COHESITY

CloudArchive & CloudRetrieve Deployment & Recovery Guide for S3-Compatible Object Storage

Use S3 Buckets to Store Your Protected Data for Long-Term Retention and Disaster Recovery

Version 2.0

December 2025

ABSTRACT

Long-term data and application retention is critical to organizations that seek to prevent data loss and meet security, legal, and compliance requirements. Cohesity Data Cloud™ and DataProtect™ offer robust on-premises solutions for enterprise data protection and storage. Cohesity's CloudArchive™ and CloudRetrieve™ bring data protection and recovery together with any S3-compatible storage, in the cloud or on-premises.

Table of Contents

Figures

Tables

CloudArchive Connects S3-Compatible Storage to Cohesity Data Cloud

Long-term data and application retention is critical to organizations that seek to prevent data loss and meet security, legal, and compliance requirements. Cohesity Data Cloud (previously known as “Cohesity Platform” and hereafter referred to as “Cohesity Data Cloud”) offers robust on-premises solutions for enterprise data protection and storage. Cohesity’s CloudArchive and CloudRetrieve bring data protection and recovery together in a single coherent solution, both on-premises and in the cloud, with any S3-compatible storage.

Figure 1: CloudArchive Connects S3-Compatible Cloud and On-Prem Storage to Cohesity Data Cloud

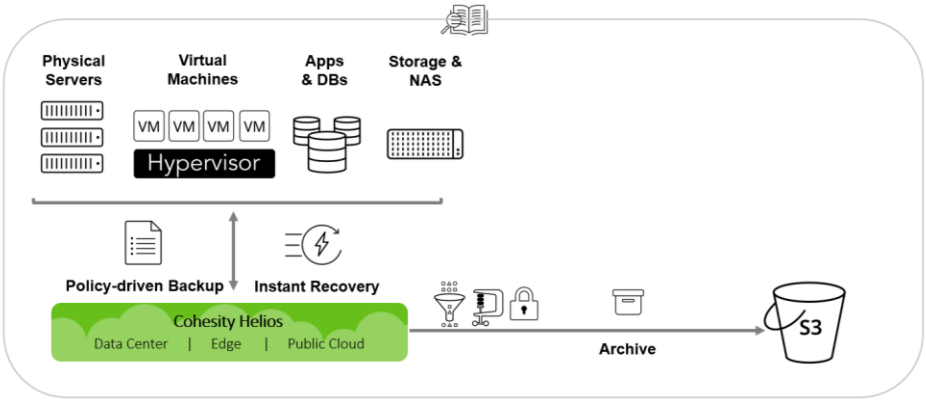

With Cohesity, IT organizations save time by quickly archiving data to multiple targets—public clouds, private clouds, any S3-compatible device, as well as NAS-NFSv3 from storage vendors, and QStar managed tape libraries, eliminating the need for cloud gateways and point solutions to connect to the cloud, while increasing operational efficiency and lowering total cost of ownership (TCO).

NOTE: This document covers only Cohesity Data Cloud operations for archiving to S3-compatible storage, on-premises, and in the cloud, and not tape, or Object-based Tape or NAS targets.

For archiving to tape, see Long Term Retention to Tape with Cohesity DataProtect solution guide. For archiving to NAS, see CloudArchive & CloudRetrieve Deployment & Recovery Guide for NAS. For archiving to public cloud vendors, see guides for AWS, Azure, and Google Cloud Platform.

CloudArchive Versions

Cohesity CloudArchive has two versions:

- CloudArchive Incremental with periodic full
- CloudArchive Incremental forever

Before you configure CloudArchive; Review the differences between the versions and the supported sources

CloudArchive Features and Benefits

Cohesity's CloudArchive provides many key features, each of which delivers several benefits to organizations and their IT administration staff. Specifically:

Table 1: CloudArchive Features and Benefits

FEATURE	BENEFITS
Policy-based archival	• Easy to use. • Archive unique data differently by mapping Protection Policies to the required SLA. • Reduce bandwidth and storage costs.
Off-site copies	• Geo-redundancy • Disaster recovery
Deduplication and compression	Efficient data transfer and storage.
Granular recovery	• Instantly locate VMs, files, and folders. • Recover just what you need.
Encryption	Data is secure both in flight and at rest.

Classes of Supported S3-Compatible Object Storage – Standard Tier

Cohesity supports any S3-compatible device as an S3-Standard tier (Regular as storage class) external target with similar capabilities to an AWS S3-Standard. Any S3-compatible targets that use S3-Standard must be registered as a regular storage class. The S3-compatible uses S3-Standard APIs for archival and recovery. This Standard tier is supported for both incremental with periodic full and incremental forever archival formats. Compression and deduplication are supported for both archival formats.

With incremental forever archival format, the Cohesity cluster will download a portion of the data from the S3-compatible Standard tier external target for the space reclamation process and reupload it after compaction. This leads to increased network bandwidth usage. Cohesity recommends having sufficient network bandwidth between the cluster and the external target; for example, hosting the external target within the same data center or network as the Cohesity cluster.

Review the support matrix to understand storage classes for CloudArchive: External target support matrix.

CloudArchive Terminology

It is important to understand the following terms as you learn about how CloudArchive works.

Table 2: CloudArchive Terminology

TERM	DEFINITION	NOTES
Cohesity Data Cloud	Cohesity Data Cloud consolidates secondary data and applications, including backups, files, objects, test/dev, and analytics on a single, software-defined platform. Inspired by web-scale architecture. Cohesity Data Cloud is a scale-out solution based on a unique distributed file system, SpanFS®.	
Archive	Completely self-contained copy of the backup (data and metadata) that is stored outside the Cohesity cluster.	
Archive Chain	The set of a Full Archive and the Incremental Archives that depend on it and the preceding Incremental.	If the Full Archive is lost for any reason, the entire archive chain becomes unusable. If an Incremental Archive is lost, the restore points that follow it are lost as well.
CloudRetrieve	The process of retrieving an archived Protection Group and its Protection Group Run details from an External Target to a different cluster. Used for geo-redundancy and disaster recovery.	CloudRetrieve cannot be performed on the same cluster that performed the archive.
Cluster	An instance of Cohesity Data Cloud.	
Deduplication Chain	The set of a Reference Archive and all the archive chains that depend on it for deduplication. This includes the Scheduled Full and Incremental Archives for each archive chain in the deduplication chain.	These dependencies determine when Cohesity Data Cloud can retire and eventually delete Reference Archives.
External Target	Any storage to which data is sent outside the source Cohesity cluster.	Archive to Cloud, Tape, NAS, and replication targets are all External Targets in Cohesity Data Cloud.
Full Archive	A full copy of the Protection Group that is archived.	

TERM	DEFINITION	NOTES
Incremental Archive	An archive that records just the changed data since the most recent archive.	
Protection Group	Defines operational requirements, such as which source objects to protect, the Protection Policy to use, and operational considerations like indexing, exclusions, and inclusions, alerts, app consistency, and more.	Each Protection Group has a schedule of Protection Group Runs, and each archive is a collection of those Protection Group Runs.
Protection Policy	Reflects business needs of Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs) by defining the frequency and retention requirements of backup, archival, and replication.	
Scheduled Full Archive	A Full Archive that runs at regular intervals (configurable, 90 days by default).	The Scheduled Full Archive does not send the same amount of data, as it is deduplicated against the Active Reference Archive. In those cases when there is no Active Reference Archive, the data sent for the Scheduled Full is deduplicated only with itself and not against any other archive. For example, if the Active Reference Archive size is 100GB and the Scheduled Full deduplication usage is 60%, then only 40GB is sent. If there is no Active Reference Archive, then the size of the Scheduled Full is 100GB.
Source-Side Deduplication	The process of eliminating redundant copies of data to reduce storage use before sending over the network. Redundant data blocks are replaced with a pointer to the unique data copy, which ensures that only unique instances of data are transferred over the network and retained on storage media.	Reduces storage as well as network bandwidth requirements and, in doing so, saves time and money.

TERM	DEFINITION	NOTES
Recover	Retrieve an entire data object, such as a VM or database, or granularly recover files and folders from an External Target onto the original cluster.	
Reference Archive	The Full Archive against which all subsequent Incremental Archives (in the archive chain) and Scheduled Full Archives as well as their Incrementals are deduplicated.	All Reference Archives are full archives. A new Reference Archive is created when Cohesity detects that deduplication with it is below 50%. **NOTE**: 50% is the default threshold. This is internally configurable, but changing this value only delays *when* (and not whether) the full data set is sent.
Retired Archive	A Reference Archive that is no longer used for deduplication.	

CloudArchive High-Level Workflow

At the highest level, leveraging CloudArchive involves several sequential tasks:

1. Create an S3 bucket with the S3-compatible object storage provider of your choice.
    a. Create a user and assign the necessary permissions to the object storage for Cohesity Data Cloud to access it. Capture the Access Key ID and Secret Access Key to register.
2. Register your S3 bucket to Cohesity Data Cloud as an External Target.
3. Archive your data to the S3 bucket.
    a. Create a Cohesity Protection Policy with a CloudArchive configuration.
    b. Create a Cohesity Protection Group.
4. Recover your data from the S3 bucket.

Figure 2: Leverage S3-Compatible Storage with Cohesity

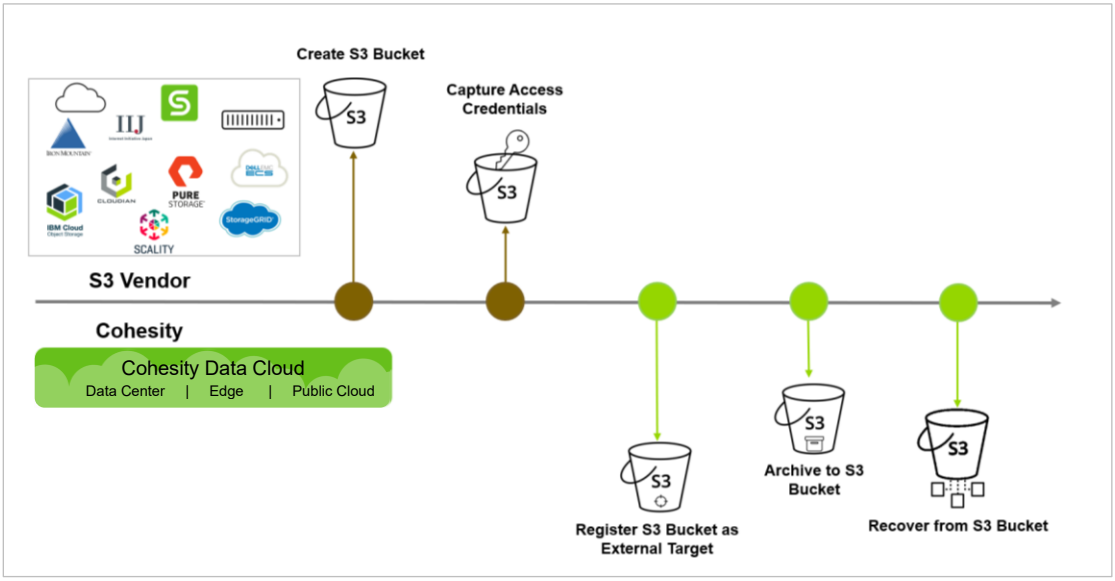

NOTE: Any S3-compatible targets that use Tape storage in the backend are not supported with the Standard tier; instead, you must register them as S3-compatible with the Tape-based Storage Class.

Create Your S3 Bucket

The first thing you'll do is create a bucket with your S3-compatible storage vendor. Though the process is slightly different for each vendor, it always involves creating an S3 bucket and a user account that has access to it. Finally, you'll need to capture the Access Key ID and Secret Access Key that gives that account access.

Figure 3: Create Your S3 Bucket

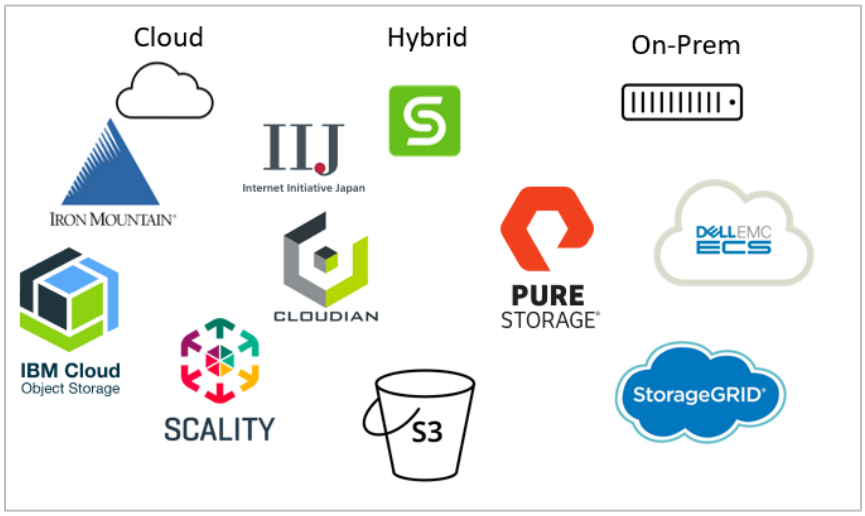

Connect Your S3 Bucket

Next, you need to connect the new S3 bucket to Cohesity by registering it as an External Target in Cohesity Data Cloud. For this, you'll need the bucket name, access key ID, secret access key, endpoint, port and geographic region (optional).

Figure 4: Register S3 Bucket with Cohesity

Archive Your Data to Your S3 Bucket

With your S3 bucket now registered with Cohesity Data Cloud, the next step is to archive your data by creating a Protection Policy (which reflects your business needs, like frequency and archival retention requirements) and running a Protection Group (where you define operational requirements, such as which data objects to protect, the Protection Policy to use, indexing, alerts, and SLA requirements).

Figure 5: Archive Data to S3 Bucket

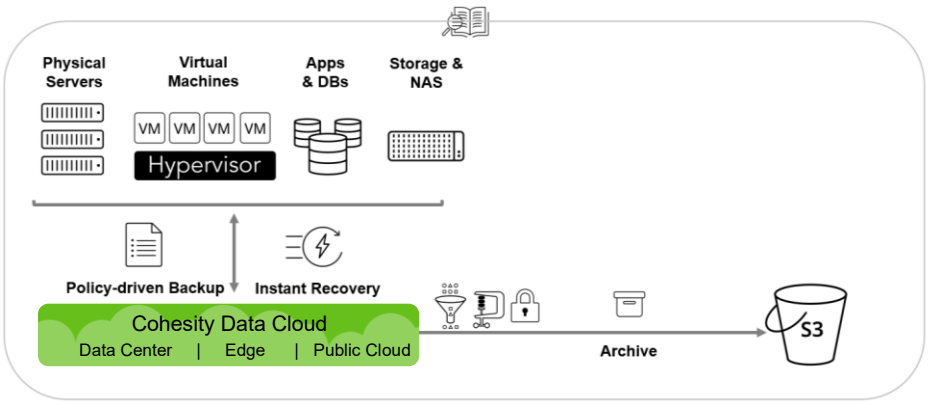

Recover Your Data from Your S3 Bucket

In most organizations, customers use on-premises storage for data that only has a short retention period, but store data with long-term retention requirements in the cloud. When a need for some or all of that cloud-stored data arises, the challenge is to locate, identify, and recover it quickly and reliably.

Cohesity Data Cloud includes an indexing engine that enables rapid search and recovery of files and virtual machines from archives stored both on-premises and in the cloud. As virtual machines and physical servers are backed up, Cohesity's indexing engine opens the underlying files and indexes the metadata. This enables extremely fast, wild-card search results that are then used for granular recovery.

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

- **Cloud Recover to source cluster**: Recover entire objects (VMs, databases, NAS, etc.), or individual files and folders, to your original cluster.
- **CloudRetrieve to new cluster**: Retrieve your previously archived data onto an entirely new cluster, for disaster recovery and geo-redundancy.

