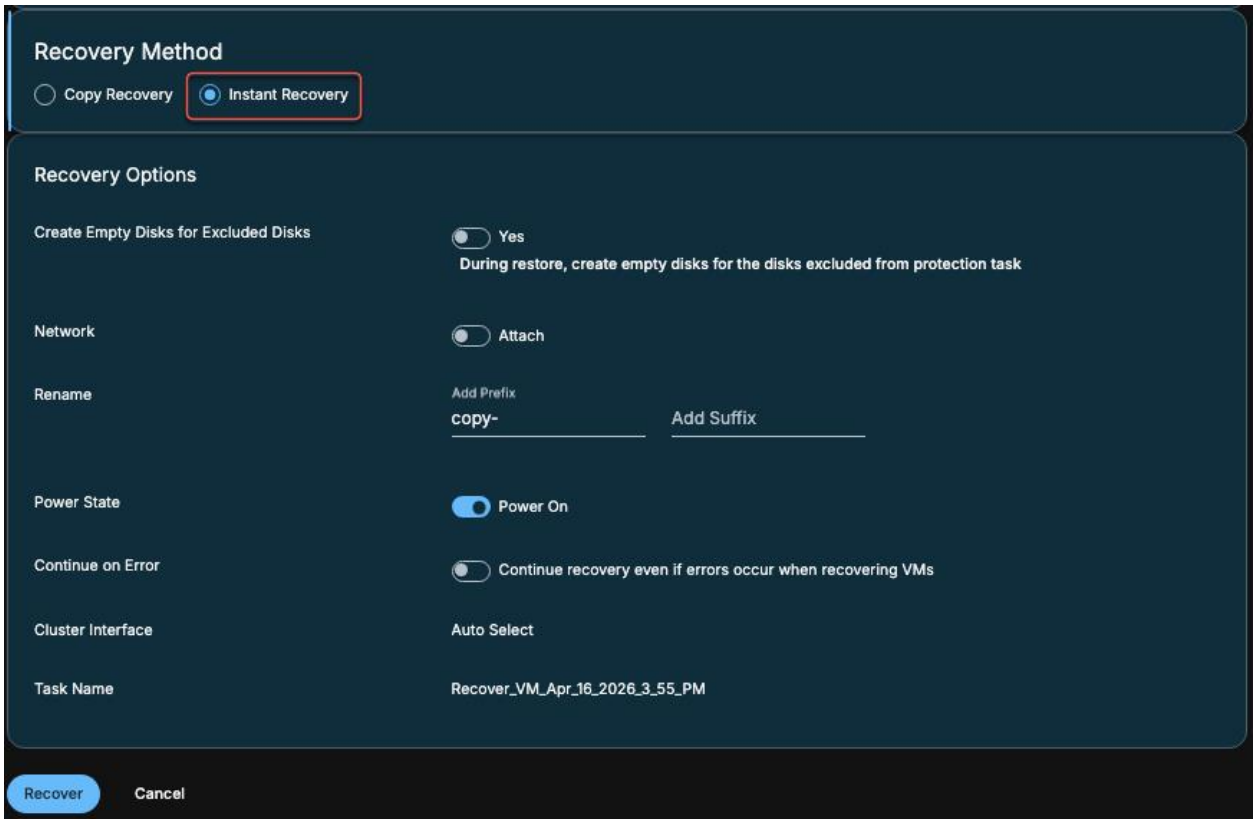

4. Click **Recover** to start the recovery process and monitor its progress. Once the recovery job finishes, the recovered VMs will appear in your Nutanix Acropolis cluster.

Restore Nutanix VM Files and Folders

Cohesity provides the ability to recover VM files and folders from a snapshot created earlier by a Protection Group. Files and folders can be recovered to their original location, a newly specified location, within the source, or a different one. You can choose to retain the original permissions and attributes of the recovered files and folders (at the time of the backup). You can also download files and folders from selected snapshots that a Cohesity Protection Group created.

File or folder-level recovery of VMs in a Nutanix Acropolis Hypervisor requires the installation of a Linux or Windows Cohesity agent on the target VM.

There are several tasks involved in restoring files and folders:

1. Mount the backup disk that contains the files and folders locally.
2. Find the disk and file system difference between the source and the target.
3. Parallel stream the differential data from the local mount point to the agent on the target VM.

