COHESITY

Archive Data to PoINT Archival Gateway S3-Compatible Object-based Tape Storage Using Cohesity DataProtect

Version 1.1

July 2025

ABSTRACT

Tape storage remains cost-effective for long-term data retention and is crucial for security, compliance, and legal requirements. Cohesity Platform™ and DataProtect™ offer robust on-premises enterprise data protection and storage solutions. Cohesity Platform™ extends its archival capabilities with the PoINT Archival Gateway S3-Compatible Object-based Tape Storage solution, simplifying the solution efficiently

Table of Contents

Figures

Tables

.

The Need for Archival to Tape

Organizations must adopt robust long-term data retention strategies to manage new data influx effectively. According to studies, most of the data is not frequently used but must be preserved for business or compliance reasons. Software-defined object storage offering an S3-Compatible REST API using standardized tape technology as the storage medium, like LTO, provides a solution to this challenge.

Archival to tape is crucial in enterprise backup solutions due to its cost efficiency, data security, and long-term reliability. Tape storage is highly durable, meeting mandatory compliance, regulatory, long-term storage requirements, and legal requirements, with minimal incremental costs of backing up to tapes in most environments. You can safely move the sensitive data that you want to retain for future reference or regulatory compliance to an archive environment. It also serves as an “air-gap” strategy to secure your data from threats and reduce primary storage consumption and ownership costs.

When choosing a tape archival solution, consider the following:

- Data Security
- Backup and Disaster Recovery
- Performance
- Reliability and Durability
- Capacity Requirement and Interoperability
- Scalability

Cohesity has partnered with PoINT Archival Gateway to deliver a comprehensive, data and tape-agnostic archival solution that simplifies long-term data retention and archival goals. This collaboration combines PoINT Archival Gateway’s enterprise-class tape archival systems with Cohesity’s cutting-edge, hyperconverged platform for managing data backups and archives.

The solution ensures seamless setup, scheduling, and management of archives on a PoINT Archival Gateway tape library through the Cohesity Platform, offering enterprises a scalable, secure, and cost-effective approach to long-term data storage and reliable data recovery.

This guide offers the solution to integrate PoINT Archival Gateway S3-Compatible Object-based tape solution with Cohesity, to ensure streamlined operations and enhanced data protection.

Long-Term Retention to Object-based Tape Features and Benefits

Cohesity's Archival to Object-based Tape provides many key features, each of which benefits organizations and their IT administration staff. Specifically:

Table 1: Long-Term Retention to Object-based Tape Features and Benefits

Features	Benefits
Policy-based archival	• Easy to use. • Archive unique data differently by mapping Protection Policies to the required SLA. • Reduce bandwidth and storage costs.
Cost Efficiency	Significantly lower cost ($/TB)
Off-site copies	• Geo-redundancy • Disaster recovery • "Air-Gap" enhancing data protection
Compression	Efficient data transfer and storage.
Recovery	• Instantly locate objects to Recover • Recover just what you need.
Encryption	Data is secure both in flight and at rest.

NOTE: This document covers only Cohesity DataProtect operations for archiving to tape using S3-Compatible Object-based Tape storage. For archiving to public cloud vendors, see guides for AWS, Azure, and Google Cloud Platform, and for archiving to NAS and S3-Compatible Storage, please refer to the above guides.

Use PoINT Archival Gateway S3-Compatible Object-based Tape with Cohesity DataProtect Platform

Long-term data and application retention is critical to organizations that seek to prevent data loss and meet security, legal, and compliance requirements. Cohesity DataProtect offers robust on-premises solutions for enterprise data protection and storage. Cohesity's CloudArchive and Cloud Recover bring data protection and recovery together in a single coherent solution, both on-premises and in the cloud, with S3-Compatible Object-based Tape Storage.

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 S3-Compatible managed tape libraries, eliminating 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 S3-Compatible Object-based Tape Storage Class uses S3 Glacier API interface to communicate between PoINT Archival Gateway (PAG) and Cohesity Cluster for Archival and Recovery.

Archival Process for S3-Compatible Object-based Tape Storage

Object-based Tape Storage Archival has two phases to perform Data archival from Cohesity backup.

Figure 1: Archiving Cohesity backup data to PoINT Archival Gateway S3-Compatible Object-based Tape Storage

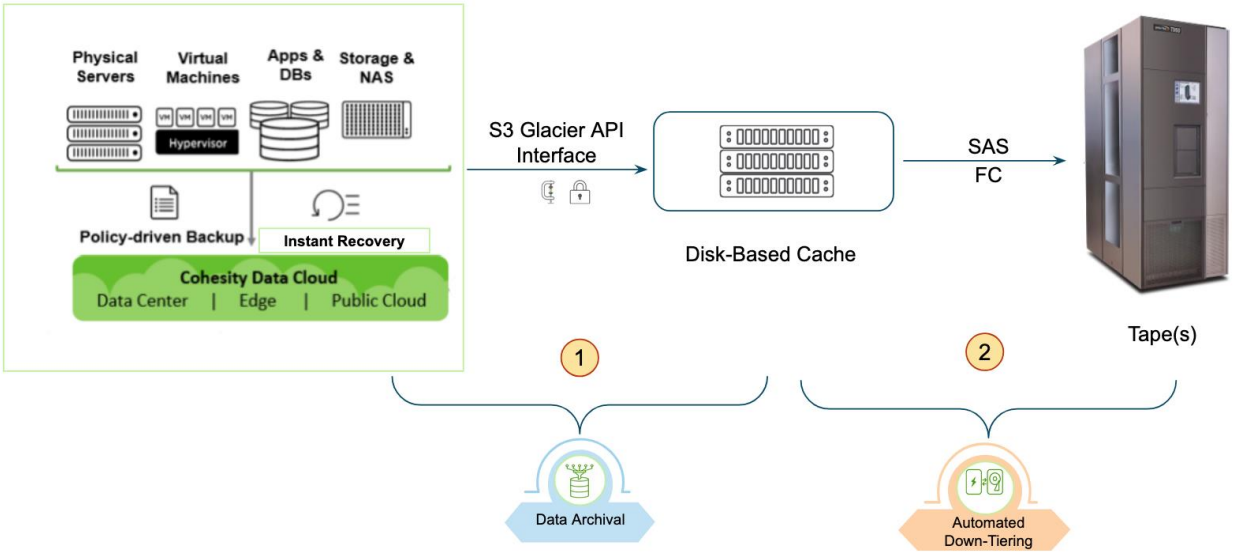

1. **Data Archival**: Once the backup to Cohesity Cluster is complete, the Archival data is ingested using S3-Compatible APIs (S3-Glacier) and lands in the Cache of PoINT Archival Gateway (PAG) Interface Node.
2. **Automated Down-Tiering**: Data is moved in parallel from the Cache to Tape drives.

Restore Process for S3-Compatible Object-based Tape Storage

When performing a restore, we perform the following actions to read the data.