Figure 6: Recover Data from Your S3 Bucket—Cloud Recover and CloudRetrieve

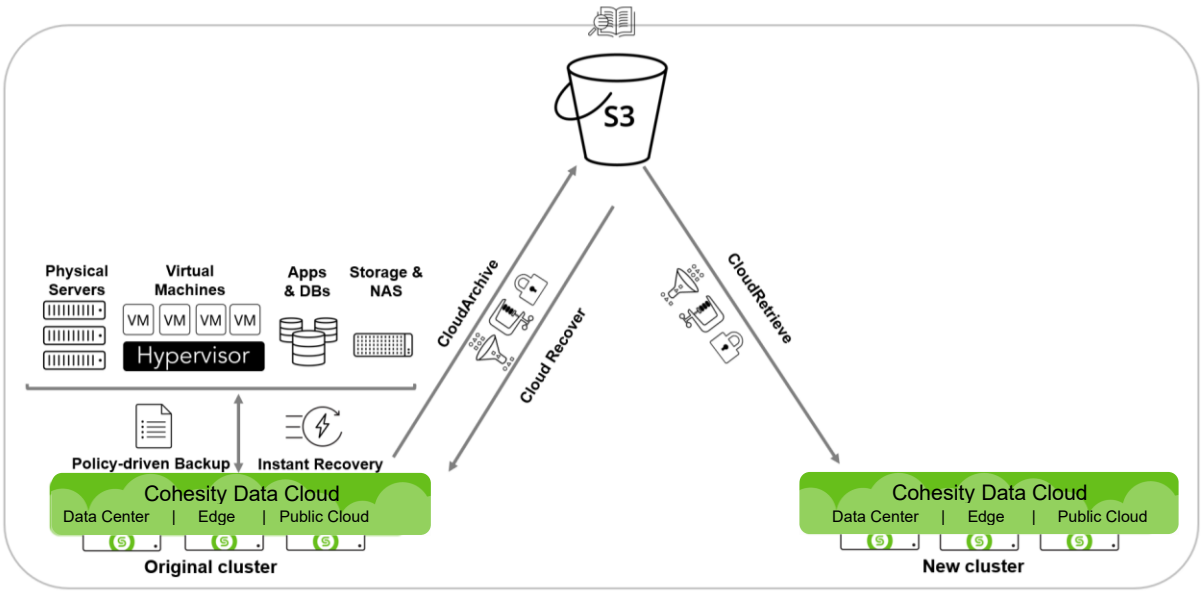

In the next chapter, we cover the individual steps that are involved in each of these tasks. Following that, we walk through the specific procedures for connecting your S3 bucket to Cohesity Data Cloud, archiving your data to your S3 bucket, and recovering and restoring your data from the S3 bucket.

Leverage Your S3-Compatible Storage with Cohesity

This chapter provides a quick overview of the sequence of actions that you will be undertaking to use your S3 bucket as an External Target in Cohesity Data Cloud. In the chapters that follow, you'll find step-by-step instructions for each CloudArchive feature you can use with your S3 bucket.

Create and Register S3 Bucket

Start by setting up your S3 bucket. Note that the same S3 bucket can be registered multiple times in the same cluster as different External Targets, as well as in multiple clusters.

1. Create an S3 bucket and an account that can access it with the S3-compatible storage vendor of your choice.
2. Capture the bucket Name, Access Key ID, and Secret Access Key.
3. Using the S3 bucket information and access credentials, register the S3 bucket as an External Target in Cohesity Data Cloud.

> **IMPORTANT**: Customers should never manually edit, change, or delete Cohesity Data Cloud archives directly in the S3 bucket.

When you register your External Target in Cohesity Data Cloud, you will be able to enable or disable:

Table 3: External Target Options

FEATURE	DESCRIPTION
Encryption	By default, Cohesity Data Cloud writes the data into External Targets in encrypted format in real time. You can disable it, but Cohesity recommends you leave it enabled in almost all cases, except when the data is already encrypted. You can choose to keep your encryption key in the S3 bucket with your archive, or, for additional security, to manage it manually. **NOTE**: If you choose the manual option, you will need to download the key after registering the External Target and store it outside the Cohesity cluster.
Compression	Reduces the impact on data transfers and data storage. Useful except when the data format doesn't compress well, such as with databases and large image files.
Source-Side Deduplication	The process of eliminating redundant copies of data to reduce storage use. Redundant data blocks are replaced with a pointer to the unique data copy, which ensures that only unique instances of data are sent across the network and retained on storage media, and dramatically reduces the impact on bandwidth and storage utilization. Cohesity strongly recommends it in all cases.

FEATURE	DESCRIPTION
Incremental Archival	An archive that records just the changed data since the most recent archive. This allows you to return to any restore point without having to create, transfer, and keep a backup copy of your whole dataset each time. Cohesity strongly recommends this setting in all cases. If this option is not enabled, it will send a full archive on every archive run.
Bandwidth Throttling	If needed, you can throttle the upload and download bandwidth that is consumed by network traffic between Cohesity Data Cloud and an External Target. You can also limit bandwidth throttling to a specific time range, if there are particular days and times when it is needed. **NOTE**: You cannot set Bandwidth Throttling to lower than 1Mbps.

NOTE: The same S3 bucket can be registered multiple times in the same cluster as different External Targets, as well as in multiple clusters.

Required S3 Vendor Fields

To register your S3 bucket as an External Target, Cohesity Data Cloud requires the following fields:

- Bucket Name
- Access Key ID
- Secret Access Key
- Endpoint

NOTE: You'll also need to know the **port** on which your S3 bucket is exposed, and whether your bucket is compatible with *AWS Signature Version 2 or 4*.

Configure Your Policy-based Archive

Once Cohesity registers your S3 bucket as an External Target, you will create a Protection Policy to define your business needs. The Protection Policy allows you to incorporate the S3 External Target that you created earlier as an archive target with a specific retention period.

In the Policy, you configure how virtual and physical servers, databases, and unstructured data are protected:

- Backup frequency and retention period.
- Whether to have your backups archived, how often, and how long to retain.

  NOTE: You can add more than one archival schedule to the same Policy, and you can use the same or a different External Target, with the same or different frequency and retention.

- Which External Target to use (in this case, your newly registered S3 bucket).

Protect Your Data

Protection Groups combine operational requirements with the business requirements that are defined in a Protection Policy.

In the Job, you select the source, which data objects from that source to store, the Protection Policy and the storage domain (the named storage location) to use, and operational details such as Start Time, End Date, QoS Policy, Pre & Post Scripts, and more. See the all the advanced Protection Group settings in the Appendix A.

Once you save a Protection Group, it will run on the schedule you define.

> **NOTE**: Multiple Protection Groups can use the same Protection Policy, but each Job can have only one policy.

Recover Data from Your Archive

When the time comes to recover your archived data, Cohesity gives you three options:

- Restore entire data objects (VMs, databases, NAS, etc.).
- Recover individual files and folders.
- Retrieve your data onto an entirely new Cohesity cluster (for disaster recovery, etc.).

For instructions, see Recover Data from CloudArchive below.

Connect S3-Compatible Storage to Cohesity Data Cloud

Cohesity's CloudArchive enables customers to connect seamlessly to any vendor of S3-compatible storage (like IBM Cloud Object Storage™, NetApp StorageGRID®, Iron Mountain®, and more) as an extension of the data center infrastructure. Customers are using CloudArchive to reduce their reliance on tape for cost-effective, long-term data retention, as well as a low-cost, disaster-recovery solution.

Figure 7: Cohesity CloudArchive Works with S3-Compatible Storage

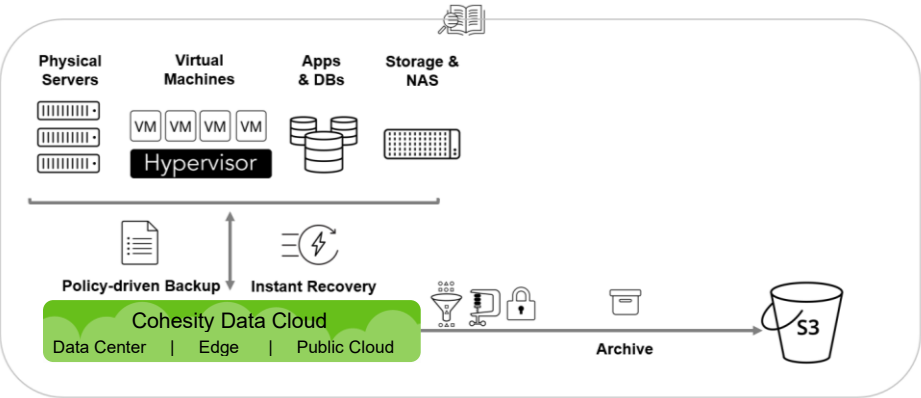

Let's get started!

Create Your S3 Bucket for CloudArchive

Cohesity supports any S3-compatible storage. Choose your S3 vendor and follow their procedure for creating a bucket for storing your data.

Enable Bucket versioning and Object Lock feature only if you are planning to use WORM/Object Lock feature in CloudArchive. Review the prerequisites and considerations for Archive Object lock.

> **NOTE**: The WORM/Object Lock feature can be utilized starting from Cohesity version 7.2 onwards. Review the support workflow matrix : S3-Compatible with Periodic Full and Incremental Forever.

Remember to make note of:

- Bucket Name
- Access Key ID
- Secret Access Key
- Endpoint
- Port on which S3 bucket is exposed
- AWS Signature version: Does your S3 bucket use AWS v2 or v4?

Register S3 Bucket with Cohesity

Now that you have the bucket that you need, you're ready to connect it to Cohesity Data Cloud (whether your cluster is on-premises, Cloud Edition, or Virtual Edition).

To register an External Target with your cluster:

1. Log in to **Cohesity Data Cloud**.
2. Click **Infrastructure** > **External Target**.

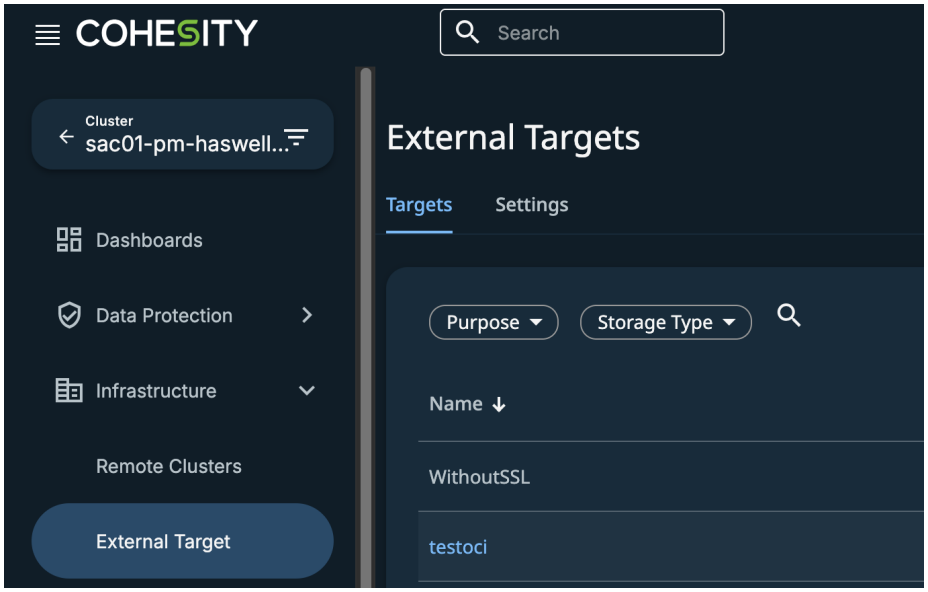

3. Then click **Add External Target**.

4. Select the purpose as **Archival**, Storage Type as **S3Compatible** and Storage Class as **Regular**.

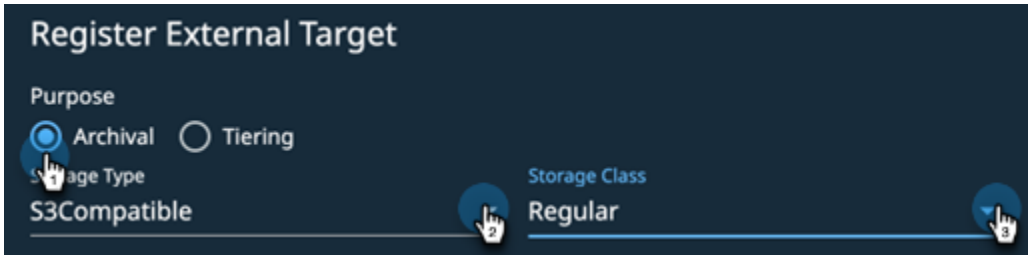

5. In the form that opens:
   a. Enter **Bucket Name**.
   b. Enter **Access Key ID**.
   c. Enter **Secret Access Key**.
   d. Enter the **Endpoint IP/FQDN**.

> **NOTE**: Do not include 'http://' or 'https://' in the mount path.

   e. Enter the **Port** number.

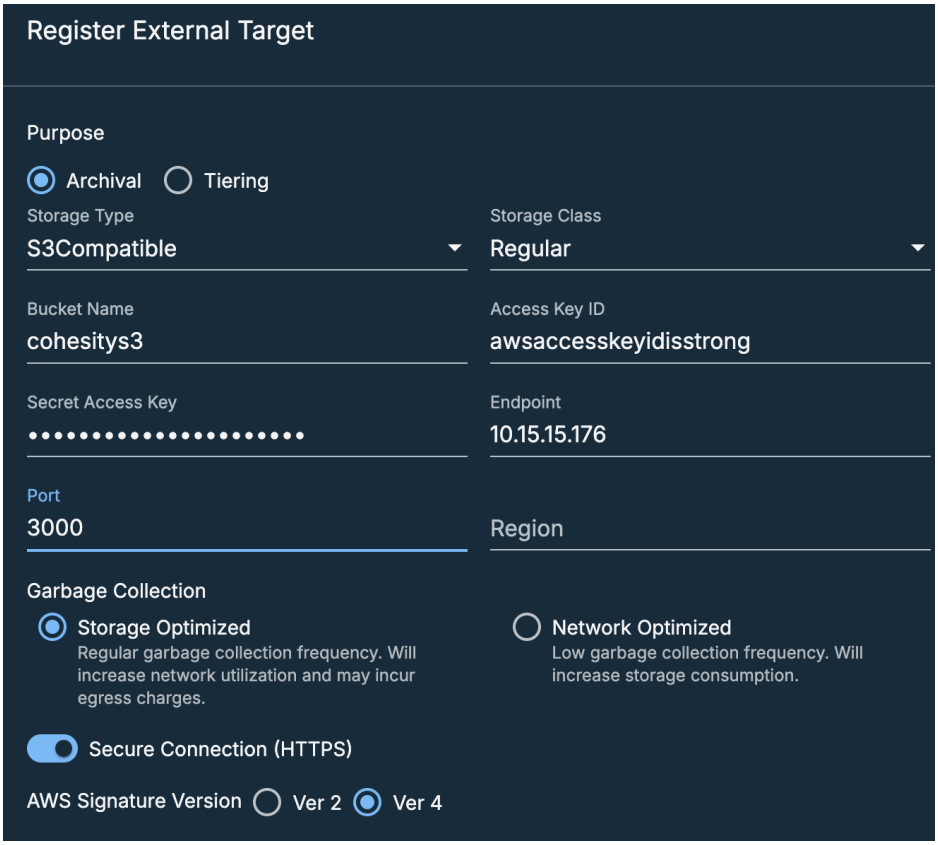

   f. Enter the Region (Geographic area to use for S3-Compatible)
   g. Choose Garbage Collection method.
      - Storage Optimized
      - Network Optimized

h. **Secure Connection (HTTPS)** is enabled by default.

i. Specify whether your S3 bucket uses AWS Signature Version 2 or 4

j. Enter a unique External target Name.

k. Select the Archival Format: Incremental Forever

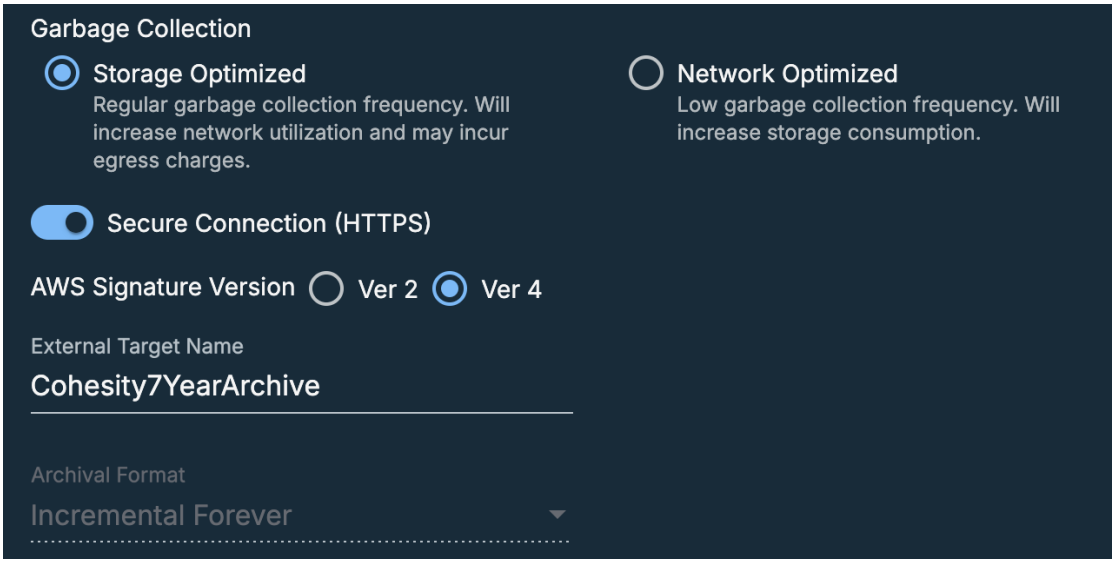

l. Archive Object Lock: If enabled, archived snapshots stored on external targets are locked until their expiration to prevent modification.