NOTE: FileLevel Recovery (FLR) is supported for disks belonging to directly attached Volume Groups if:

- The disks were mounted in the guest OS at backup time (for example, listed in /etc/fstab on Linux)
- The mounted disks are supported by the Cohesity agent.

To restore your Nutanix files and folders:

1. Log in to Cohesity Data Cloud and select **Data Protection** > **Recoveries**. Click **Recover.** In the drop-down, select **Virtual Machines > Files or Folders.**

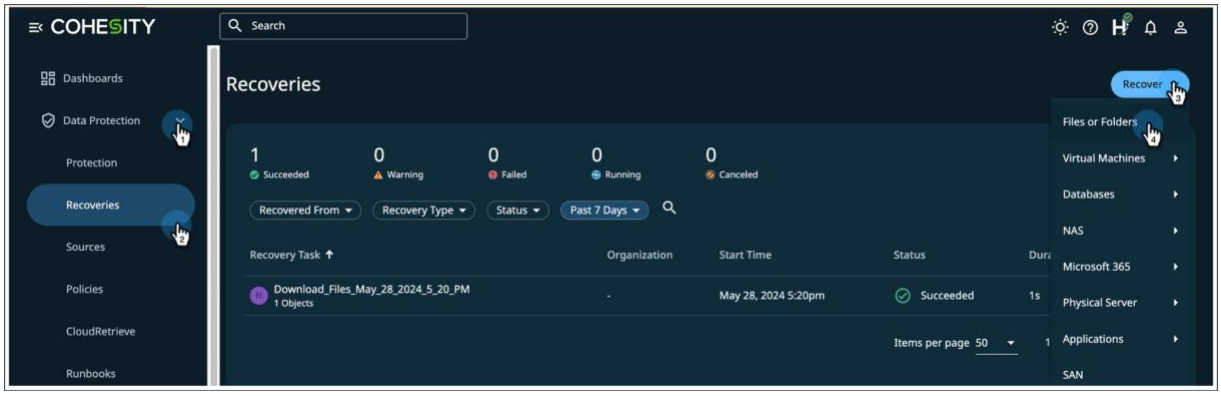

2. Search for and click the file or folder you need to restore. Click **Next: Recover Options** for recovery options or select **Downloa**d file to download the file.