Figure 2: Restoring data from PoINT Archival Gateway S3-Compatible Object-based Tape to the Restore Target

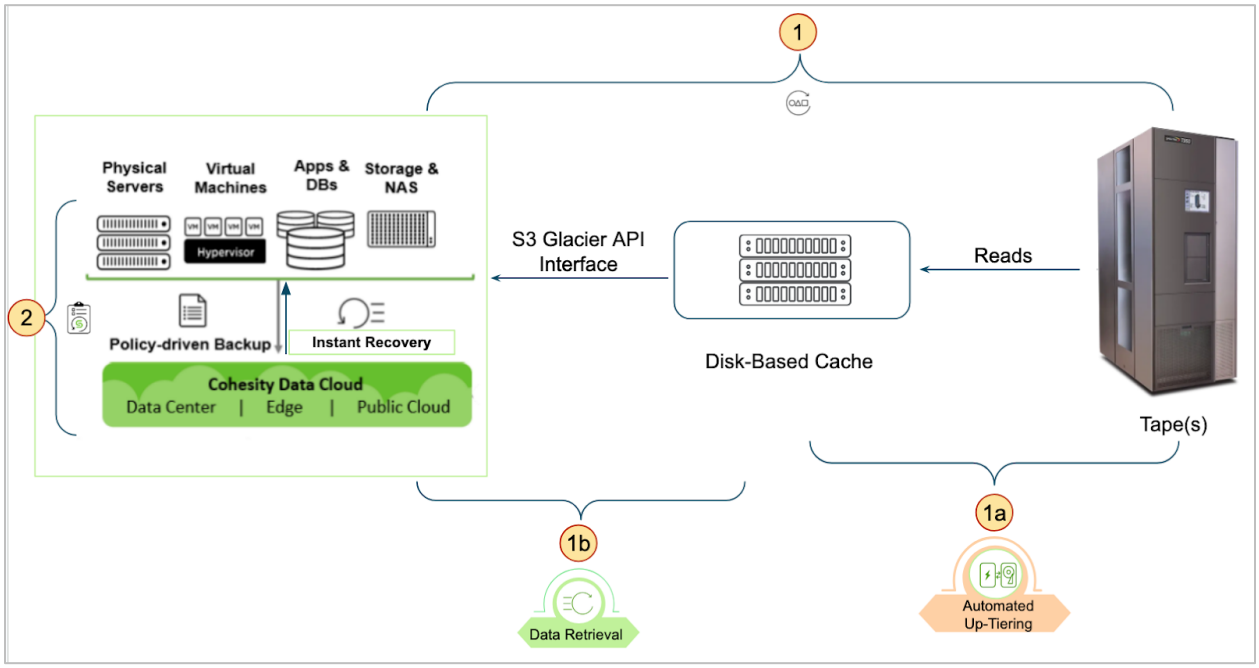

1. Data is Read from the Tape and Written to the Cohesity Cluster using S3-Glacier APIs
   a. **Automated Up-Tiering**: Data is Up-tiered from the Tape to the Cache.
   b. **Data Retrieval**: Data is Read from the PAG Cache and Written to the Cohesity Cluster.
2. **Data Restore**: Data is Read from the Cohesity Cluster and Written to the Restore Target.

Prerequisites

Archiving to PoINT Archival Gateway S3-Compatible Object-based Tape Storage is supported from the below versions:

Table 2: Minimum Cohesity and PoINT Archival Gateway software version needed

Software	Cohesity DataPlatform Version	PoINT Archival Gateway Version
Minimum Version Required	7.1.1	4.0 Update 3 (4.0.217)

Supported Storage Class for Archive to S3-Compatible Object-based Tape External Target

Cohesity Archive supports Archival to S3-Compatible Object-based Tape External Target from Cohesity version 7.1.1 onwards.

- Any S3-Compatible Object-based Tape Storage target that uses tape storage in the backend must be registered as an **Object-based Tape** Storage Class.
- These targets have capabilities similar to those of AWS S3 Glacier Storage Class.
- The Object-based Tape Storage Class uses S3 Glacier APIs for archival and recovery.

Considerations

The following are the critical considerations with S3-Compatible Object-based Tape Storage Solution.

1. The archival format is ***Always Full***, with Encryption and Compression enabled by default.
2. The recovery workflow retains the restored data from the tape on the S3 cache for a minimum of seven days. Please contact Cohesity Support to reduce the retention period of restored data on the PoINT Archival Gateway Cache.
3. The number of parallel streams is set to 10 by default. Contact Cohesity Support to change the values.
4. Does not support Deduplication and Incremental archives.
5. Does not support CloudRetrieve.
6. Does not support file-level recovery.

CloudArchive High-Level Workflow

At the high level, leveraging Archival to S3-Compatible Object-based Tape Storage involves below 4 tasks:

1. Create an S3 bucket on the PoINT Archival Gateway UI.
   - Capture your S3 bucket's Access Key ID and Secret Access Key.
2. Register the S3 bucket created above with the Cohesity Cluster as an S3-Compatible Object-based Tape External Target.
3. Archive your data to the External Target.
   a. Create a Cohesity Protection Policy.
   b. Create a Cohesity Protection Group.
4. Recover your data from the External Target.

Figure 3: Leverage PoINT Archival Gateway S3-Compatible Object-based Tape Storage with Cohesity

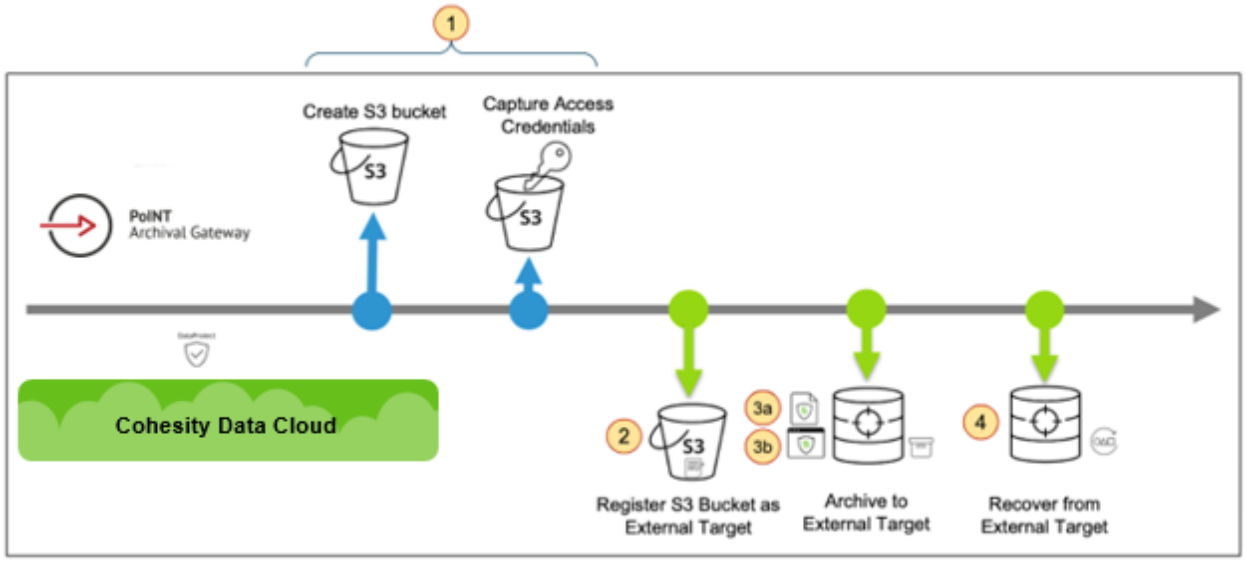

Create Your S3 Bucket on PoINT Archival Gateway

Creating an S3 Bucket on PoINT Archival Gateway involves creating an IAM user, the S3 bucket (Object Repository), and capturing the below required external target fields to register the S3 bucket as an External Target on the Cohesity Cluster.