> **NOTE**: Snapshots stored on this external target are made immutable using S3 Object Lock until their expiration. Note that Object Lock increases storage used and once enabled, this option cannot be disabled.

m. By default, Encryption and Compression are enabled, while Additional security by managing key manually and Bandwidth throttling are disabled.

> **NOTE**: For more on Encryption, Compression, Source-Side Deduplication, Incremental Archival, and Bandwidth Throttling, see Create and Register S3 Bucket above.

- If you want to enable manual key management for extra security, turn on **Additional security by managing key manually.** A pop-up window appears to confirm the change.

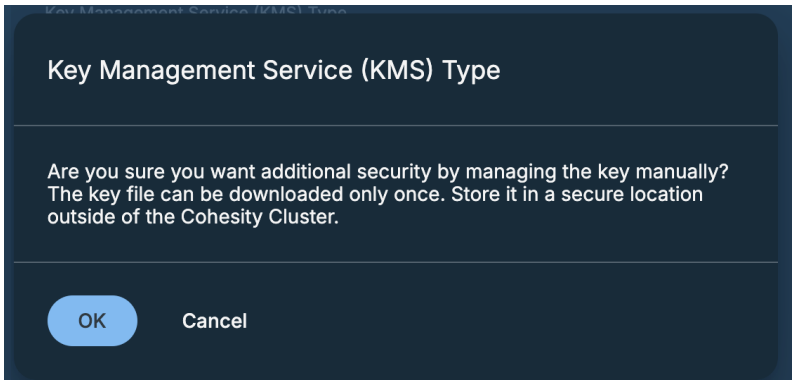

> **IMPORTANT**: With this option on, a cluster must have the key to access data from the archive. You can download the key file (only once) after you register your bucket. This key is required when you use CloudRetrieve. If you do not have it, you will still be able to recover data to its original cluster, but you will not be able to retrieve it onto a new cluster (in a disaster-recovery scenario, for example).

n. Enable **Bandwidth Throttling** if needed. You can throttle upload and download speeds separately and apply throttling all the time or only specific days and times.

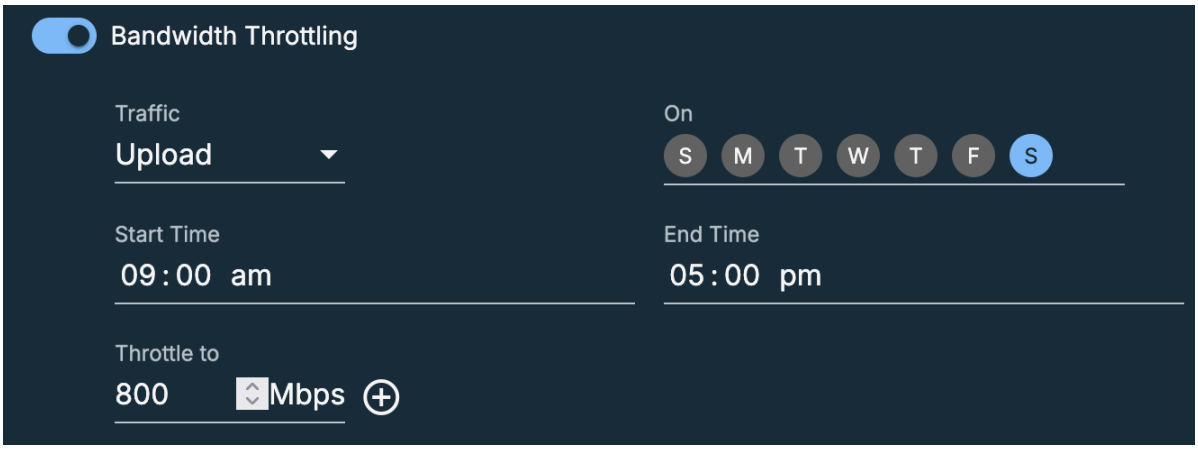

6. Click **Register**.

   Your S3 bucket is now an External Target in Cohesity, and is available to select when you create a Cohesity Protection Policy for use in Protection Groups.

Rotate Access Keys (Optional)

For security, it is important to rotate the access keys to your External Target in your S3-Compatible Object Storage. Depending on your corporate policy for changing keys and passwords, when the time comes, you will have to rotate your Object storage target's access keys and update your Cohesity External Target with the new keys.

To rotate the access keys:

1. Log in to the S3 vendor's management console (or use CLI/API).
2. Go to the **user management** section (IAM user in AWS, or equivalent in other vendors).
3. Create a **new access key pair** (Access Key ID + Secret Key).
4. Now you need to update the External Target on Cohesity with the new access keys. Log in to Cohesity Data Cloud and select **Infrastructure > External Target.**

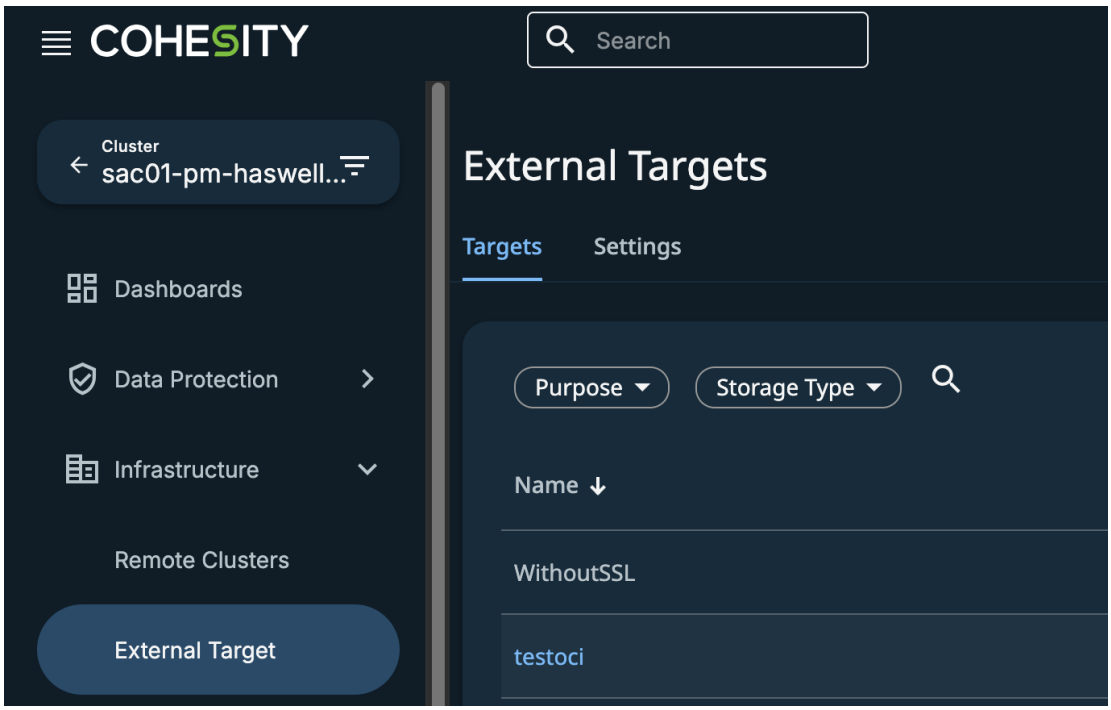

5. Find your External Target in the list and click **Edit** on the right.
6. Enter the new **Access Key ID** and **Secret Access Key**, then click **Save**.

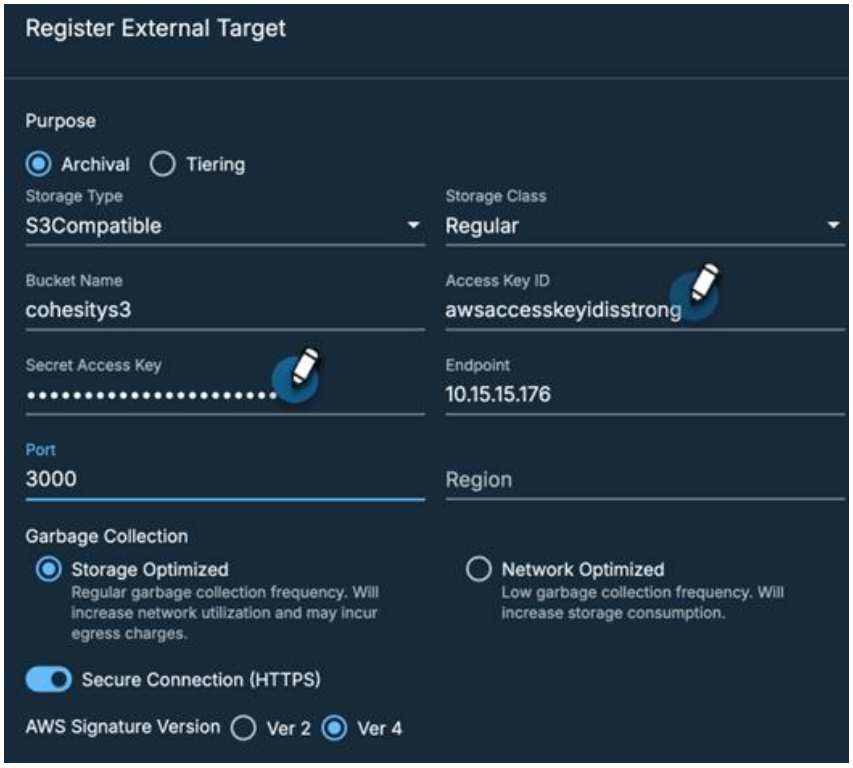

7. Verify that the new archival is completing successfully.
8. Once you confirm the new key is working, **deactivate/make inactive** the old key so that you can reactivate if necessary.

You have successfully rotated your Access Keys.

Create a Protection Policy

In Cohesity Data Cloud, Protection Groups use Protection Policies. Protection Policies reflect business needs, such as backup and archival frequency, retention requirements for each, as well as Recovery Point Objectives and Recovery Time Objectives, while a Protection Group defines operational requirements, such as which source objects to protect, the Protection Policy to use, and operational considerations like indexing, exclusions and inclusions, alerts, app consistency, and more. This process of combining business needs (Policy) with the objects to protect (source data) and the operational requirements (Job) provides rich flexibility to customers.

A Protection Policy defines:

- How source data (like virtual and physical servers, databases, unstructured data, etc.) will be backed up and then archived.
- Where and how frequently they will be archived.
- How long the archives will be retained.

This list addresses parameters that affect CloudArchive operations. For the complete list of Protection Policy parameters, see Create or Edit a Policy in the online Help.

In the Protection Policy, you can select the cloud-based External Target you just created and registered as an External Target.

To create a Protection Policy:

1. Log in to Cohesity Data Cloud.
2. Click **Data Protection > Policies.**

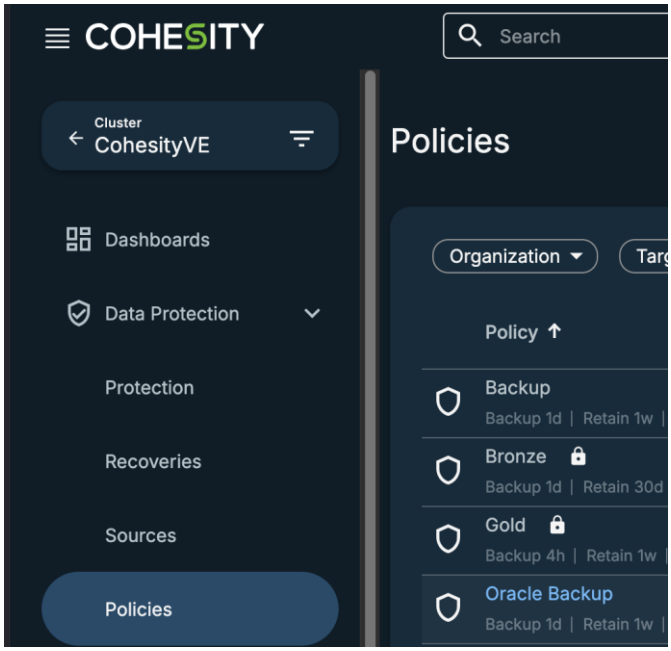

3. Click **Create Policy.**

4. In the form that opens, enter a **Policy Name.**

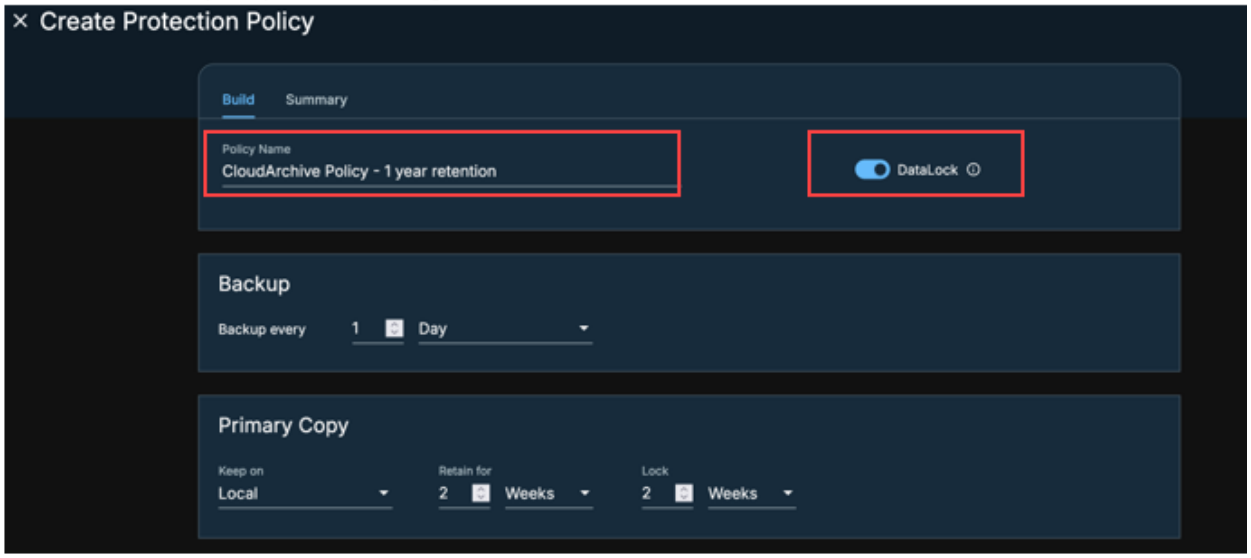

Add a **DataLock** for compliance and regulatory requirements, to ensure that your protected data, including local backups, archives, and replication, cannot be modified until the DataLock expiration.

Once applied, a **DataLocked Snapshot will be deleted only after its retention period expires.** A DataLock prevents all users, including those who have the Data Security role in Cohesity Data Cloud, from modifying or deleting any Snapshots that were generated by the Protection Groups that use this policy. Only users with the Data Security role can add, modify, or remove a DataLock from a Policy. See online Help for more information.

> **NOTE**: You can also add a legal hold to a specific Protection Group run (a *Snapshot*) to preserve it for legal reasons. See Apply Legal Hold to Completed Protection Group Run below.

5. Under **Backup**, set the **Backup** interval (every day, by default). Select **Primary Copy** as Local and specify retention period and lock period (if you want to apply data lock).