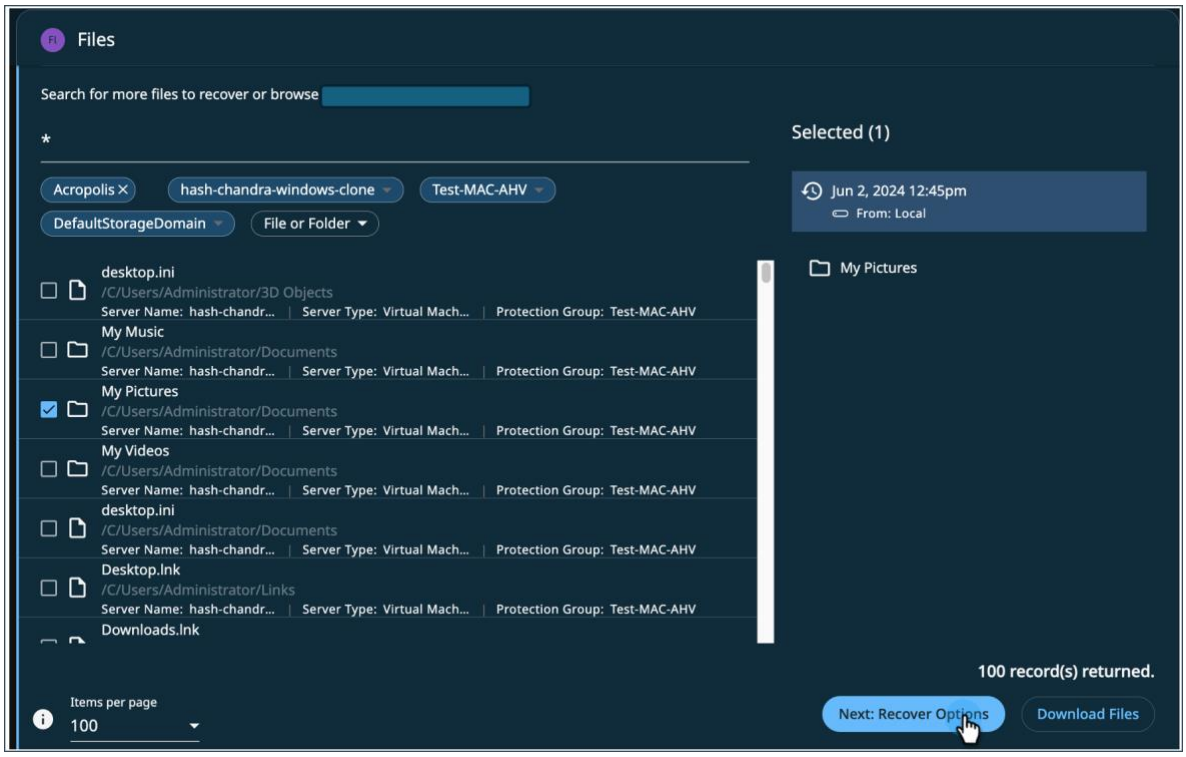

3. In the Recover Options dialog box, choose the location and additional Recovery Options and click **Recover.**

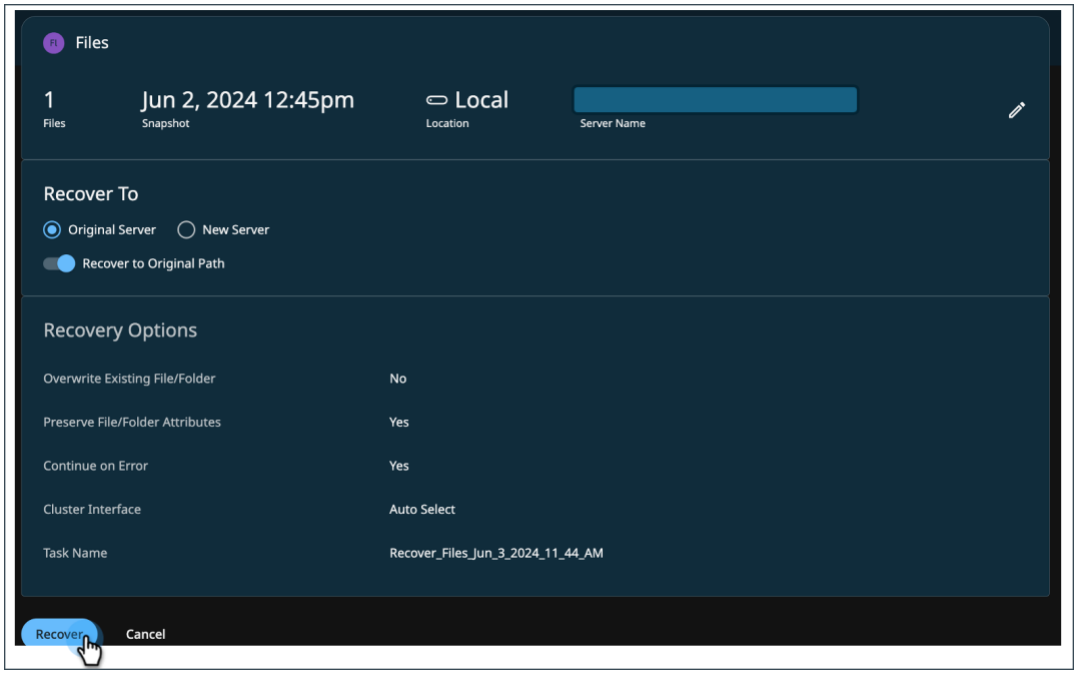

TIP: You can also save the selected file or folder to any new location from this window. To do so, click **New Server**.

For details, see Recover Files or Folders in the Cohesity documentation.

Use CloudArchive for Long-term Retention

In addition to backup and recovery, Nutanix Acropolis cluster administrators can use Cohesity CloudArchive to address long-term data retention requirements. Cohesity Data Cloud provides a policy-based method to archive any VMs or files to public clouds (AWS, Azure, GCP), S3-compatible storage, or NFS mount points. Cohesity CloudArchive's cloud-native integrations with AWS, Azure, and GCP eliminate the need for cloud gateways and point solutions to connect to the cloud while increasing operational efficiency and lowering the total cost of ownership (TCO). The archived data is efficiently transferred and stored by sending only deduplicated, compressed, incremental backups, reducing network and storage utilization.