Required External Target Fields to Register Your S3 Bucket

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

- Access Key ID
- Secret Access Key
- Bucket Name
- Endpoint
- Port on which the S3 bucket is exposed
- AWS Signature Version

To get started, we need to create an IAM user on the PoINT Archival Gateway UI console.

Create IAM Users on PoINT Archival Gateway Management Console

1. Log in to the PoINT Archival Gateway UI.
2. In the upper right corner of the UI, click the **Security & Auditing** button and select **Users.**

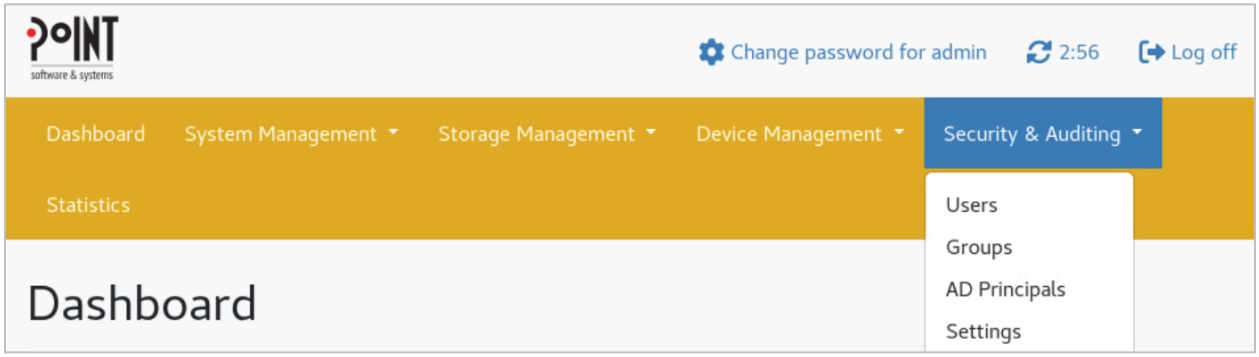

3. Under the Users banner, click **Create User**.

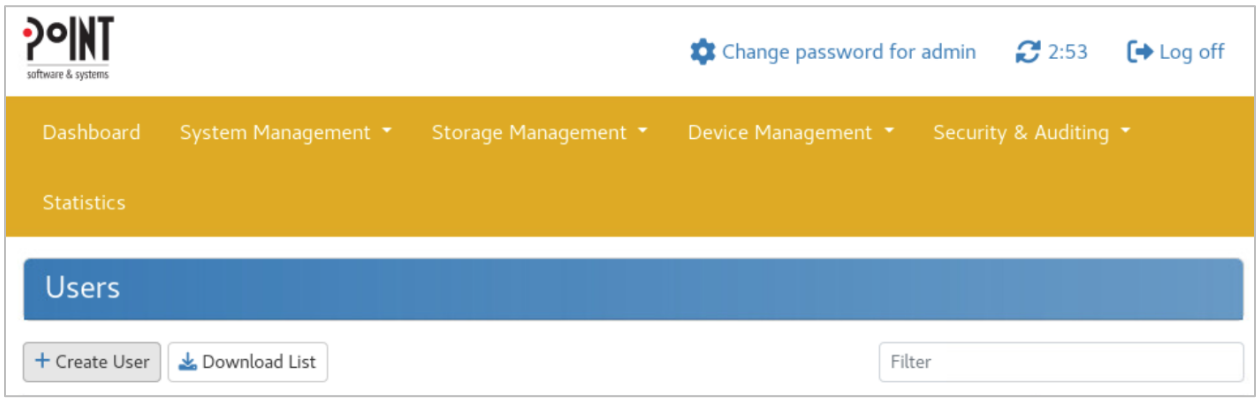

4. Enter the **Username and Password** for the new user. Click **OK.**

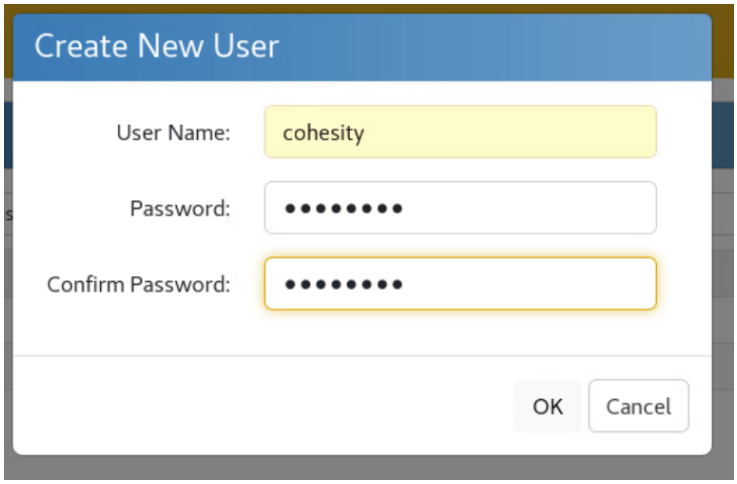

5. Select the user created, and click on the user, to configure user login settings.
6. Enable **Allow S3 Access**.

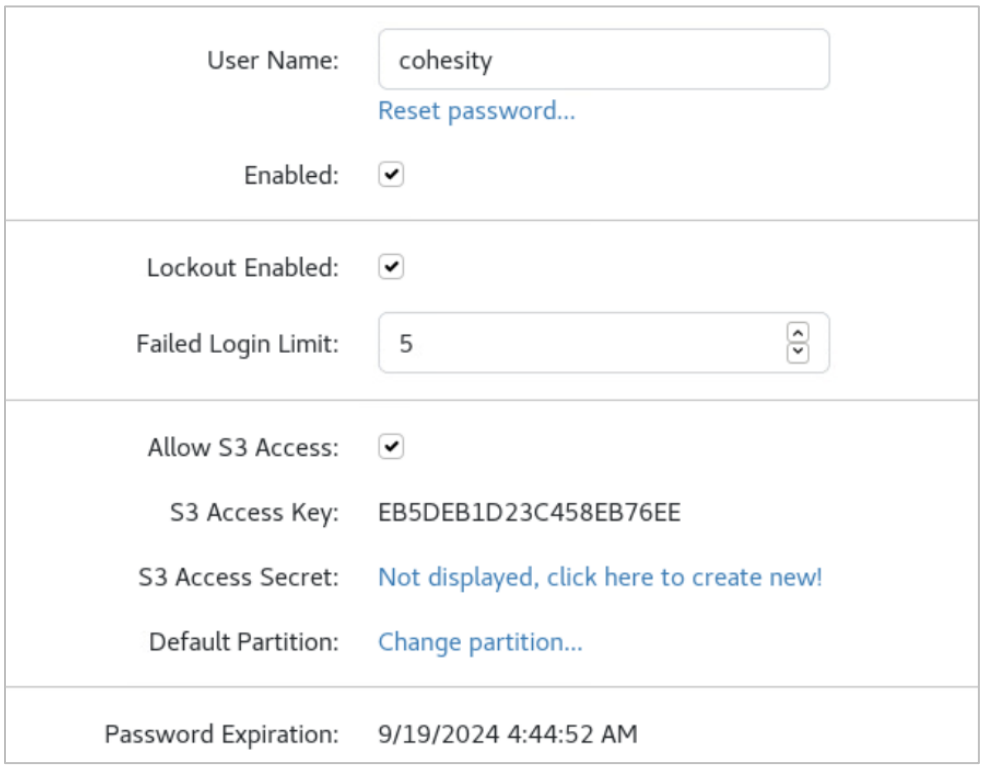

7. Copy the **S3 Access Key** and **S3 Access Secret Key** which is generated, and Click **Apply.**

> **NOTE**: This is the **only** time this secret access key can be viewed. You cannot recover it. Store the Secret Access Key safely.