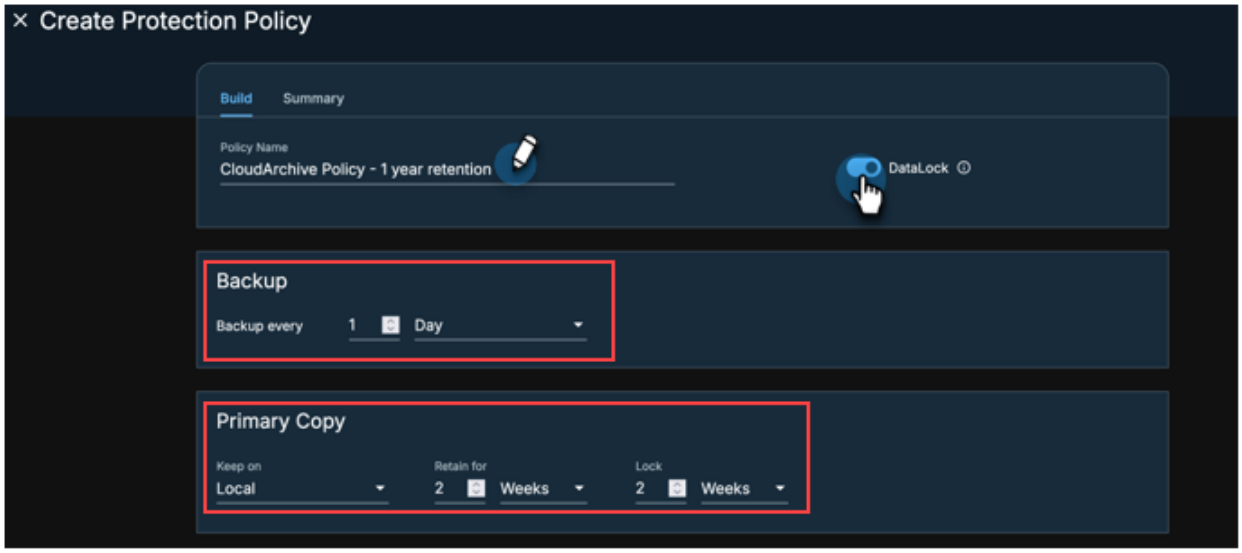

6. Click **Add Archive**, and for **Archive to**, select the External Target you just created. Set the **Archival** interval (every day, by default) and **Retain for** period. If the selected bucket has been enabled with versioning and object lock, then the UI will provide another option to specify the lock period for the archived data. You can also enable **Archive only fully successful Runs** in the checkbox on the right.

   Click **Add Archive** again if you need additional archival schedules.

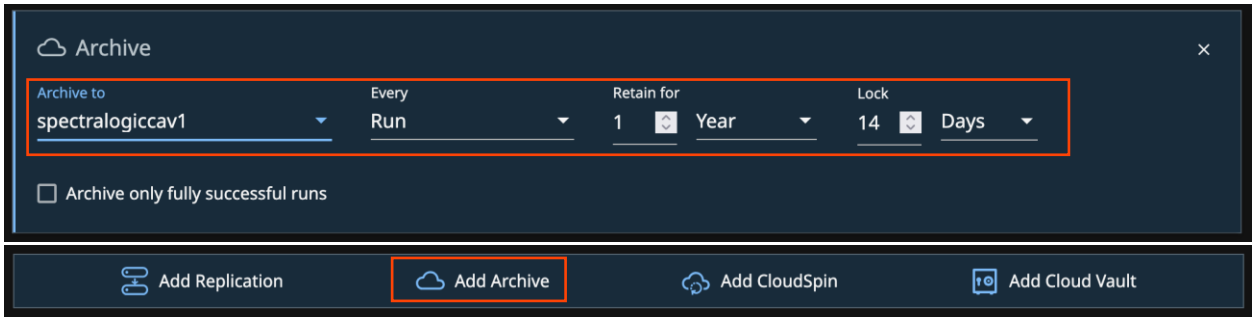

NOTE: You can add multiple archival schedules that use the same or different External Targets, as well as the same or different intervals and retention periods, to a given Protection Policy. When you add more schedules and send them to the same External Target with different retention and schedule times, the schedules rationalize among themselves and only the necessary archive is sent, with the longest retention.

For example, if you add these three archival schedules to the same External Target:

1. Once a day, retain for 90 days
2. Once every 7 days, retain for 180 days
3. Once every 30 days, retain for 365 days

Then:

- On Day 7, only one archive is sent, meeting both Schedule 1 and Schedule 2 (and retained for 180 days, per Schedule 2, as it is the longer of the two).
- On Day 30, only one archive is sent, meeting both Schedule 1 and Schedule 3, but is retained for 365 days, to meet the Schedule 3 retention requirement.

By contrast, if you send the archives to different External Targets, then:

- On Day 7, per Schedule 1, the archive is sent to the first External Target and retained for 90 days. Per Schedule 2, the archive is also sent to the second External Target and retained for 180 days.
- On Day 30, per Schedule 1, the archive is sent to the first External Target and retained for 90 days. Per Schedule 3, the archive is also sent to the third External Target and retained for 365 days.

When you use multiple schedules with different External Targets, the schedules don't rationalize, and you accrue network and storage usage for each scheduled run.

7. Click **Create**.

Your new Policy can now be used in Protection Groups. For the complete list of Protection Policy parameters, see Create or Edit a Policy in the online Help.

Create a Protection Group

Protection Groups combine operational requirements with the business requirements that are defined in a Protection Policy. Multiple Protection Groups can use the same Protection Policy, but each Job can have only one Policy. Protection Groups protect specific source objects, such as virtual servers, physical servers, Views, SQL servers, Oracle databases, Remote adapters, Pure Storage Volumes, or network-attached storage (NAS).

For this example, we look at the steps to create a Protection Group for NAS data, but the steps to protect other source objects are very similar.

To create a Protection Group:

1. Log in to Cohesity Data Cloud.
2. Click **Data Protection** > **Protection**.

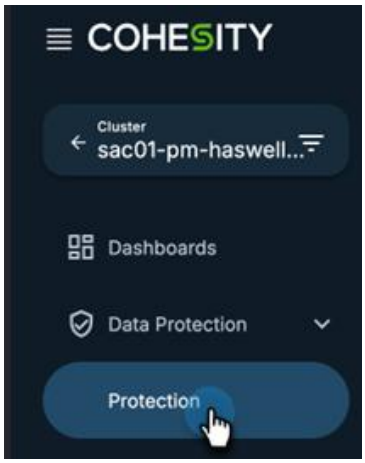

3. Click **Protect** and choose the type of data to protect.

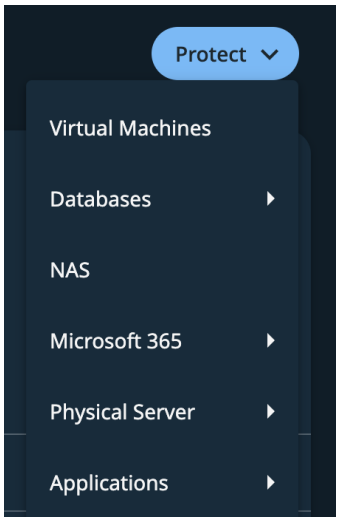

4. Click **Add Objects** to select a Source to protect.

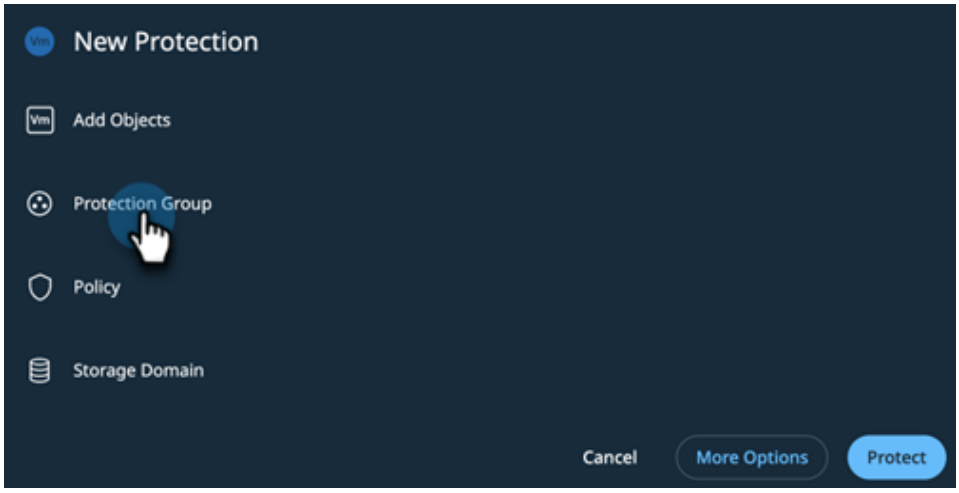

5. Select the specific objects you wish to protect, and click **Continue**.

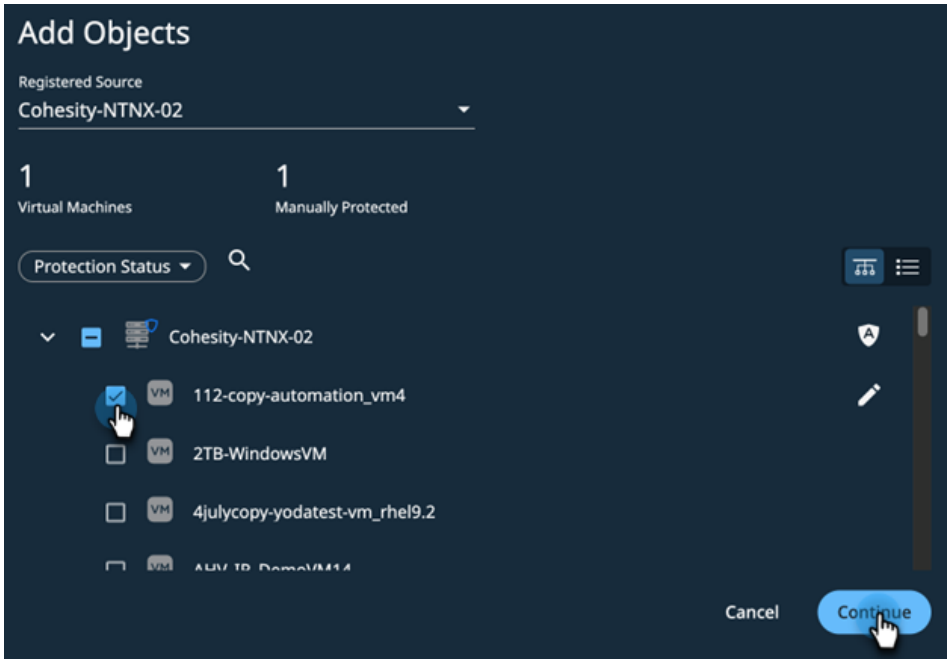

6. Name the **Protection Group.**

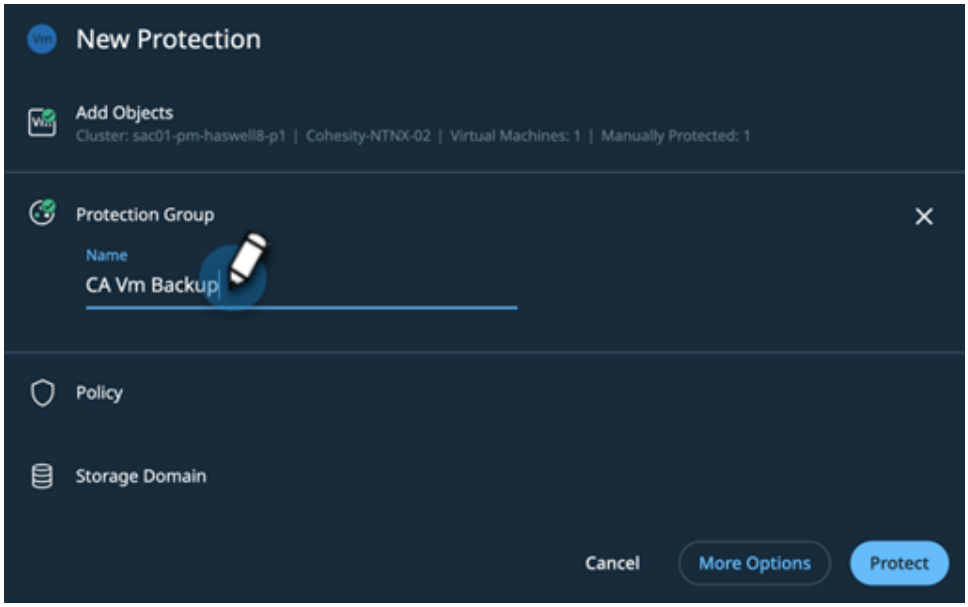

7. Select a **Protection Policy** which we have created earlier, or create a new Protection Policy.

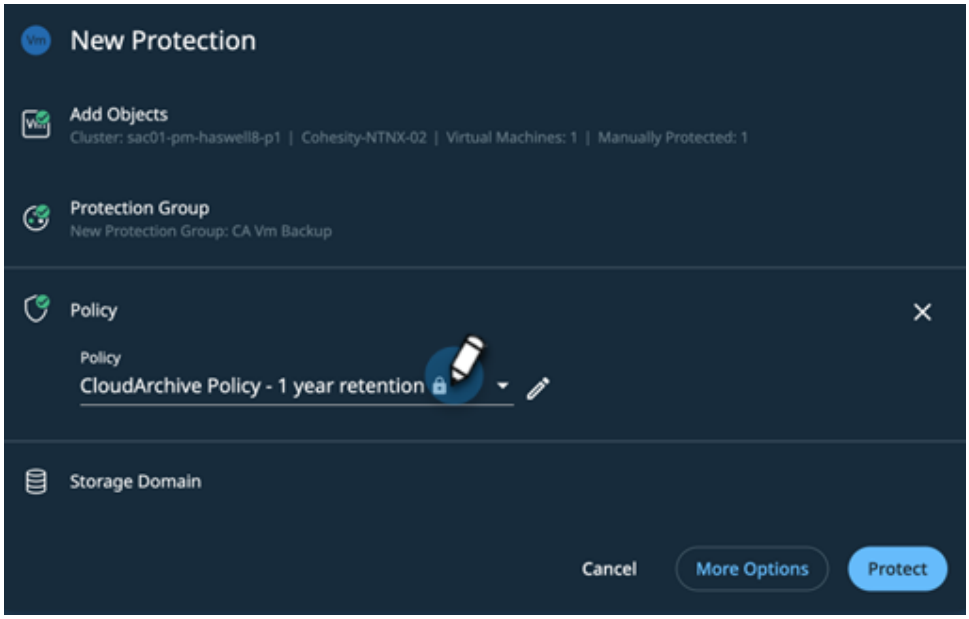

8. On the same screen, select a **Storage Domain**. Click **More Options** if you need to change any of the **Advanced** settings. When you're done, click **Protect**.

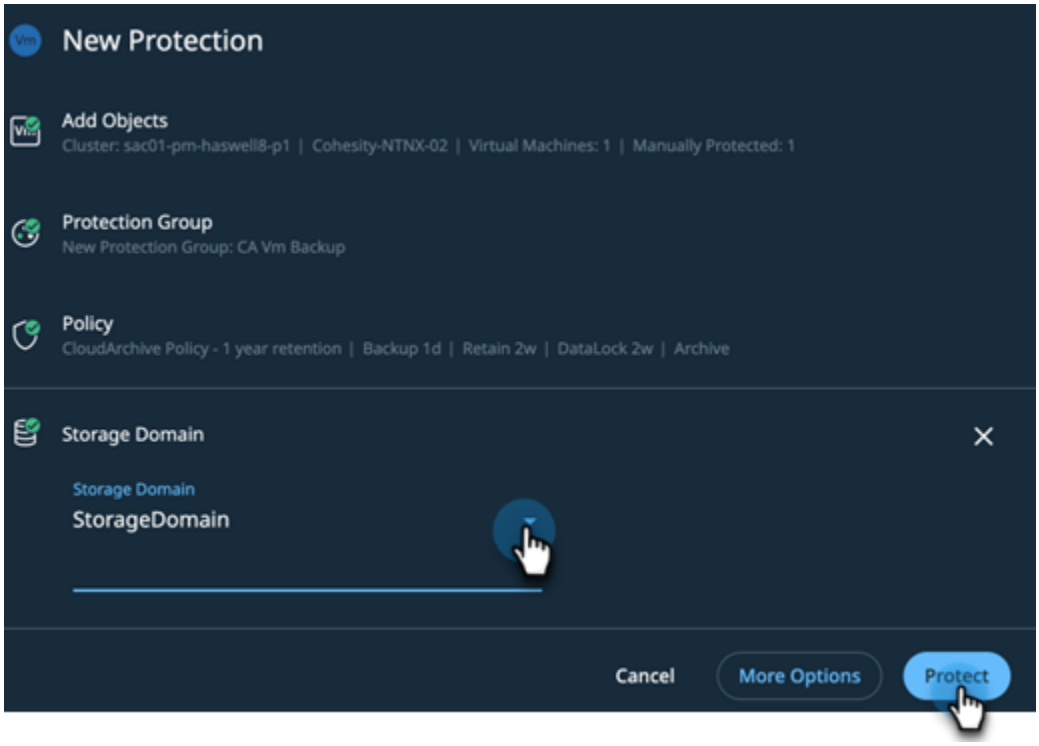

> **NOTE**: See the complete list of Advanced settings and the Job types that contain them in the Appendix A.

9. Select **Data Protection > Protection** to verify that your new group is on the list.

Your new Protection Group is now active and running. To manage Protection Groups, see online Help.

Apply Legal Hold to Completed Protection Group Run

Only users who are assigned the Data Security role can put a legal hold on existing Snapshots (Protection Group runs), to preserve them for legal purposes. Once a legal hold is applied, the retention period is ignored, and the Snapshot is preserved until the legal hold is removed. Legal hold Snapshots can only be deleted by a user with the Data Security role.