Once your data is archived with CloudArchive, when you need to access it again, you can get it back using Cloud Recover (to your original cluster) or CloudRetrieve (to a new cluster).

1. **Cloud Recover to source cluster:** You can recover entire objects or individual files and folders to your original cluster.
2. **CloudRetrieve to new cluster:** For disaster recovery, geo-redundancy, and business continuity, retrieve your previously archived data onto an entirely new cluster.

Figure 5: Leverage Public Cloud Infrastructure for Long-term Data Retention and Archival

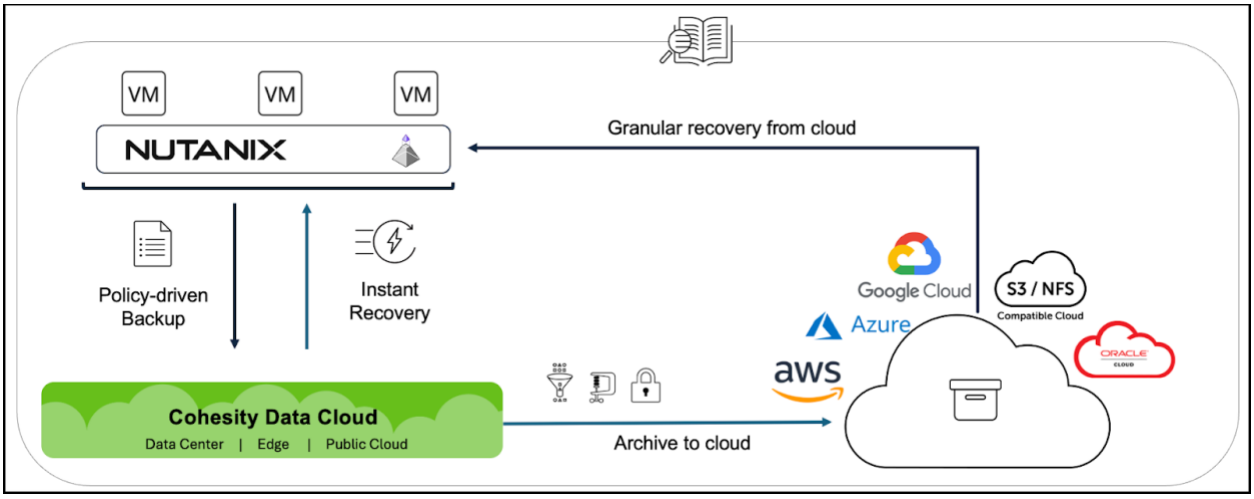

Cohesity CloudArchive automatically copies an existing snapshot created by Protection Groups in Cohesity Data Cloud and stores it on a registered External Target.

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

To learn how to archive your AHV VMs to cloud object storage, see the CloudArchive guides for AWS, Azure, GCP, NAS, and S3-Compatible cloud object storage. For more details on how to set up, check the Cohesity Cloud Solution Guides.

Access Your Cloud-Stored Data

Once the data is archived, Nutanix AHV administrators can also use Cloud Recover and CloudRetrieve features. Use:

- **Cloud Recover:** to recover VMs to your *original* cluster.
- **CloudRetrieve:** to retrieve the previously archived data onto an *entirely new cluster* as a cost-effective alternative for disaster recovery, geo-redundancy, and business continuity.