This will be used to register this bucket as an external target archival.

Create Your S3 Bucket

1. In the PoINT Archival Gateway UI, click **Storage Management > Storage Partition**.
2. If there is no Storage Partition created, click **Create Partition.**

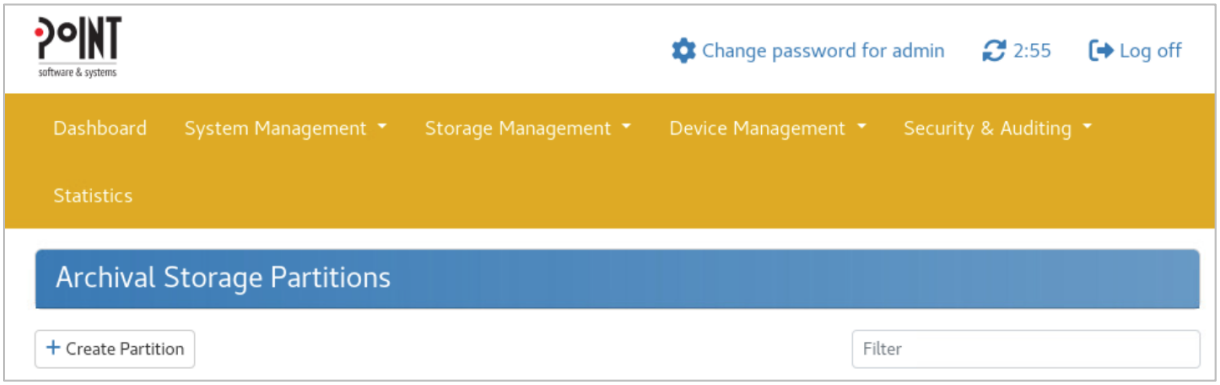

NOTE: Please refer to the PoINT Archival Gateway documentation for a detailed guide on Storage Partition and settings to be configured.

3. Enter the desired **Partition Name** and click **OK.**
4. Choose the Storage Classes as **Disk & Tape** and click **Apply.**

NOTE: The Storage Class must be Disk & Tape for the Archival Storage Partition, to enable S3 Glacier Storage Class for Tape.