> **NOTE**: A legal hold can be added to both regular and DataLocked Snapshots.

You can add a legal hold to a Protection Group run or to individual objects in a Protection Group Run:

- If you add a legal hold to a Protection Group Run, it applies to all the Snapshot objects that were backed up by that Protection Group Run, and the legal hold is propagated to replicated and archived objects.
- If you add a legal hold only to selected objects in a Protection Group Run, the legal hold is propagated to archived objects, but not to the replicated objects. You must manage the legal hold status on the remote replication cluster manually.

NOTE: A legal hold prevents Snapshots from being deleted until the legal hold is removed. Using a legal hold for long periods of time can result in the cluster running out of space.

To add or remove a legal hold from a Protection Group Run, see Adding a Legal Hold to a Snapshot in the online Help.

The Difference Between Legal Hold and DataLock

While both a legal hold and DataLock are features that empower the Data Security role in Cohesity Data Cloud to prevent backed up and archived data from being deleted, they differ in purpose and function.

Table 4: The Difference between Legal Hold and DataLock

PURPOSE	LEGAL HOLD	DATALOCK
Business need	Reactive: Set on a specific Snapshot (i.e., Protection Group Run), usually prompted by legal requirements.	Planned: Set on all Protection Group Runs that use a Protection Policy with DataLock, usually for compliance.
Expiration period	No expiration. Removal managed by the user.	Defined in the Protection Policy.
Granularity	Set on individual Protection Group Runs and at the Object Level.	Applies to all Protection Group Runs of any Protection Groups that use a Policy with DataLock.
Deletion	Can be deleted to recover storage space, but only by a user with the Data Security role.	Cannot be deleted before the DataLock expiration date, even by a user with the Data Security role.

Recover Data from CloudArchive

Cohesity Data Cloud provides two ways to get your data back from your S3 bucket: Cloud Recover and CloudRetrieve.

Figure 8: Cohesity Protection, CloudArchive, Cloud Recover, and CloudRetrieve

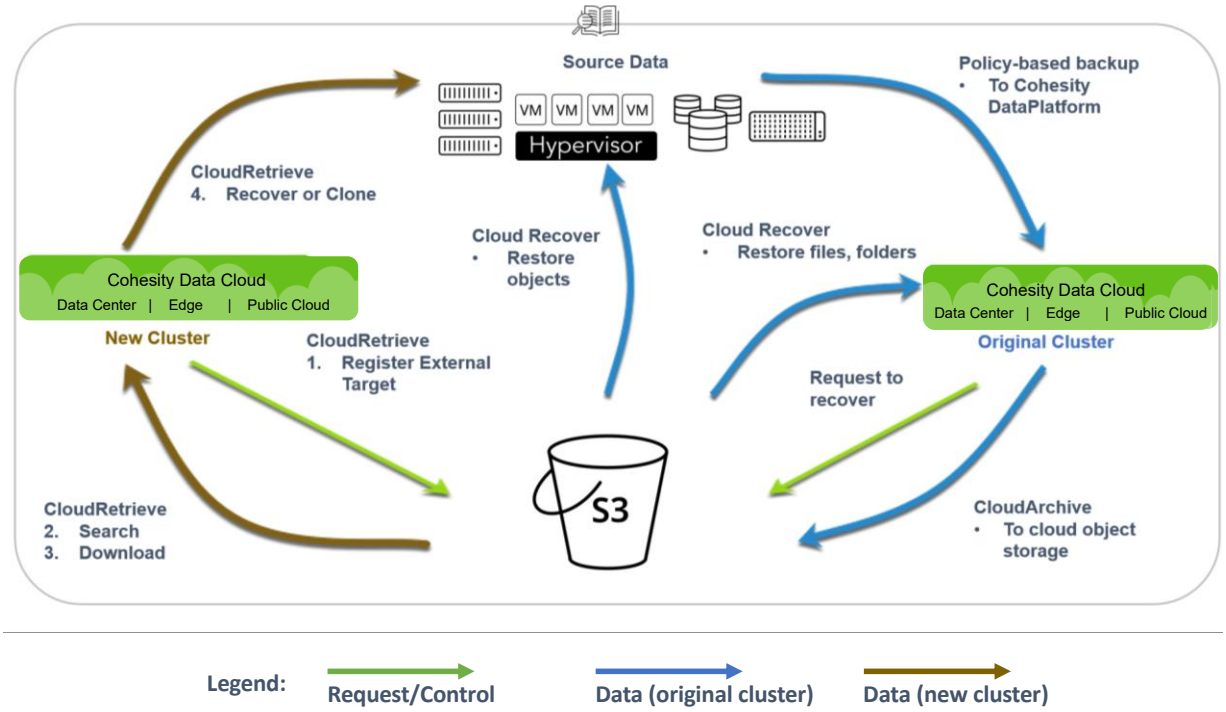

- **Cloud Recover**: Recover entire objects (such as VMs, databases, NAS, etc.) or individual files and folders back onto the Cohesity Data Cloud that archived them.

  NOTE: When you recover a complete object (such as a VM or database), it is restored to its original location once it is downloaded to the Cohesity Data Cloud from the cloud, and restored via the Instant Volume Mounting capability in Cohesity Data Cloud.

- **CloudRetrieve**: CloudRetrieve allows you to extract your Protection Group and its metadata, including Protection Group Run details, from the archive in the cloud, so you can search it and recover the data you need onto a new or different cluster. This approach involves several steps:
  - Register the External Target containing your archived data.
  - Search the archive in the cloud.
  - Select and download metadata for the archived Protection Groups.
  - Recover objects from the downloaded Protection Group

But first, let's start with recovering data onto your original Cohesity cluster.

Recover Your Data to Original Cluster

To locate and recover a file, a folder, or an entire virtual machine to the original cluster:

1. Log in to Cohesity Data Cloud.
2. Select **Data Protection** > **Recoveries**.

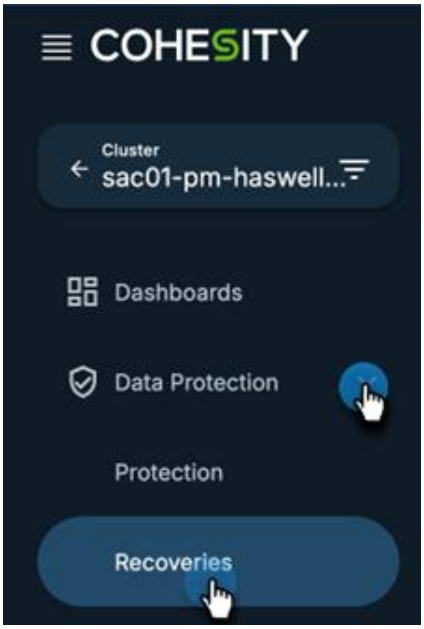

3. Click **Recover** and select the type of object you seek — a file or folder, VMs, physical server, and more.

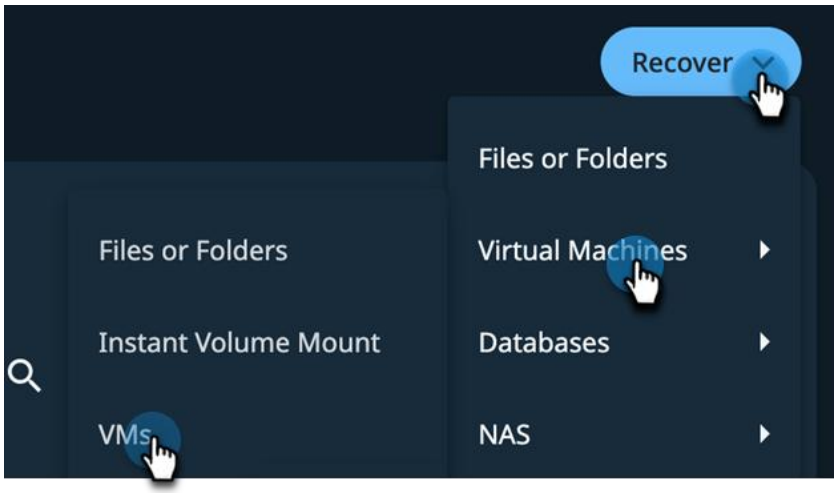

4. To retrieve a list of virtual machines, for example, select **VMs** and enter part or all of the VM names.

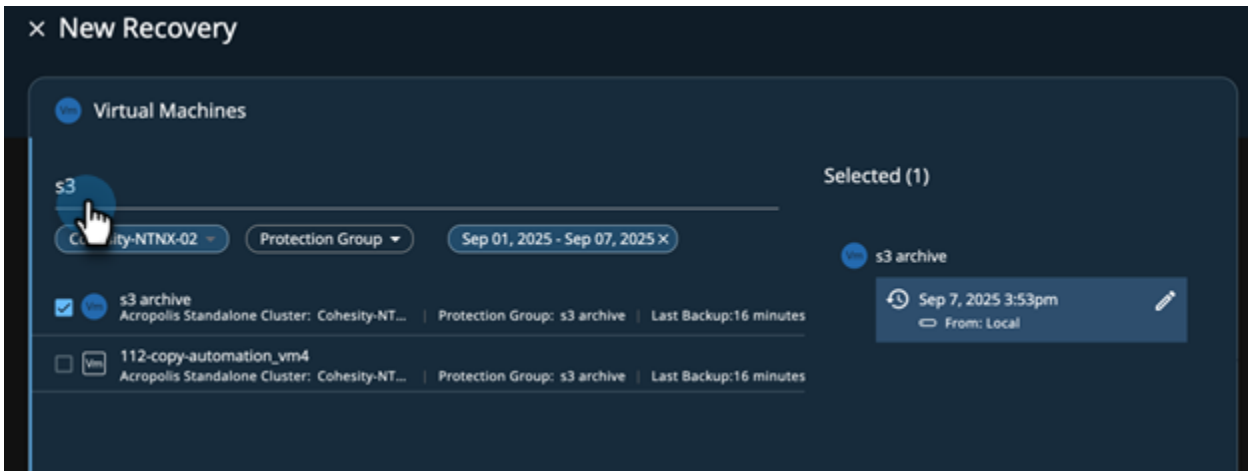

5. Select the VMs you need or select an entire Protection Group to recover all the VMs it archived, and then click **Edit** to select a recovery point. Next: **Recover Options**.

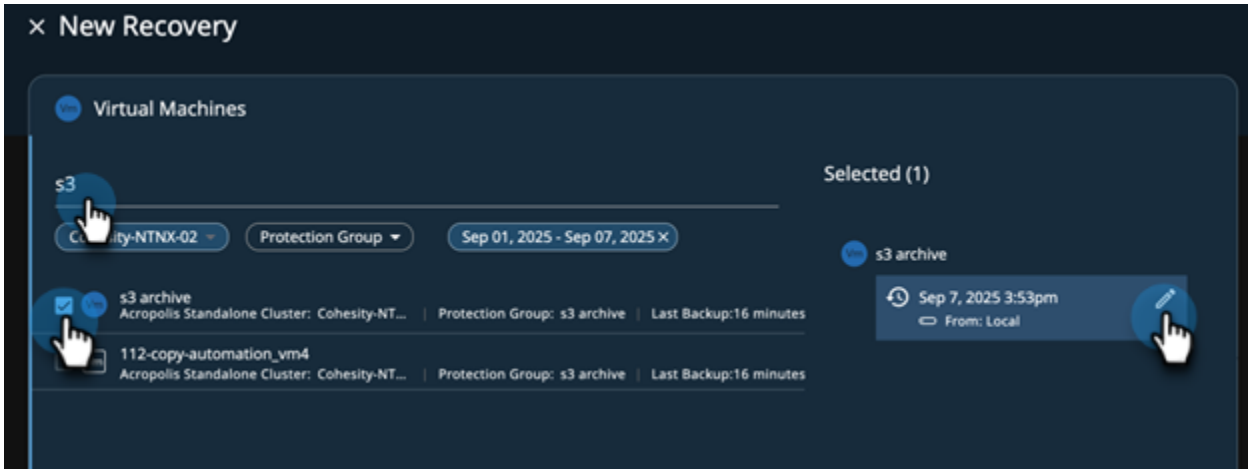

6. Choose a date. To recover data from the external target, change the location from local to cloud. Click **Select Recovery Point.**

7. Click **Next: Recover Options**

8. Choose a **Recovery Location and Recovery Method** and specify how to handle the existing VM.

9. In the Recovery Options, attach Network, **Rename** Recovered VMs with appropriate Prefix and Suffix. Select the **Power State** of the VM and enter a **Task Name.** Click **Recover** to start the recovery process.

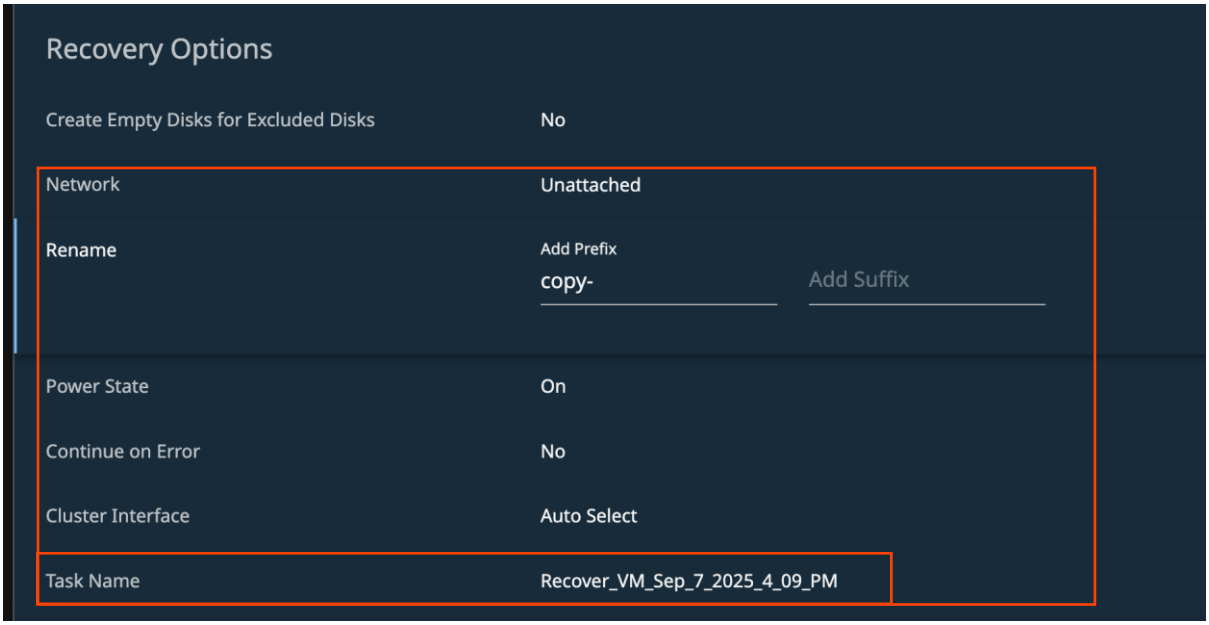

Table 5: Recover Task Options

RECOVERY OPTIONS	DETAILS
Recover to an Original location	Specify this option to recover the VM files (such as the VMDK files) to their original datastores and create new instances of the VMs in the original location in the original source. For more, see Recover to Original Location in the online Help.
Recover to New Location	Recover the VM files (such as the VMDK files) to an alternate datastore and create new instances of the VMs in an alternate Resource Pool of a registered Source. For more, see Recover to Original Location in the online Help.
Recovery Method	**Instant Recovery**: The VM(s) will be usable instantly in the target environment and will be moved to target storage later.
	Copy Recovery: Recovered VMs will be usable in the target environment only after all the data has been copied over from Cohesity to the target storage.
Detach network	For each recovered VM, the virtual Network Interface Card is removed from the VM.
Existing VM Handling	Specify how to handle the existing VM.
Network	For each recovered VM, connect to the original or new network when the VM reboots. **IMPORTANT**: If this option is not selected, the VMs are not connected to any network on reboot.
Rename	Rename recovered VMs with appropriate Prefix and Suffix.
Power State	The recovered VMs remain powered on after they are created.
Continue on Error	With this option, if one of the VMs cannot be created, Data Cloud will still attempt to create the other VMs.
Task Name	Specify a task name

NOTE: This example is for recovering a VM. The recovery options vary by Protection Group type.

For more on the many capabilities and choices in our recovery process, see Recovery in the online Help.

CloudRetrieve Your Data to New Cluster

CloudRetrieve provides the ability to download data that was archived from a cluster to an alternate (non-original) cluster. In other words, you have Cluster A, which archives data to an External Target, but you need to download that archived data to Cluster B, for geo-redundancy or disaster recovery.