Figure 6: CloudArchive, Cloud Recover, and CloudRetrieve for Disaster Recovery & Geo-redundancy

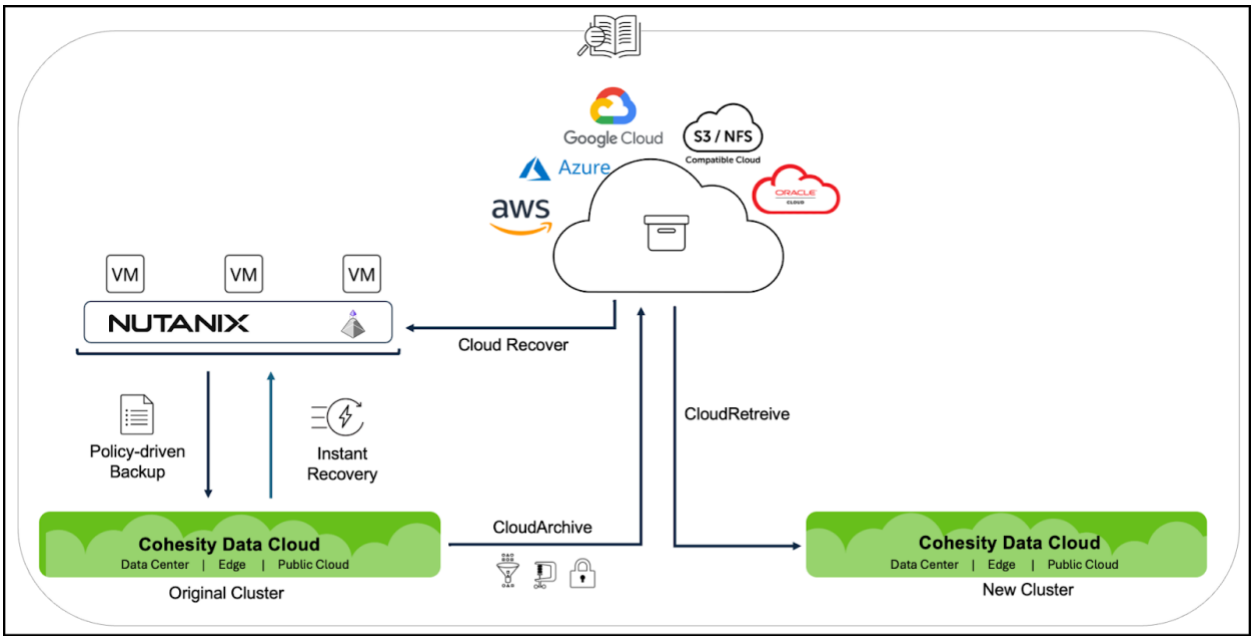

Cloud Recover and CloudRetrieve are very flexible. You can use them with AWS, Azure, GCP, NAS, and S3-Compatible cloud object storage.

Maintain Business Continuity with Disaster Recovery

Cohesity delivers a highly flexible, enterprise-grade data replication and disaster recovery (DR) solution that provides near-instant data recovery from core to edge to cloud and back as often as needed. This simplifies backup and DR operations on the same web-scale data management and application platform. Cohesity Data Cloud performs source-side deduplication and compression and sends only the changed data over the network as part of replication. If the primary site becomes unavailable, application and backup admins can fall over to the DR site for backup and recovery of their data.

Cohesity Data Cloud provides two mechanisms for protecting your data from disruptions and disasters.

1. **Replication**. This feature provides a simple way to store and retrieve data during major business disruptions, such as natural disasters and IT failures.
2. **CloudRetrieve**. CloudRetrieve works with CloudArchive to restore your data to an alternate Cohesity cluster.

Replicate Backups to Other Cohesity Clusters

Nutanix administrators can use Cohesity replication for cost-effective disaster recovery (DR). Cohesity provides a policy-based data replication solution from the core to the cloud to the edge, from one cluster to another cluster in your DR site. Cohesity performs source-side deduplication and compression as part of replication and sends only the changed data over the network. If the primary site becomes unavailable, application and backup admins can fail over to the DR site for backup and recovery of their data.