5. Encryption will be enabled in External Target Registration on Cohesity UI, so leave it unchecked.

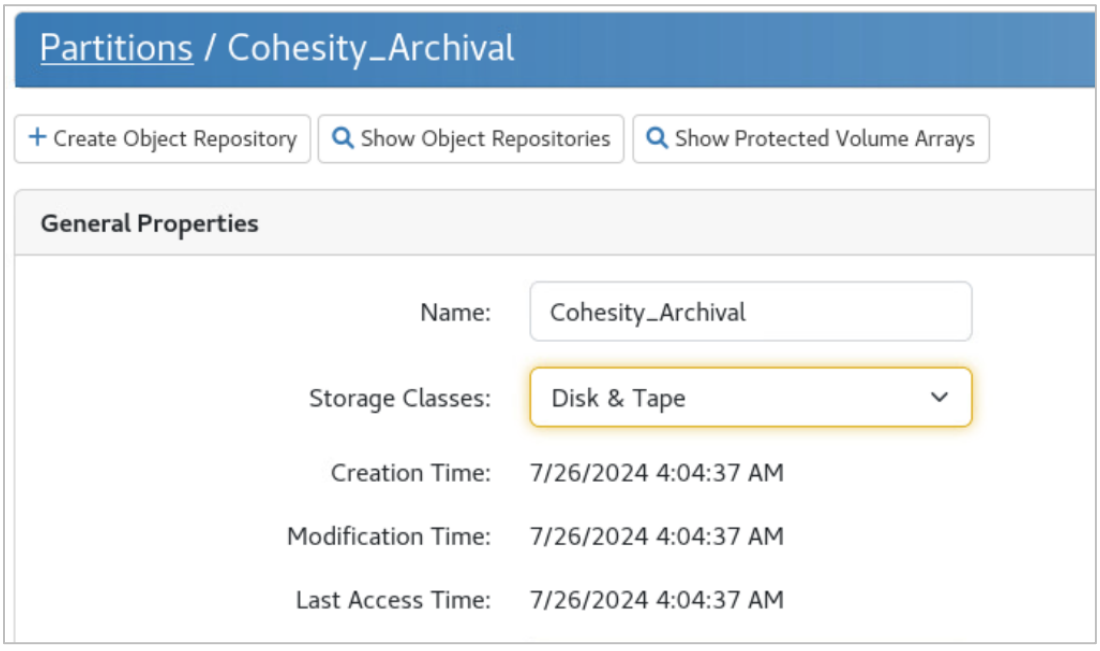

6. Once the Storage Partition is created, click on **Create Object Repository.**

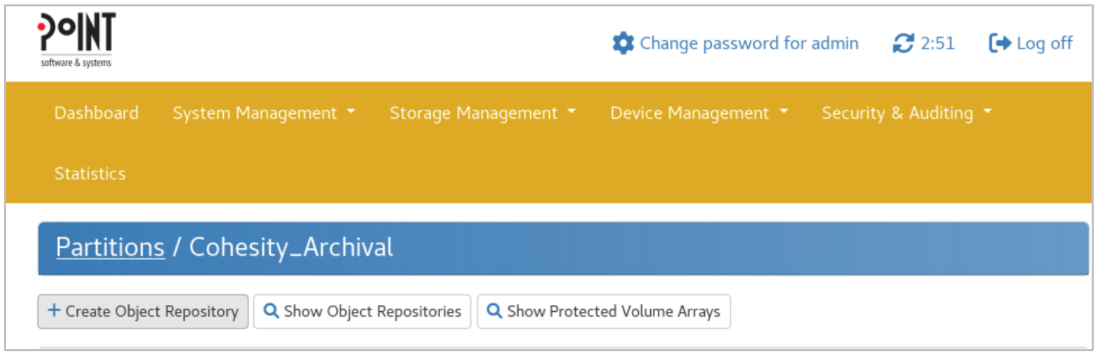

7. Enter the desired **Repository Name** and click **OK.**

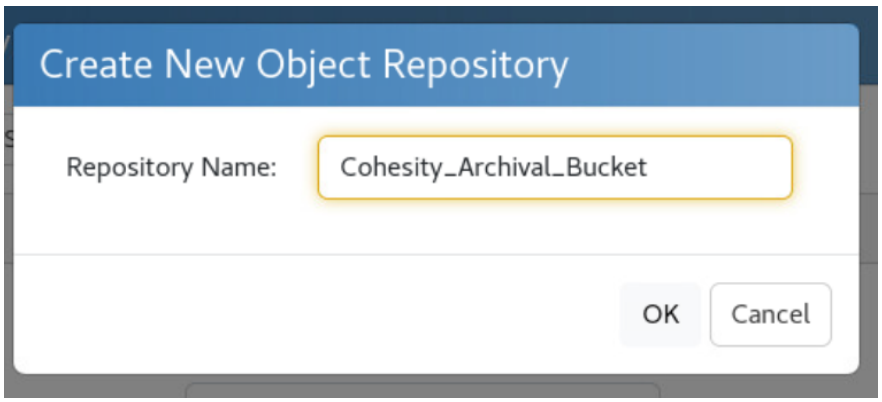

8. Click on the newly created **S3 Bucket (Object Repository)** and change the **S3 Storage Class for Tape to GLACIER**.

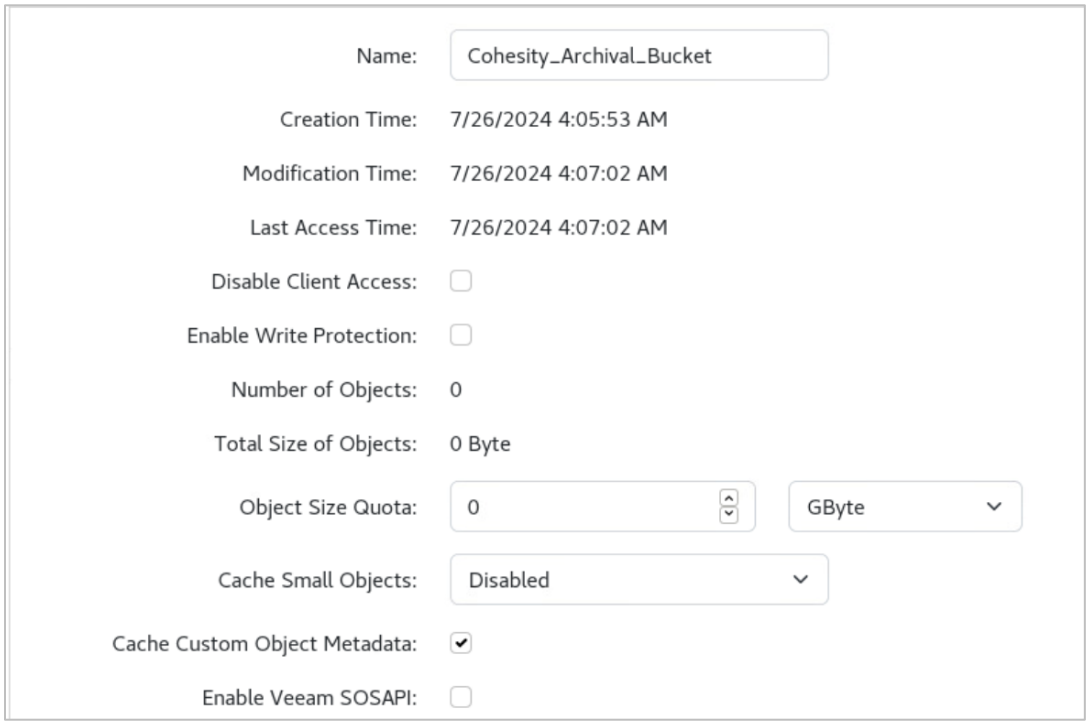

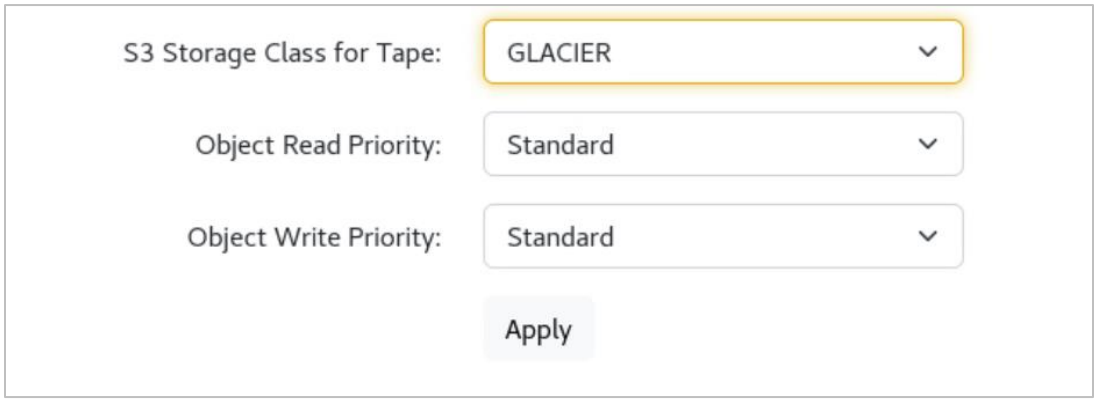

9. Click on the **Change Owner** button to select a user to own the bucket. The bucket owner sets permissions for the bucket.

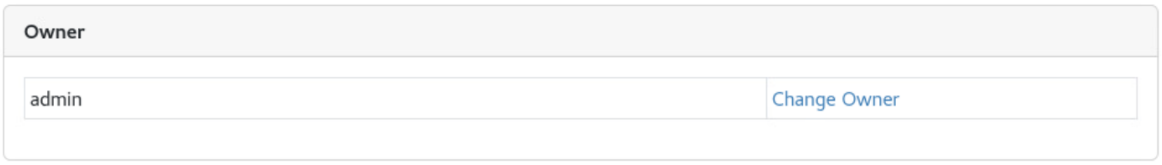

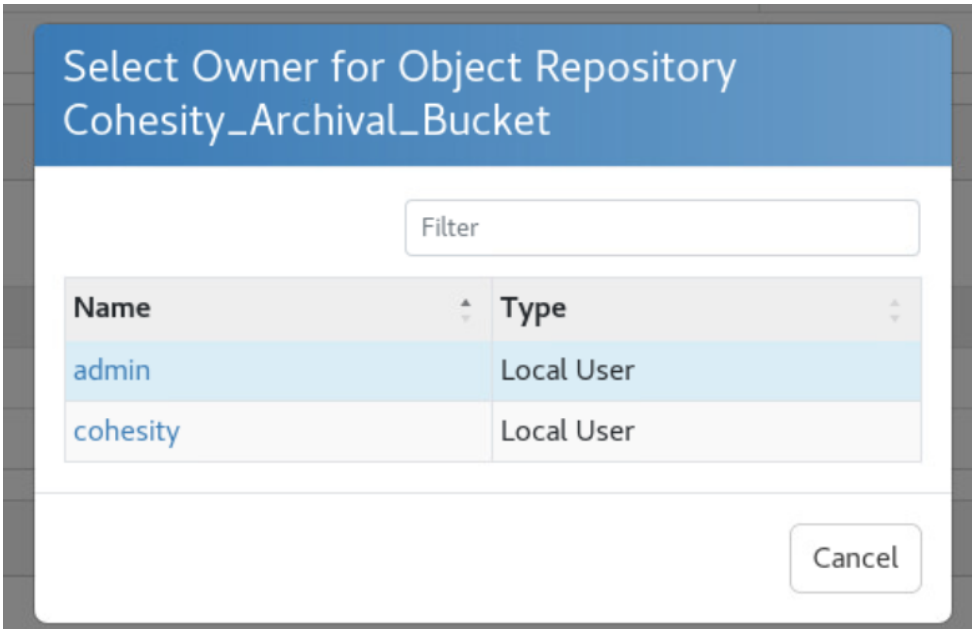

NOTE: Please refer to PoINT Archival Gateway documentation for the detailed guide on Object repository settings to be configured.

NOTE: The tape-based Storage Class uses S3 Glacier APIs for archival and recovery. So, no lifecycle needs to be configured to transfer the data.

10. Review the configuration, then click **Apply** to create the bucket.

NOTE: PoINT recommends following Amazon AWS documentation naming conventions and restrictions to maintain consistency and avoid potential issues. See AWS Bucket Naming Rules for more information

Protect Your Data

Now that you have the bucket created, and the IAM user created, you're ready to connect the S3 Bucket to the Cohesity Cluster as an External Target.