Complete the following steps to recover data from your S3 bucket to a different Cohesity cluster:

Figure 9: CloudRetrieve Workflow

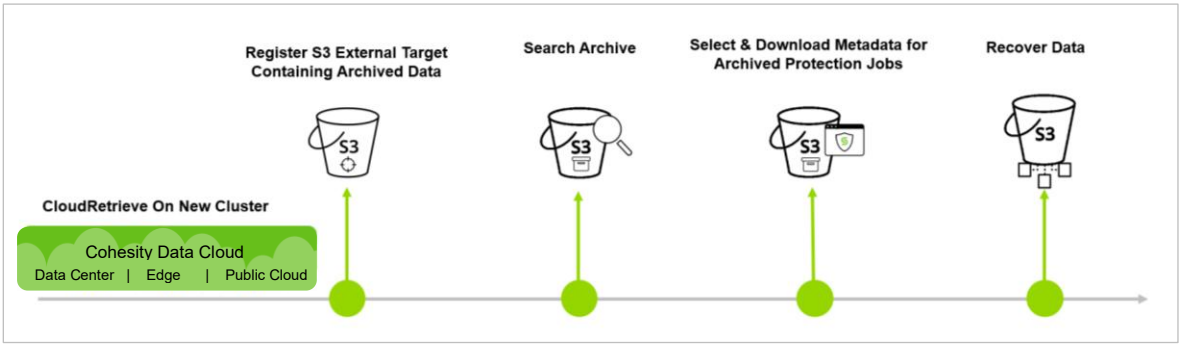

The sections below describe the steps to:

1. Register the External Target containing your archived data to the new cluster.
2. Enter the retrieve parameters (cluster name, date range, Protection Group name) to search the archive in the external target s3 bucket. (The search can take from minutes to several hours, depending on the data-retrieval SLA for the class of storage you are using from your S3 provider. For example, some storage classes have a standard retrieval SLA of up to several hours.)

   NOTE: If your External Target is protected by a manually managed key, before you can search it, you will need to upload the External Target's access key.

3. From your search results, select and download the metadata (the Protection Group Run details) for the archived Protection Groups onto the new cluster, so that you can review Protection Group Run details and choose just the specific you need to recover or clone.

   NOTE: In this step, you are prompted to select a date range, and if you know exactly which Protection Group Run (Snapshot) you need, you can also choose to download it along with the metadata, to be able to recover your data objects as soon as it completes.

4. After the metadata download completes, select the necessary Protection Group Run from the archived Protection Group to recover or clone your objects.

Register S3-Compatible External Target Containing Archived Data

To register your S3 bucket as an External Target on the new cluster:

1. Log in to a cluster other than the cluster that archived your data, or stand up a new cluster.
2. Log in to Cohesity Data Cloud on your new cluster.
3. Follow the steps in Register S3 Bucket with Cohesity Data Cloud above to register your archived S3 bucket to the new cluster.

Search Archived Data from S3 Bucket

To submit a search request for a list of archived clusters and Protection Groups:

1. Log in to Cohesity Data Cloud on the new cluster.
2. Select **Data Protection** > **CloudRetrieve**.

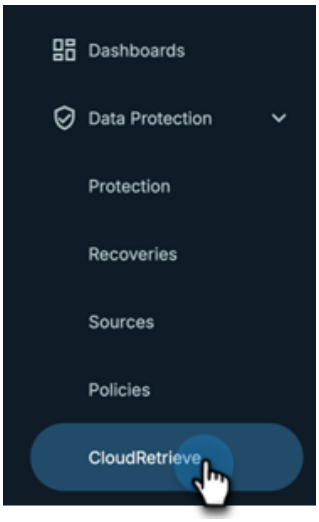

3. If this is your first use of CloudRetrieve, the CloudRetrieve summary screen appears. Click **Start Search**.

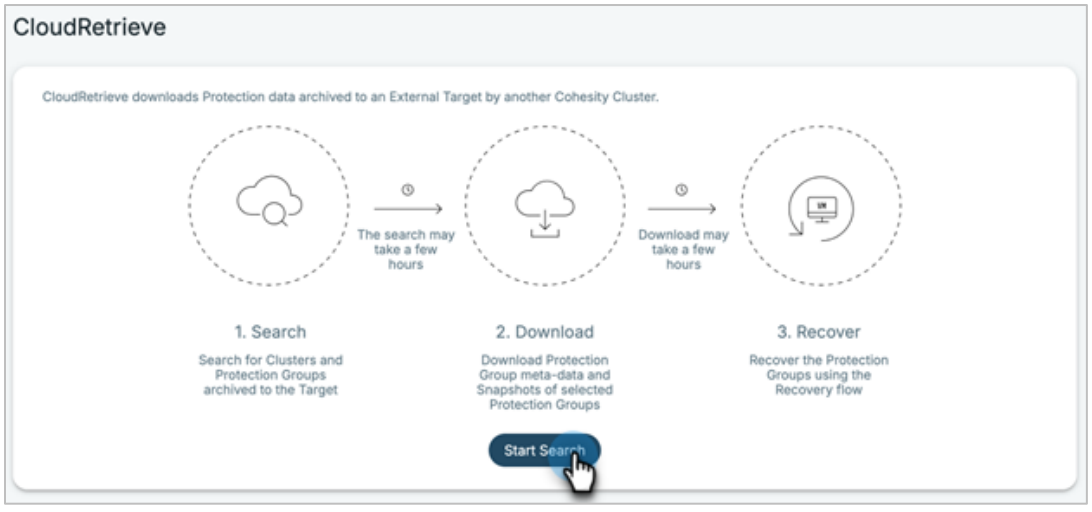

a) If this is not your first visit, the list of downloaded Protection Groups appears. In that case, click **New Search**.

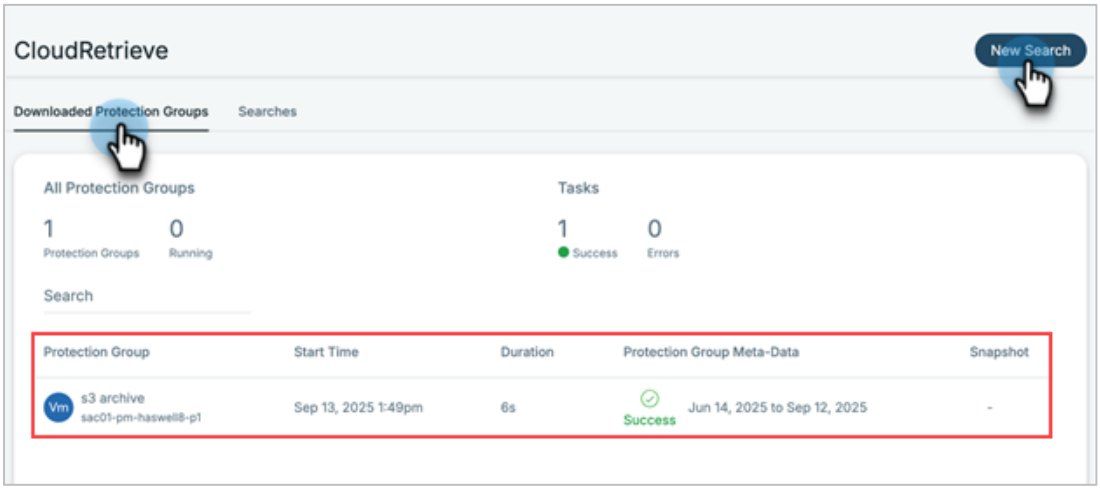

b) Click on **Searches**, to view previous Cloud Retrieve Searches performed

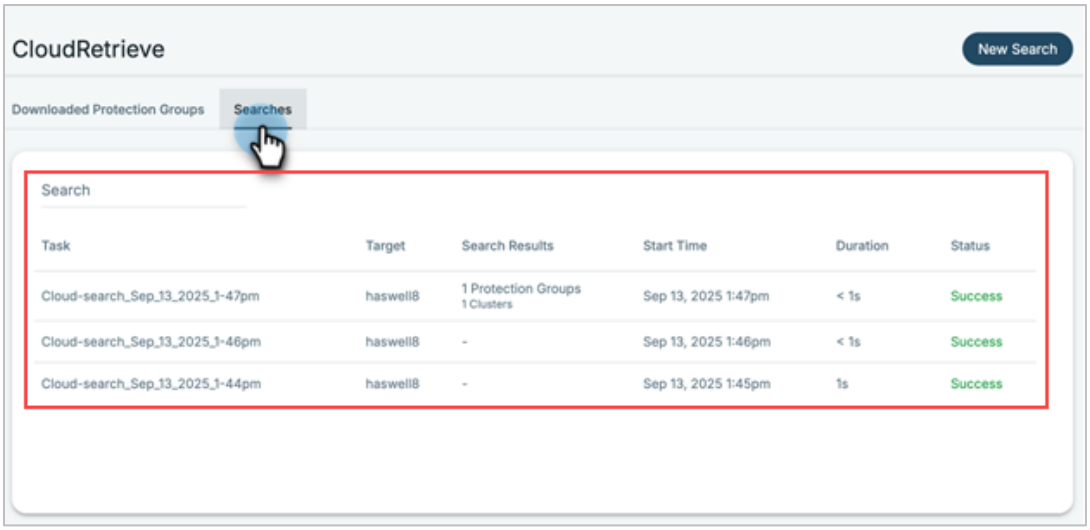

4. Select your **External Target** from the drop-down list.

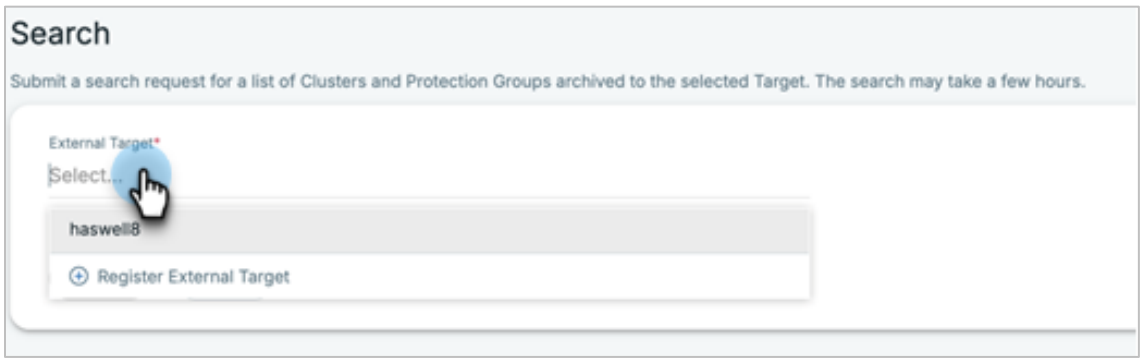

NOTE: If you skipped the first step and have not yet registered your External Target, you can register here. To do so, click **Register External Target** from the drop-down menu and follow the steps in Register S3 External Target with Cohesity Data Cloud.

5. In the form that opens, enter:

Table 6: CloudRetrieve Search Options

FIELD	DESCRIPTION	NOTES
Date Range (required)	Select a Date Range (past year by default) to limit the scope of your search.	
Cohesity Cluster Name (optional)	To narrow your search to a specific cluster, enter a cluster name. This is especially helpful if the same storage is used with more than one cluster. To broaden your search to match more than one cluster, use a partial name (for example, 'Acme' instead of 'Acme_Raleigh').	**IMPORTANT**: Wildcard characters (like '*') are NOT supported. If you enter search terms for both **Cluster Name** *and* **Protection Group Name**, your search must find matches for the Protection Group *within* clusters that match. If your search is too narrow, try entering a search term for just **Cluster Name** *or* **Protection Group Name**, or leave one or both empty.
Protection Group Name (optional)	To narrow your search to a specific Protection Group, enter a Protection Group name. This is especially helpful if the same storage is used for more than one Protection Group. To broaden your search to match more than one Protection Group, use a partial name (for example, 'NAS' instead of 'NAS-Bronze').	
Upload key file (optional)	If your External Target is protected by a manually managed key, click **Attach**.	
Task Name (required)	By default, Cohesity uses the current timestamp to name the task automatically (for example, 'Cloud_search_***<CurrentTime>***'). Cohesity recommends that you replace the automatic **Task Name** with terms that will make it easy to identify (for example, '***<ExternalTarget>***_From_***<SourceCluster>_<Purpose>***').	

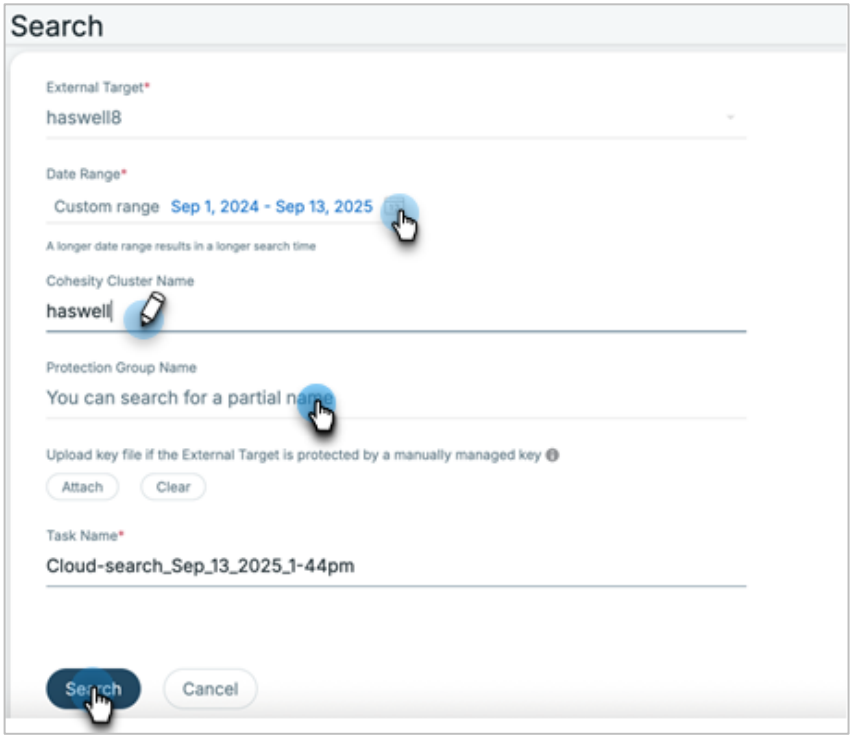

6. Click **Search**.
7. Wait while the search runs.

> **NOTE**: The search can take from minutes to several hours, depending the data-retrieval SLA for the class of storage you are using from your S3 provider. For example, some storage classes have a standard retrieval SLA of up to several hours.

The success of a CloudRetrieve search does not guarantee that the search found any matches. It means only that the search operation completed successfully. If your search results came up empty, broaden your search with partial names for the cluster and/or Job, leave them blank, and/or extend the date range.

Select and Download Metadata for the Archived Protection Groups

Once you have your search results, choose the Protection Groups to download to your new cluster. After the download, you will be able recover your data from the downloaded archive. See Figure 8 above.

When your search completes:

1. Select the Protection Group(s) you wish to recover from the search results and click **Edit**.

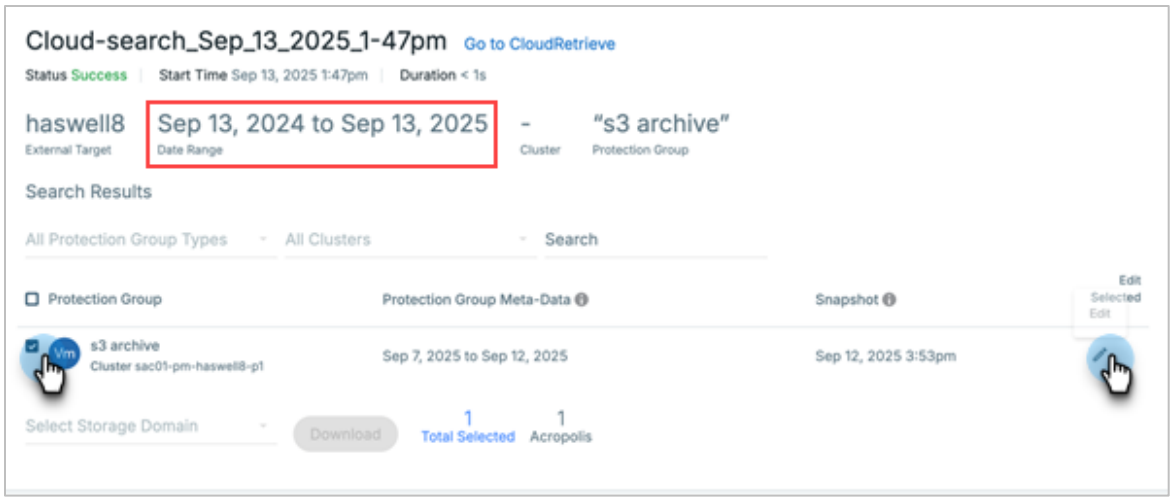

2. In the form that opens, you can choose to **Download Protection Group Meta-Data** (that is, the details of each Protection Group Run in the archived Protection Group), **Download Snapshot** (a specific Group Run), or both. Select a snapshot date range.