Figure 7: Replication Protects Nutanix Off-site

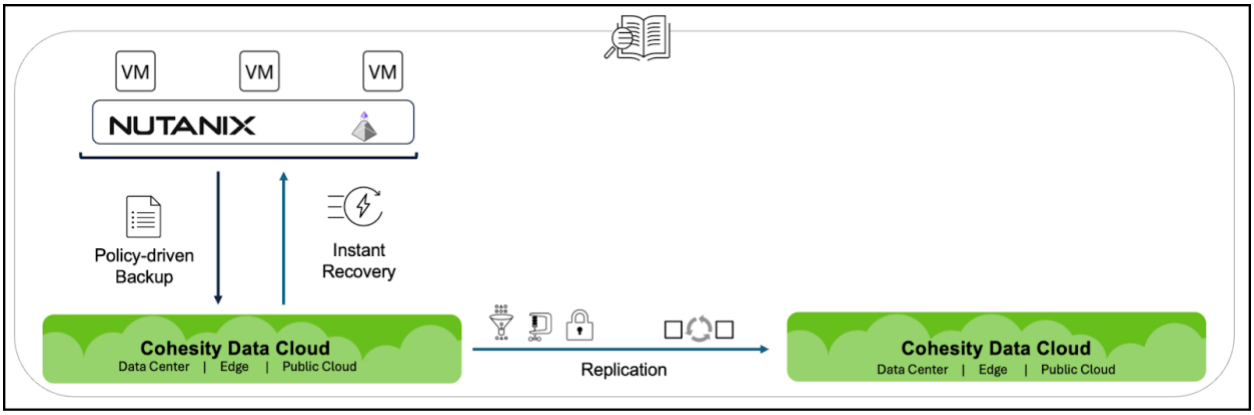

For more, see About Replication in the Cohesity documentation.

Monitoring

Monitoring events and reviewing reports are essential to making data-driven decisions about the day-to-day operations of any system in your data center.

Cohesity Data Cloud provides:

1. Performance monitoring of the Cohesity cluster, Storage Domains, and Nodes with respect to attributes such as CPU, memory, IOPS, throughput, and latency.
2. Detailed information about the health of your Cohesity cluster.
3. Several reports to understand cluster capacity and usage metrics and help you address your organization's specific needs most effectively.
4. Details about an event, such as date and time, category, type, and user, using Audit Logs.
5. Common pre-canned statistics in charts and graphs using Advanced Diagnostics.

Refer to the Monitoring section on the Cohesity documentation to know more about each topic.

Developer Extensions and Integration

Cohesity is built on an API-driven architecture. Our developer portal represents a brand-new developer experience, helping organizations customize, automate, and orchestrate secondary data and application workflows using the Cohesity REST API and the third-party tools of your choice.

The Cohesity Developer Portal provides enhanced developer experience to help you interactively build out REST API requests and then copy them over for further use.

Cohesity provides a PowerShell Module and Python for interacting with Cohesity Data Cloud. The PowerShell Module can be used on Windows, Linux, or macOS. It includes cmdlets that are useful for automating common tasks and orchestrating workflows in your environment.

The Cohesity Management SDK provides an easy-to-use language binding to harness the power of Cohesity REST APIs in your Python applications.

For more on Cohesity integrations, see Enable Infrastructure Automation with Cohesity and Ansible and Enable Self-Service Infrastructure Automation with Cohesity and ServiceNow.

Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Sadik Sayed is a Snr. Solution Architect at Cohesity. In his role, he focuses on Virtualization Data Protection – VMware vSphere, VMware Cloud Director, VMware Cloud Foundation, Microsoft HyperV and Nutanix AHV.

Other essential contributors include:

- Gautam Bhasin, Dir, Product Management
- Anand Arun, Technical Director
- Ashish Patwardha, Product Manager
- Adaikkappan Arumugam, Sr Director Product Solutions

Document Version History

VERSION	DATE	DOCUMENT HISTORY
3.3	April 2026	• Support for Volume Group backup and recovery • Cap Concurrent streams per datastore
3.2	Feb 2026	Restore Network Configuration and screenshot updates
3.1	Nov 2025	Support for VM Tags/Categories and vTPM metadata backup.
3.0	Sep 2025	Support for Prism Central
2.0	July 2024	New Release
1.0	Dec 2019	First Release

ABOUT COHESITY

Cohesity is a leader in AI-powered data security and management. Aided by an extensive ecosystem of partners, Cohesity makes it easier to protect, manage, and get value from data – across the data center, edge, and cloud. Cohesity helps organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring for malicious behavior, and rapid recovery at scale. Cohesity solutions are delivered as a service, self-managed, or provided by a Cohesity-powered partner. Cohesity is headquartered in San Jose, CA, and is trusted by the world's largest enterprises, including six of the Fortune 10 and 44 of the Fortune 100.

Visit our website and blog, follow us on Twitter and LinkedIn, and like us on Facebook.