Register Your S3 Bucket as an External Target

To Register an External Target with your cluster, follow these steps.

1. Log in to Cohesity Cluster UI.
2. Click **Infrastructure** > **External Targets**.

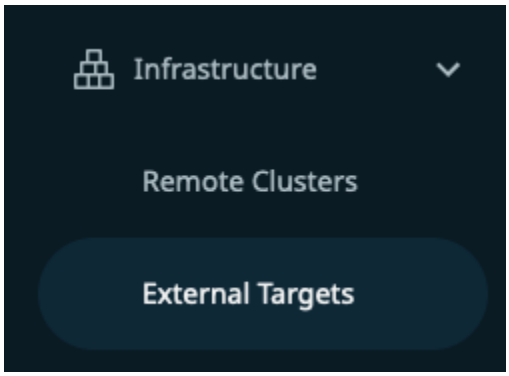

3. Click "**Add External Target**".

4. Choose Purpose as "**Archival**", Storage Type as "**S3Compatible**", and Storage Class as "**Tape Based**".
   In the form that opens:

- Enter the **Bucket Name**.
- Enter the **Access Key ID** and **Secret Access Key** that you captured when you created the users in PoINT Archival Gateway UI.
- Enter **Endpoint**, **Port**, and **Region**.
- **Secure Connection (HTTPS)**: Enabled by default.
- **AWS Signature Version 4**: PoINT Archival Gateway exposes the S3 bucket on Port 4443 and supports AWS Signature Version 4.
- Enter a unique **External Target** Name.

NOTE: Archival Format will be Always Full as it is S3-compatible object-based Tape Storage.

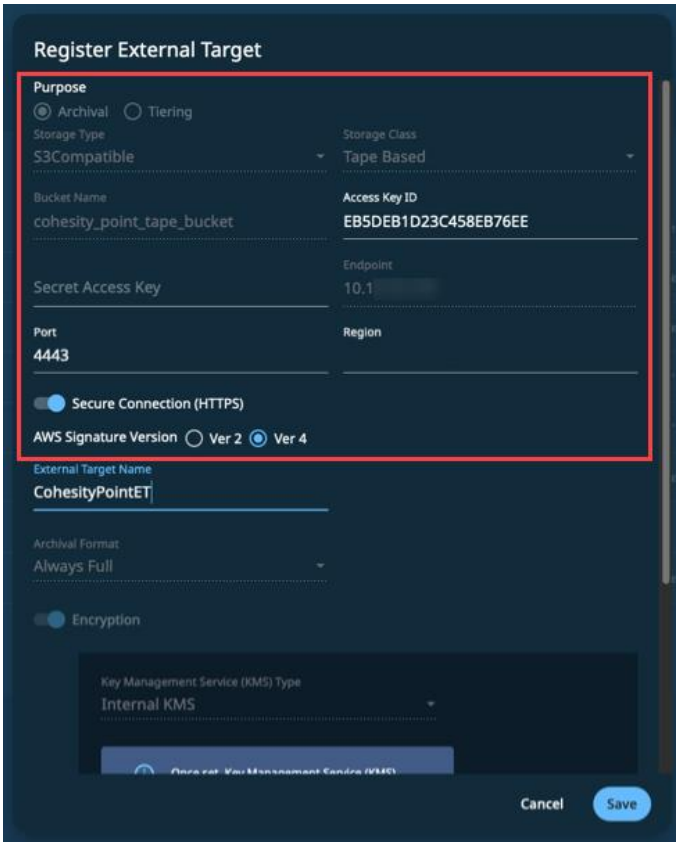

NOTE: The Region field can be left blank, as this is ignored in the backend.

5. Enable **Encryption** and choose between the options provided.

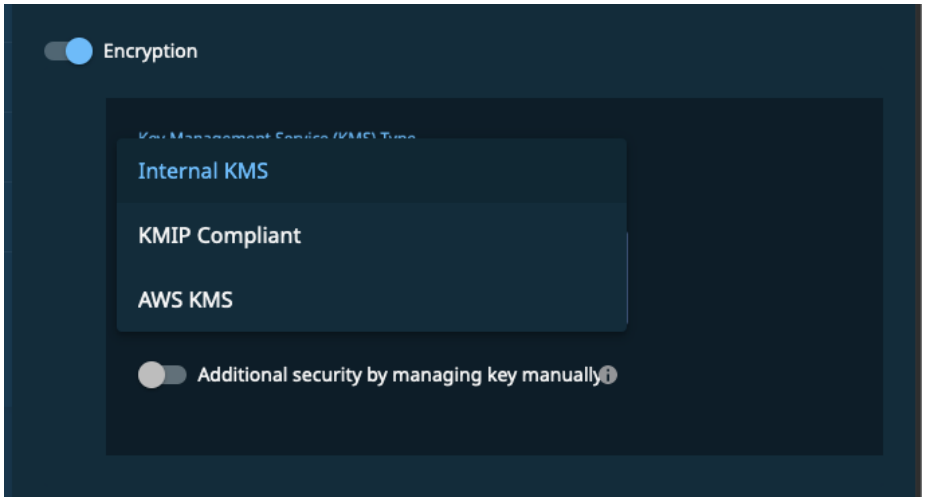

6. **Compression** is enabled by default**.**

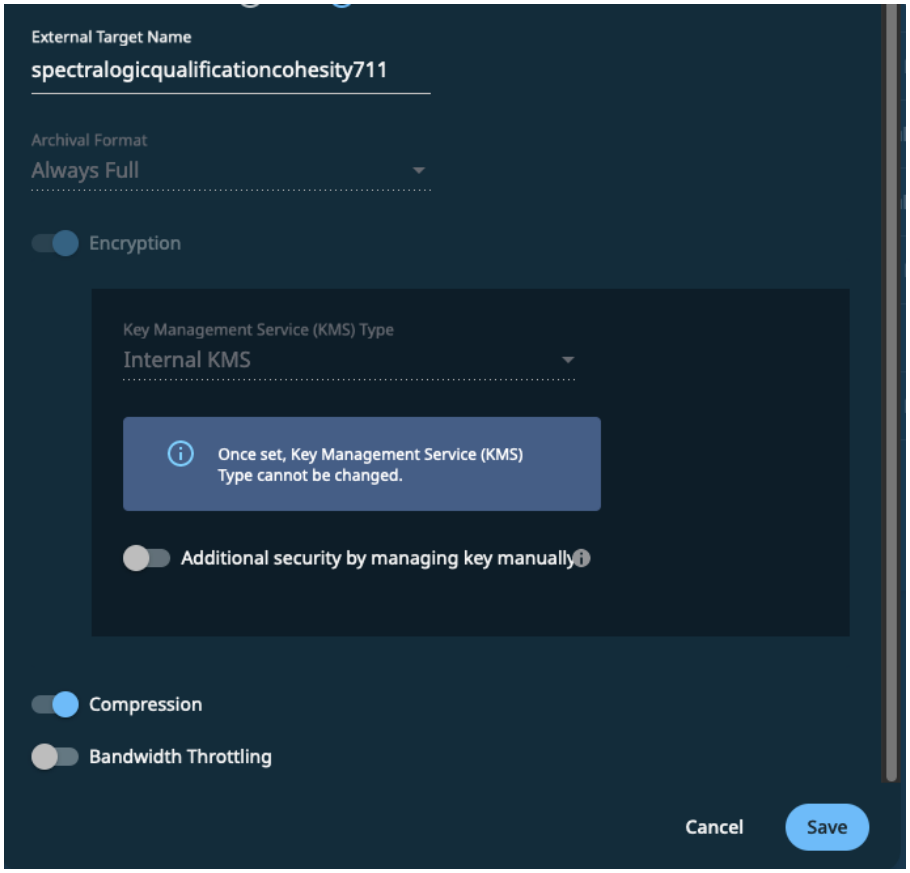

7. Click **Save.**

Your registered External Target is now available to select when creating a Cohesity Protection Policy for Protection Groups.