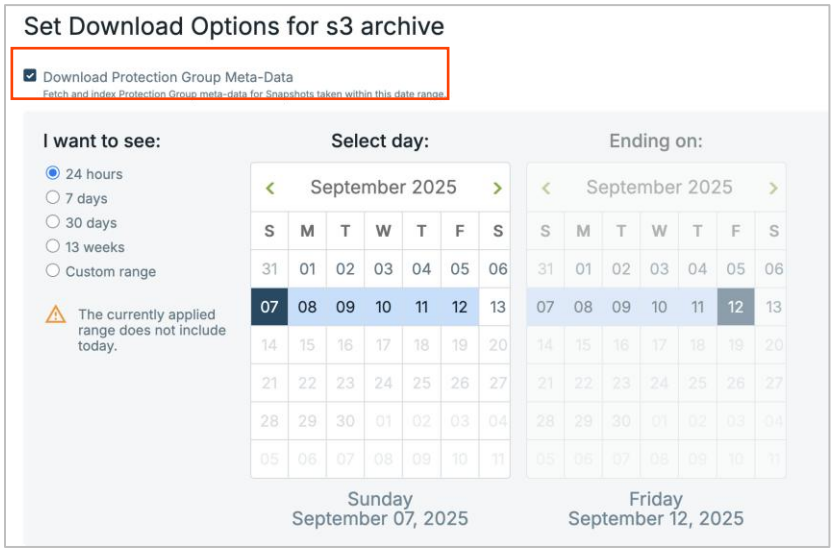

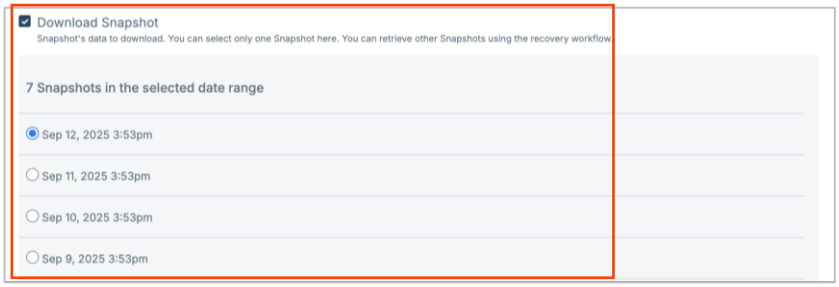

NOTE: If you are not certain which Snapshot contains the objects you need to restore, Cohesity recommends you deselect **Download Snapshot**. Once you have the Protection Group metadata, you will be able to review the details of each Snapshot in the Protection Group, to help you narrow the download to just the specific data you need.

3. Set your download options and click **Save**.
4. Select the **Storage Domain** and click **Download**.

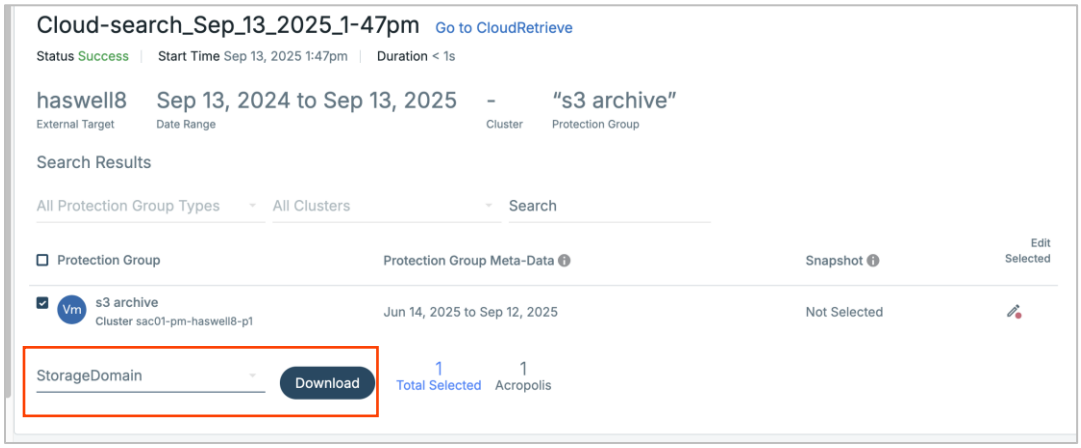

5. The downloaded Protection Group(s) will be accessible as **Failover Ready** under **Protection Groups**.

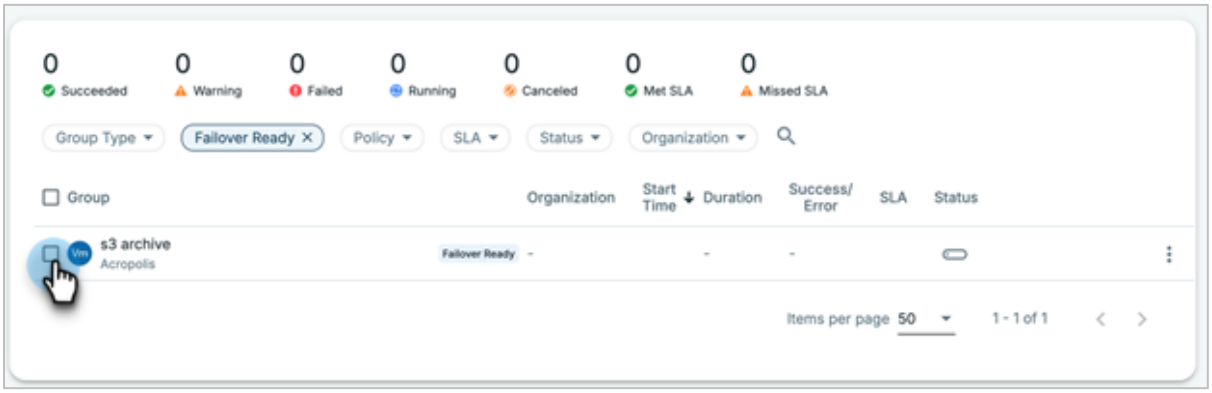

Wait for the download to be completed.

The downloaded Protection Group will be listed on the Cloud Retrieve Page.

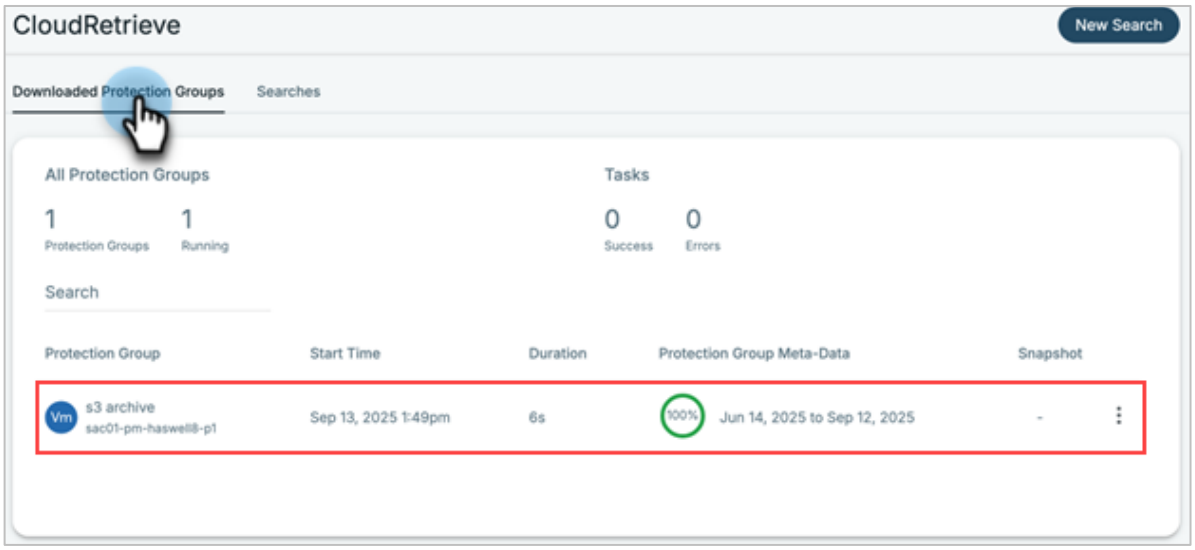

The Protection Group is now available on your new Cohesity cluster, and can be used to recover your archived data.

NOTE: CloudRetrieved Snapshots are not made to expire automatically by the new cluster. Once you have recovered the data you need, if you need to reduce your external target storage expenses, you will have to delete the archived data from your S3 bucket manually. Do NOT do this if the original cluster is still intact.

Recover Source Objects from Retrieved Archive on New Cluster

Now that you have downloaded the archived Protection Group Runs metadata onto the new cluster, you can recover whole objects or individual files from the downloaded archive.

To recover an entire data object from a CloudRetrieved archive:

1. Log in to Cohesity Data Cloud on the new cluster.
2. Select **Data Protection** > **CloudRetrieve**.

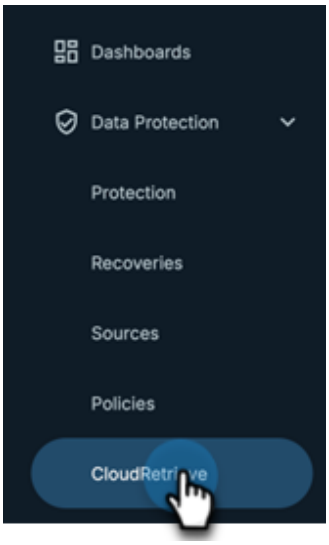

3. On the **Downloaded Protection Groups** tab, find the Protection Group you retrieved and click on it.

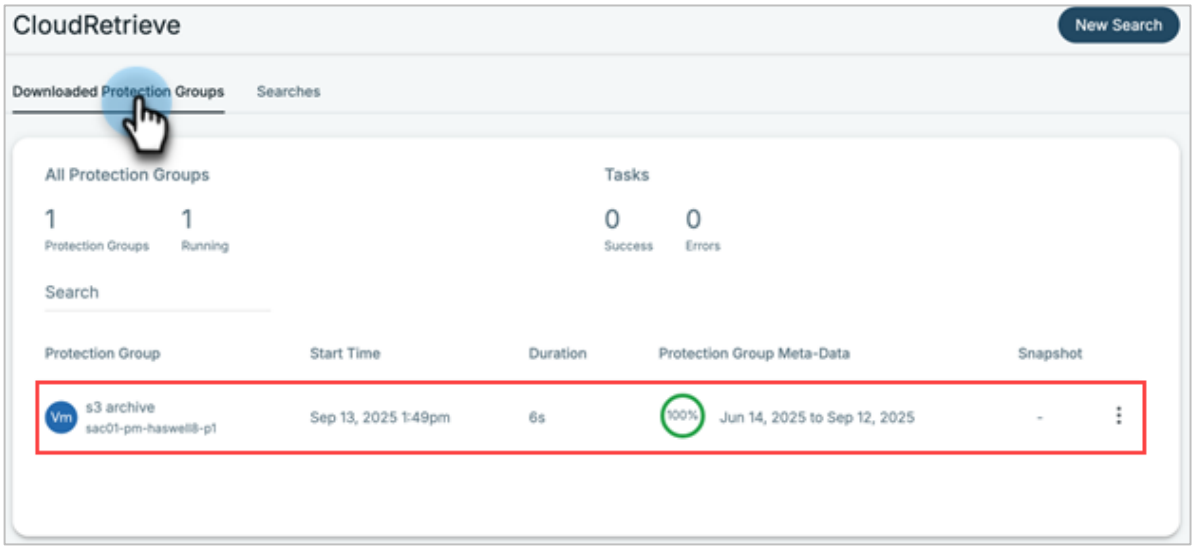

4. If the list of Protection Group Runs is empty, click the date range and select a longer range.
5. When the list of Protection Group Runs in the retrieved archive appears, inspect the details for each Run (**SLA**, **Schedule Type**, **Logical** and **Data Read**, **Success/Error**, and **Run Status**) and click the most appropriate Protection Group Run.

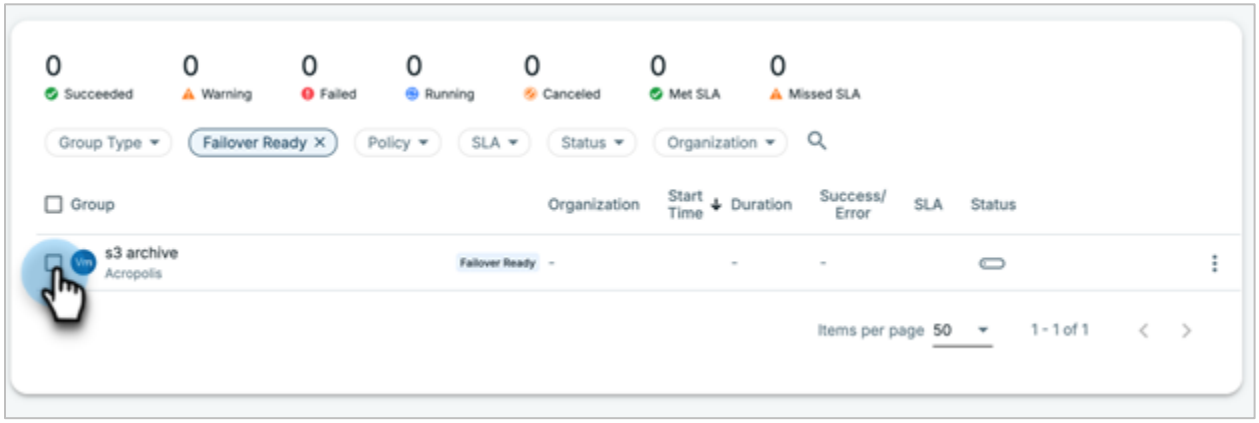

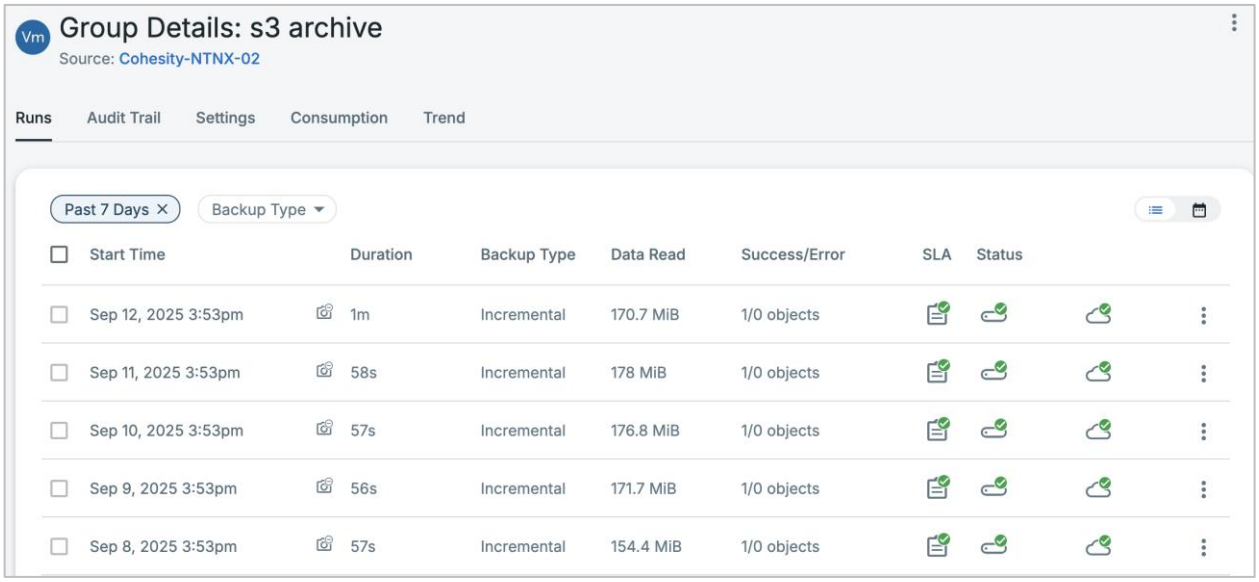

6. In the list of data objects included in that Protection Group Run, find the object you need to recover (for example, a particular VM), hover over the Action menu on the right, and select **Recover VM** or **Clone VM** (If the local snapshot has expired, we need to Recover using Recoveries workflow)

7. Edit the **Task Name** and **Recover** as fields, if necessary, and then follow the rest of the standard procedure for recovery above to complete your recovery task.

See About CloudRetrieve in the online Help for more.