Create a Protection Policy

Once Cohesity Cluster registers your S3 bucket as an External Target, you will Create a Protection Policy according to your business needs. A protection policy is a reusable set of settings that define how and when objects are protected, replicated, and archived. It allows you to incorporate the External Target that you created above 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 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)?

Create a Protection Group

Once an external target is registered on Cohesity Cluster UI, we need to create a Protection Group to protect an object, which will be used to archive. Protection Groups combine operational requirements with the business requirements that are defined in a Protection Policy.

In the Protection Group, 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. 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 Protection Group can have only one policy.

Performance: Optimize Archival Throughput

The archival process involves several steps, including reading data from the source, compressing it, storing the primary copy in the Cohesity Cluster, and then writing it to archival storage. If not managed effectively, these tasks can increase disk I/O and CPU utilization, potentially affecting the cluster's overall performance.

NOTE: We by default allow 10 streams to perform Archival to PoINT Archival Gateway S3-Compatible Object-based Tape External Target.

To achieve optimal archival performance, collaborate with Cohesity support to take advantage of Cohesity and PoINT Archival Gateway Archival Performance.

Recover Your Data from The Tape Archive

When recovering your archived data, Cohesity Cluster allows you to Restore Entire Object(s) (VMs, databases, NAS, etc.).

Please review the considerations mentioned here.

NOTE: PoINT recommends keeping the tape in the tape library. Restore fails if the required tape is offline.

Restore Timeline and Workflow for Common Scenarios

When a user initiates the Restore on the Cohesity Cluster UI, the actions performed to retrieve the data are below. Let's take a look at the Restore Timeline.

Figure 4: Restoring data from PoINT Archival Gateway S3-Compatible Object-based Tape to the Restore Target

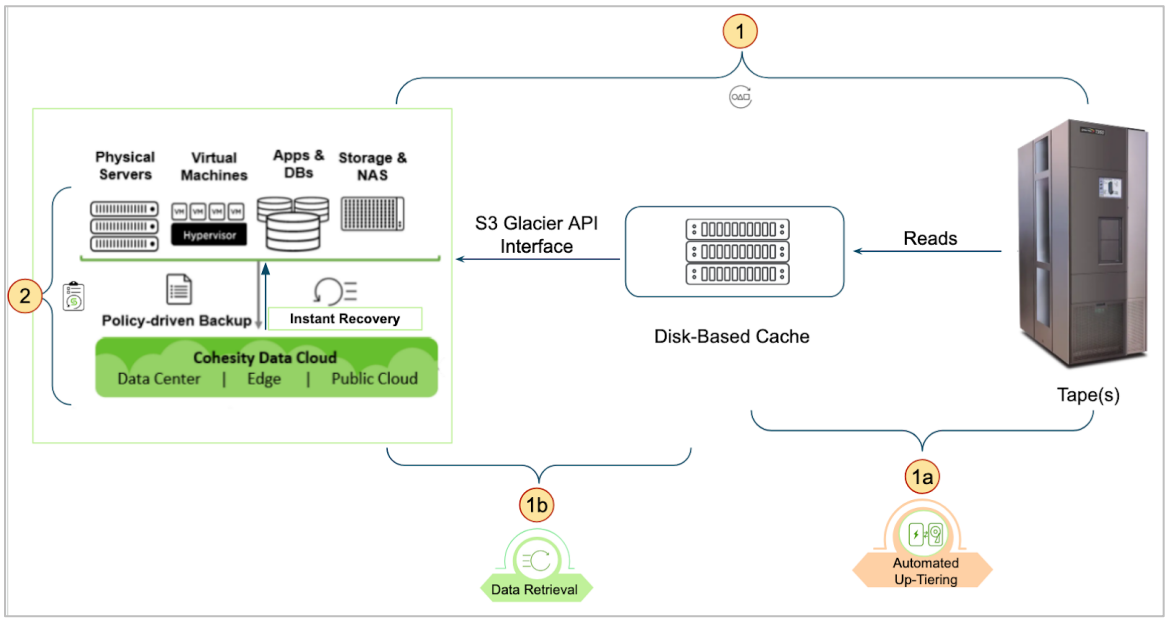

Table 3: Restore Workflow for Common Scenarios

	Tape Data in Cache	Tape is Online
User Intervention	✖	✖
Restore Time Taken	Quick	Take additional time (due to Up-Tiering Data from Tape to Cache)
Steps	1b, 2	1a, 1b, 2

Below are the illustrations for 2 different common scenarios:

1. **Tape Data in Cache**: When a user initiates a restore, the restore starts immediately as the data is present in Cache, therefore there is no need to read data from Tape.

   There are 2 Steps involved:

   a. **Data Retrieval**: Data is Read from the Cache and Written to the Cohesity Cluster.

b. **Data Restore**: Data is Read from the Cohesity Cluster and Written to the Restore Target.

Figure 5: Tape Data in Cache

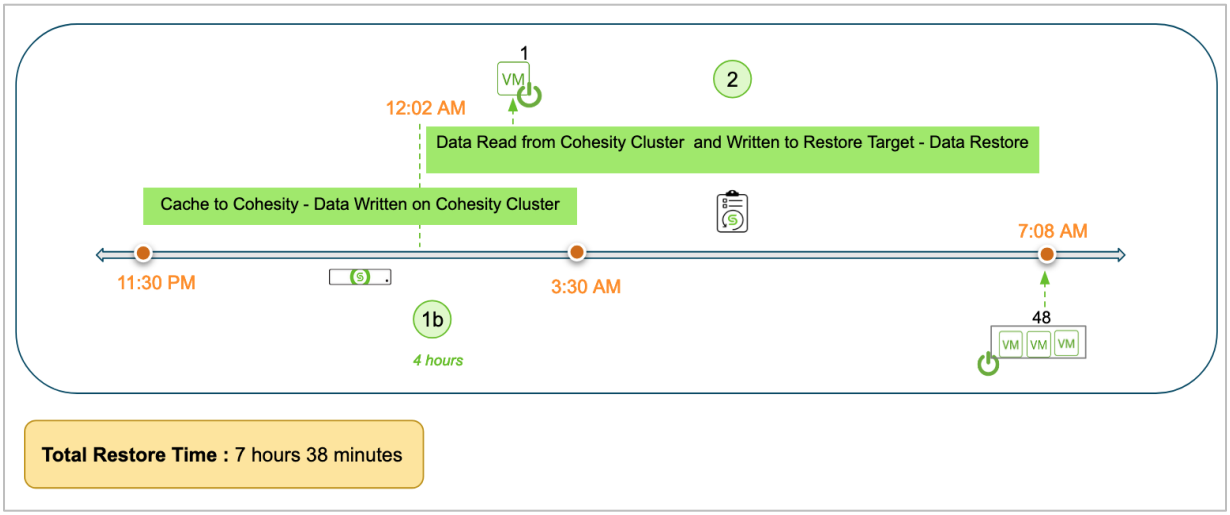

2. **Tape is Online**: When a user initiates a restore, the requested data is not present in the Cache, requiring Up-tiering of data from Tape.