Garbage Collection Update with S3-Compatible Object Storage

For Incremental Forever archival, the Cohesity cluster may download portions of data from S3-compatible and NAS external targets during each garbage collection process and reupload it after compaction. Cohesity 7.2.2 introduces two Garbage Collection Methods, to ensure efficient Garbage Collection.

Choose a **Garbage Collection** option based on your optimization preference:

- **Storage Optimized**: Selecting this option maintains a regular garbage collection frequency to optimize storage consumption on the archive. However, this results in higher network bandwidth usage. To minimize additional network usage, Cohesity recommends using S3-compatible and NAS targets within the same data center or network as the Cohesity cluster.
- **Network Optimized**: Selecting this option reduces garbage collection frequency to minimize network bandwidth usage. However, this results in higher storage consumption on the cluster due to less frequent garbage collection.

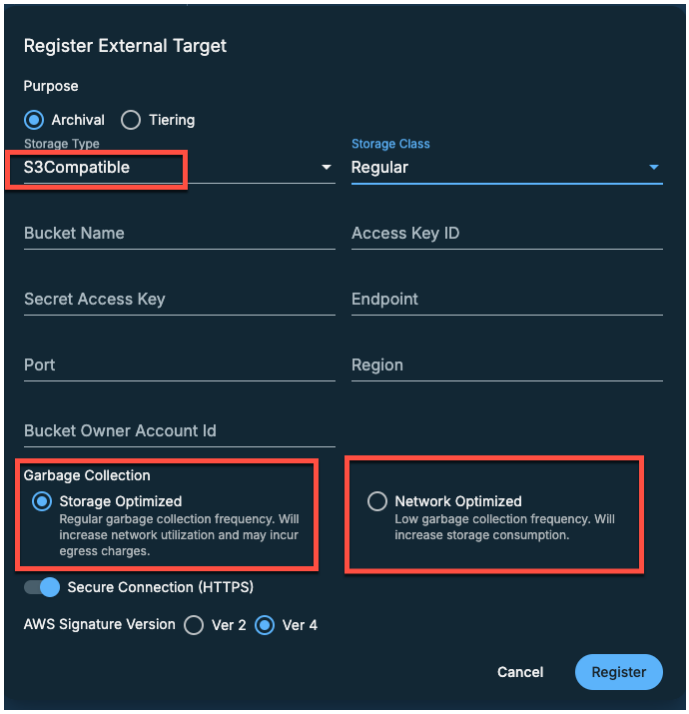

Appendix A: Protection Group Advanced Settings

Protection Groups combine operational requirements with the business requirements that are defined in a Protection Policy. See all the advanced Protection Group settings, and the Protection Group types that include them, in Table 8:

Table 7: Protection Group Advanced Settings

FIELD	DESCRIPTION	APPLICABLE GROUP TYPE
Pause Future Runs	Once enabled, no runs will be scheduled	All group types
End Date	Toggle on End Date and select the date on which the Protection Group stops capturing Snapshots. A Protection Group Run that starts prior to this date will run until completion even if it completes after this date.	All group types
QoS Policy	Select HDD or SSD. Backup HDD: The Cohesity cluster writes the data directly to an HDD drive for this Protection Group. Backup SSD: The Cohesity cluster writes the data directly to an SSD drive for this Protection Group. Only specify this policy if you need fast ingest speed for a small number of Protection Groups. Cohesity recommends HDD (the default).	All group types
Leverage Storage Snapshots for Data Protection	Toggle on to leverage storage array-based snapshots. For backups with high change rate deltas, this option can minimize the persistence time of VADP snapshots. This feature can leverage Cisco HyperFlex or Pure Flash Array Storage snapshots.	Virtual Server (VMware only)
Pre & Post Scripts	Edit this option to run scripts on the protected server before and/or after a Protection Group runs. If configured, the scripts are run every time an object is backed up by a Protection Group Run.	Physical Server, MS SQL, Oracle Database, NAS
Skip Files on Errors	Toggled on by default. The Protection Group continues to run even if it encounters errors on files, such as permissions errors. If files are skipped, the Protection Group Run details page indicates a warning status and provides additional	NAS **NOTE**: This setting is always enabled automatically for file-

FIELD	DESCRIPTION	APPLICABLE GROUP TYPE
	information. If toggled off, the Protection Group stops when it encounters an error.	based Physical Server backups.
File DataLock	Enable DataLock in Compliance or Enterprise mode	
Use Isilon Change List	Leverages the Isilon Changelist API to directly discover changed files/directories for faster incremental backup. Cohesity needs to keep one extra snapshot on Isilon after each backup, which will be deleted by the next successful backup.	Isilon
Exclusions and Inclusions	Everything is included by default. Toggle on Exclusions and Inclusions if you want to exclude or include locations. By creating exclusion and inclusion rules, you can limit the Protection Group to a specific set of files and directories and therefore minimize the disk space used to store the data. Cohesity automatically excludes the following NetApp system files: .vtoc_internal and .bplusvtoc_internal files .copy-offload directory and .tokens file WARNING: Always specify forward slashes (/) even for Windows systems. For Windows, do not specify the drive letter and colon in front of directory path.	Virtual Server, NAS, Microsoft 365
Indexing	Indexing is required for file recovery. The Cohesity Cluster will scan all the files in the Protection Group and create an internal index that can be used later by a Recover task to locate files by name. When creating a volume-based SQL group, indexing is not turned on automatically. Cohesity recommends turning the indexing on because indexing is required to restore .mdf, .ldf, and .ndf files from the cloud.	Virtual Server, Physical Server, MS SQL, Microsoft Office 365, NAS
Cancel Runs at Quiet Time Start	Cancel in-progress Protection Runs at the start of quiet times (as defined in the associated Protection Policy).	All group types
Alerts (optional)	Select one or more of the following settings if you want Alerts to be created for the following triggers:	All group types

FIELD	DESCRIPTION	APPLICABLE GROUP TYPE
	Success: Create an Informational Alert when a Protection Group completes successfully. Emails are not sent when Informational Alerts are created. Failure: Create a Critical Alert if the Protection Group fails to complete. Emails are sent when Critical Alerts are created. SLA Violation: Create a Warning Alert if the Protection Group takes longer than the time period specified in the SLA field. Emails are sent when Warning Alerts are created.	
Priority	Select a priority for the Protection Group execution. Cohesity supports concurrent backups, but if the number of Jobs exceeds the ability to process Jobs, they are executed in priority order: High first, then Medium, and then Low.	All group types
SLA	The Service-Level Agreement (SLA) defines how long the administrator expects a Protection Group Run to take. Incremental: Enter the number of minutes you expect an incremental backup Protection Group Run to complete. An incremental backup captures only the differences (changed blocks) since the last Protection Group Run. Full: Enter the number of minutes you expect a full backup Protection Group Run to complete. A full backup captures the entire object (all blocks).	All group types
Description	Specify a description for the Protection Group	All group types
Pause Future Runs	Once enabled, no runs will be scheduled	All group types
Email Recipients	You can add an email address to a Protection Group to notify the email recipient(s) when Alerts are triggered for the Job.	All group types
Incremental After Restart	For Windows physical servers, toggle on if you want the first scheduled backup performed after an intentional restart of the server to be an incremental backup instead of a full backup. If an incremental backup is not possible, for example, the server restarts after a power failure, a full backup is performed.	Physical Server (block-based only)

FIELD	DESCRIPTION	APPLICABLE GROUP TYPE
Abort in Blackouts	Available only if the selected policy has at least one Blackout period. Toggle it on to specify that all currently executing Protection Group Runs should abort if a blackout period specified for the Protection Group starts. By default, this toggle is off, which means after a Protection Group Run starts, it continues to execute even when a black period specified for this Job starts. However, a new Protection Group Run will not start during a blackout period.	All group types NOTE: This setting only applies to local backups and not to replication and archival.
Exclude Disks	By default, all the volumes on a selected server are protected. Toggle on to select disks to exclude for all VMs is this Job. Provide the type of controller, bus number and unit number for each disk to exclude. Excluded disks are not backed up and are not recovered during VM recovery. NOTE: If you exclude a disk that is part of a striped volume, the Cohesity cluster does not index the volume even if Indexing is toggled on for this Job. You will not be able to search for or recover individual files in that volume.	Virtual Server
Exclude Physical RDM Volumes	Toggle on to exclude VMs that have Physical Disks with Raw Device Mappings (RDMs). If toggled off, those VMs will not be backed up and Protection Group Run will fail. (Creating Snapshots of VMs that have Physical Disks with Raw Device Mappings (RDMs) is not supported by VMware vSphere or the Cohesity cluster.)	Virtual Server
App Consistent Backups	Toggle App-Consistent backups for a Protection Group if you want the guest Operating Systems of all the VMs in the Job to be quiesced before Snapshots of these VMs are created. If this option is selected, the Cohesity cluster makes a request to the VMware vSphere software to create a quiesced VM Snapshot by invoking VMware Tools (installed on the guest Operating Systems of the VM). The VMware Tools requests that applications on the guest OS quiesce their state so application-consistent Snapshots can be created. This quiescing of VMs prior to capturing Snapshots ensures the integrity of the data saved in the Snapshots. App-Consistent backups apply to VMs only. For physical Servers, Windows is	Virtual Server

FIELD	DESCRIPTION	APPLICABLE GROUP TYPE
	app-consistent by default and Linux is crash-consistent. For more information, see Creating Application-Consistent Snapshots in the online Help. If the App-Consistent backups toggle is on, the Take a Crash Consistent backup 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 the Cohesity cluster fails to capture an App-Consistent Snapshot. For example, the cluster may be unable to perform an App-Consistent backup when VMware Tools is not installed on the VM, the VM is powered off, or the VM cannot be quiesced. If the Take a Crash Consistent backup if unable to perform an App Consistent backup toggle is off and the Cohesity cluster is unable to perform a App-Consistent backup of a VM, a Snapshot is not captured.	
Cloud Migration	Enable this option to support the Cloud Migration of VMs between hypervisors (such as vCenter) servers and Cloud Services for failover and failback. Cohesity agents must be installed on the Windows VMs prior to backing up the Snapshots on the on-premises Cohesity cluster. Disaster Recovery using Cloud Migration is currently supported for Windows VMs backed up from a VMware vCenter Source. For more information, see Disaster Recovery of VMs using Cloud Edition in the online Help. After enabling Cloud Migration, you can download the Cohesity Agent directly from the Cohesity Dashboard.	Virtual Server
Backup Method	Two backup methods are available for MS SQL on physical Servers: Volume-based protects MS SQL at the server level, meaning all databases on a server. File-based protects only the specific MS SQL databases that you select. If you want to change the Backup Method to volume-based after Objects are selected, you	MS SQL

FIELD	DESCRIPTION	APPLICABLE GROUP TYPE
	must select a Server Object. A file-based Job can protect any Object.	
Make Full Backups Copy-only	Toggle on if you want full backups to be copy-only backups so they do not affect the differential base. Copy-only full backups do not take log backups even if they are scheduled by the policy.	MS SQL
Databases to Backup	Select which databases to back up.	MS SQL

Use these settings when you are setting up your Protection Group. Refer Product documentation for more options specific to the Protection Groups.

Appendix B: Connect IBM ICOS S3 Storage to Cohesity

If you get your S3-compatible storage from IBM Cloud Object Storage (ICOS), you can use that storage as a Cohesity target for your data protection and long-term retention needs. Wherever your ICOS S3 bucket actually resides, Cohesity's CloudArchive connects Cohesity DPlatform to that storage for archival, retention, geo-redundancy, compliance, and recovery.

To connect your ICOS bucket to Cohesity Data Cloud, you'll perform these tasks:

1. Create an IBM ICOS Bucket.
2. Create IBM ICOS Security Credentials and Service ID.
3. Register IBM ICOS to Cohesity.

Once you complete those steps, you can return to the instructions in the rest of this guide to archive your data to your ICOS bucket and recover your data from it.

Create IBM ICOS Bucket

The first step is to create an S3-compatible bucket in your IBM ICOS account.

1. Log in to your IBM Cloud Object Storage account.
2. In the form under **Create bucket**, enter a **Unique bucket name**, and select the appropriate **Resiliency**, **Location**, and **Storage class** for your needs.

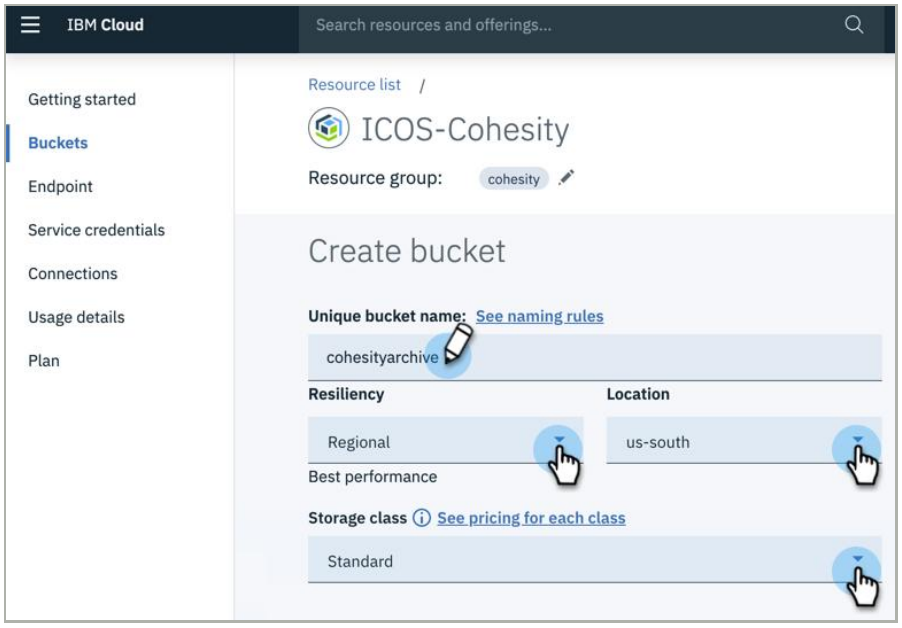

3. Click **Create bucket**.

Your new IBM ICOS bucket appears with your other ICOS buckets.

Create IBM ICOS Security Credentials and Service ID

Now that you have your IBM ICOS bucket, you need to create the credentials necessary to access and update it.

1. Navigate to your IBM ICOS account.
2. Select the IBM ICOS bucket that you created in the previous section.
3. Click **New credential**.
4. Enter:
   a. Name.
   b. Role. Select the Writer role at minimum.
   c. Service ID. Either select an existing Service ID with adequate permissions or click Auto Generate a new Service ID.
   d. Select Include HMAC Credential.

      IMPORTANT: This option is *required*, as it is the only way to capture the Access Key and Secret Access Key for Cohesity.

   e. Click **Add**.

Register IBM ICOS Bucket to Cohesity

Now that you have the IBM ICOS bucket and credentials, you're ready to connect it to Cohesity.

Capture IBM ICOS Credentials and Endpoint

Before you start, you'll need to capture these credentials: **Access-Key**, **Secret-Access-Key**, and **Endpoint**.

1. Navigate to your IBM ICOS account.
2. Select the IBM ICOS bucket that you created.
3. Click **Service Credentials**.
4. Identify the **Service Credential** name you need and click **View Credentials**.
5. Copy the values for **Access-Key** and **Secret-Access-Key**.
6. In the left menu, select **Endpoint**.
7. Navigate to the Resiliency and Location where you deployed the IBM ICOS bucket.
8. Copy the **Public Endpoint** for the desired location.

Register IBM ICOS Bucket as a Cohesity External Target

To register your IBM ICOS bucket as an External Target with your Cohesity cluster, refer Register S3 Bucket with Cohesity

Once you have registered your IBM ICOS S3-compatible bucket as an External Target in Cohesity Data Cloud, you can start protecting your data by creating a Protection Policy and Protection Group. When the need arises, you will be able use CloudArchive to recover your data from the External Target in Cohesity Data Cloud that uses your IBM ICOS bucket.

Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Jedidiah Sonavane is a Solutions Architect, STAT, and a part of Data Protect COE at Cohesity. He focuses on Service Providers/Organizations, Cloud Archive On-Prem, Gaia. His work proofs of concept, enterprise data protection, solution validation, solution design, testing, qualification, and ensuring software usability. He collaborates closely with teams to tailor solutions that meet customer needs while adhering to industry standards and best practices.

Other essential contributors include:

- Barb Abicht, Sr Technology Writer and Editor, TMSE
- Adaikkappan Arumugam, Sr Director, Solutions Architect
- Dayanand Sharma, Director, Product Management
- Kevin Hill, Cloud Solutions Architect
- Praveen Yarlagadda, Technical Director
- Radhani Guturi, Principal Engineer

Document Version History

VERSION	DATE	DOCUMENT HISTORY
2.0	Dec 2025	Content updates as of 7.2.2
1.2	July 2024	Republishing
1.1	June 2019	Content updates
1.0	May 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.