   There are 3 Steps involved:

   a. **Automated Up-Tiering**: Data is Up-tiered from Tape to Cache.

   b. **Data Retrieval**: Data is Read from the Cache and Written to the Cohesity Cluster.

   c. **Data Restore**: Data is Read from the Cohesity Cluster and Written to the Restore Target.

Figure 6: Tape is Online

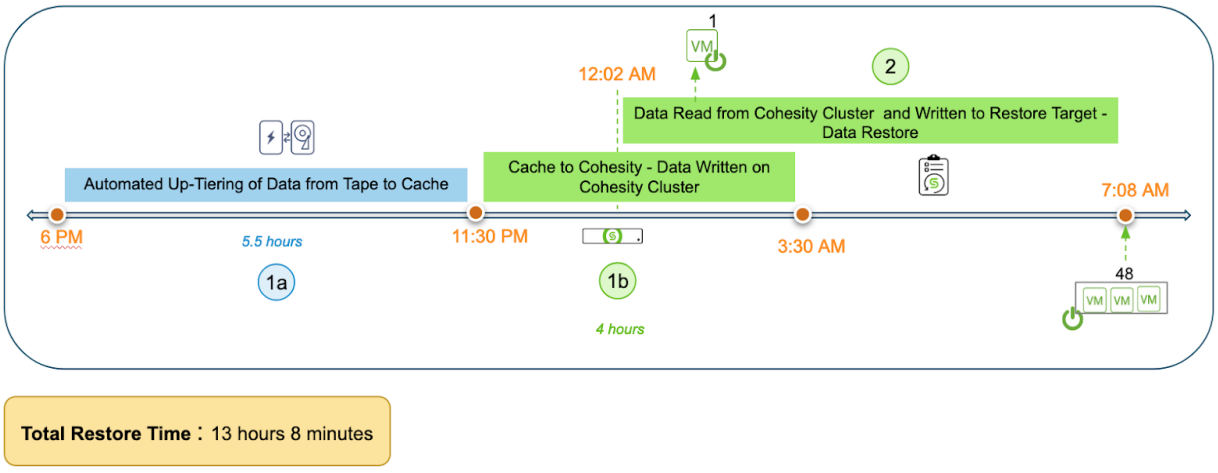

Managing PAG Cache

Effective cache management is essential when performing large restore operations from tape. When multiple large restores are initiated, cache space can quickly become exhausted, leading to restore failures. In the following scenarios, we highlight how cache capacity directly impacts the success of restoration processes and provide strategies to optimize cache usage for smoother restore operations.

Scenario 1: PAG Cache Capacity Exceeded - SQL DB Restore Fails Due to Insufficient Space on Cache

In this scenario, Restoring 100VMs of size 40 TiB and SQL Databases of size 70 TiB are triggered sequentially, but the cache becomes overloaded when Restoring SQL DBs requires more space in the PAG Cache than is available. This highlights the consequences of insufficient cache capacity and the need to manage restore operations with caution.

- Cache Capacity: 100 TiB
- Restore VMs Size: 100 VMs (40 TiB)
- Restore SQL DBs Size: 70 TiB
- Cache Retention: 7-day default period

Timeline	Action	Details	Cache Usage (100 TiB Limit)
Day 1	VM Restores Begin	Current Cache Utilization	10 TiB
Day 3	VM Restore In-Progress	Data from Tape to Cache	50 TiB
Day 4	VM Restore Completes	Data held in cache for 7 days	50 TiB
Day 4	SQL DB Restore Begins	Requires 70 TiB of Cache space	120 TiB (Exceeds 100 TiB Limit)
Day 6	Restoring SQL DBs Fail	Cache Capacity Exceeded	Restore fails

Outcome: Restoring SQL DBs fails because it exceeds the available cache space after Restoring VMs has already filled the cache. The system cannot process the restore due to insufficient capacity.

Scenario 2: Efficient Cache Reclamation Strategy Ensures SQL DB Restore Succeeds

In this scenario, Restoring of VMs completes and cache space is reclaimed after 7 days, allowing Restore SQL DB to start with adequate cache capacity. This demonstrates how cache management, including reclaiming space, ensures that subsequent restore operations can proceed without hitting capacity limits.

- Cache Capacity: 100 TiB
- Restore VMs Size: 100 VMs (40 TiB)

- Restore SQL DBs Size: 70 TiB
- Cache Retention: 7-day default period

Timeline	Action	Details	Cache Usage (100 TiB Limit)
Day 1	Restore VMs Begins	Initial Cache Load	10 TiB
Day 3	Restore VMs In-Progress	Data from Tape to Cache	50 TiB
Day 4	Restore VMs Completes	Data held in cache for 7 days	50 TiB
Day 12	Cache Reclamation	Automatic reclamation of space in Cache	10 TiB
Day 13	SQL DB Restore Begins	Requires 70 TiB of Cache space	80 TiB
Day 20	Restoring SQL DBs Succeeds	Restoring SQL DBs completes successfully	80 TiB

Outcome: Restores complete successfully, as sufficient capacity exists to hold the required data.

Important Note on Cache Retention

IMPORTANT: When a Restore is initiated, the object restored from the S3-Glacier / Tape is kept in the PoINT Archival Gateway Disk Cache for 7 days by Cohesity (this is to avoid the need to uptier the data multiple times if the restore takes longer/fails/cancels for any reason). Please contact Cohesity Support to change the configuration setting to reduce the number of days the restored data is kept in the Cache.

Takeaway and Recommendations

To prevent cache capacity exceeding and ensure smooth restore operations, consider the following strategies:

1. **Increase Cache Capacity**: For environments with frequent large restores, increasing cache capacity can accommodate more data without running into space issues.
2. **Adjust Cache Retention Period**: Contact Cohesity support to reduce the default 7-day cache retention period.
3. **Stagger Restores**: Space out restores to avoid overloading the cache, allowing it to recover between operations.

Cache space management is essential when performing large restores. By following the strategies highlighted above, you can ensure that restore operations are completed successfully without exceeding cache limits. Effective cache handling not only improves performance but also ensures business continuity by minimizing restore failures due to capacity breaches.

Appendix

CloudArchive Terminology

There are several important terms to understand as you learn how CloudArchive works.

Table 4: CloudArchive Terminology

Term	Definition	Notes
Cohesity Platform	Cohesity Platform 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 Platform is a scale-out solution based on a unique distributed file system, SpanFS®.	
Cluster	An instance of Cohesity Platform.	
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 Platform.
Protection Group	Defines operational requirements, such as which source objects to protect, the Protection Policy to use, and operational considerations like indexing, exclusions, inclusions, alerts, app consistency, and more	Each Protection Group has a schedule of Group Runs, and each archive is a collection of those 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.	
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	

Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Jedidiah Sonavane is a Solutions Architect at Cohesity. In his role, he focuses on enterprise data protection, solution validation, solution testing, solution qualification, and software usability for Cloud, Object-based Tape, Service Provider and Multitenancy Solutions.

Other essential contributors included:

- Adaikkappan Arumugam, Sr Director, Product Solutions

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.0	Nov 2024	First release
1.1	July 2025	Re-Publishing

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.