COHESITY

Version 1.3

October 2024

Integrated Data Protection with Cohesity and HPE Nimble Storage

Bringing Together Simplified VM Backup Using Native Integration with HPE Nimble Storage

ABSTRACT

Cohesity and HPE Nimble Storage integrated data protection solution stands out as a better way to back up data and applications with rich features and benefits. What's more, you'll be able to archive the backups to any public cloud or tape storage for long-term retention and replicate them to a different location for disaster recovery.

Table of Contents

Figures

Tables

Introduction to HPE Nimble Storage Integration

Cohesity and HPE have partnered to provide native integration of Cohesity's platform with HPE Nimble Storage for virtualized environments. With this integration, you can protect applications running on an HPE Nimble Storage array using array-based snapshots. This results in more efficient backup and recovery, as well as enhanced application performance.

Cohesity and HPE have certified fit-for-purpose HPE servers to run Cohesity's platform and DataProtect, backed by HPE interoperability assurance. This best-of-breed solution combines the web-scale simplicity and efficiency of Cohesity software with the power and density of HPE servers. As a result, customers can more efficiently and effectively manage unstructured data growth, acquire new insights, and reduce costs and complexity with a single, integrated solution.

This solution guide highlights Cohesity integration with HPE™ Nimble™ storage for backup and replication. With this integration, your VMware vSphere® compute resources are spared, and the snapshots are offloaded to HPE Nimble Storage to make the solution efficient and robust.

Use cases covered for this solution include:

- With Cohesity being the backup target, the hardware appliance can be on Cohesity models or HPE-qualified server models.
- HPE Nimble Storage is the primary storage source.
- VMware virtual machines are hosted on HPE Nimble Storage via the iSCSI and FC protocols.
- Using HPE Nimble Storage array snapshots during virtual machine backups.
- Backup performance comparison between HPE Nimble Storage snapshot integration and VMware snapshot.

Intended Audience

The target audience for this document includes storage and virtualization administrators, consulting data center architects, field engineers, and administrators who are designing, implementing, and maintaining application backups using Cohesity and HPE Nimble Storage. A working knowledge of VMware vSphere, servers, storage, backups, networks, and data center design is assumed but it is not a prerequisite for reading and understanding this document.

Cohesity-HPE Nimble Storage Solution Overview

Figure 1: Cohesity's Integrated Solution with HPE Nimble Storage

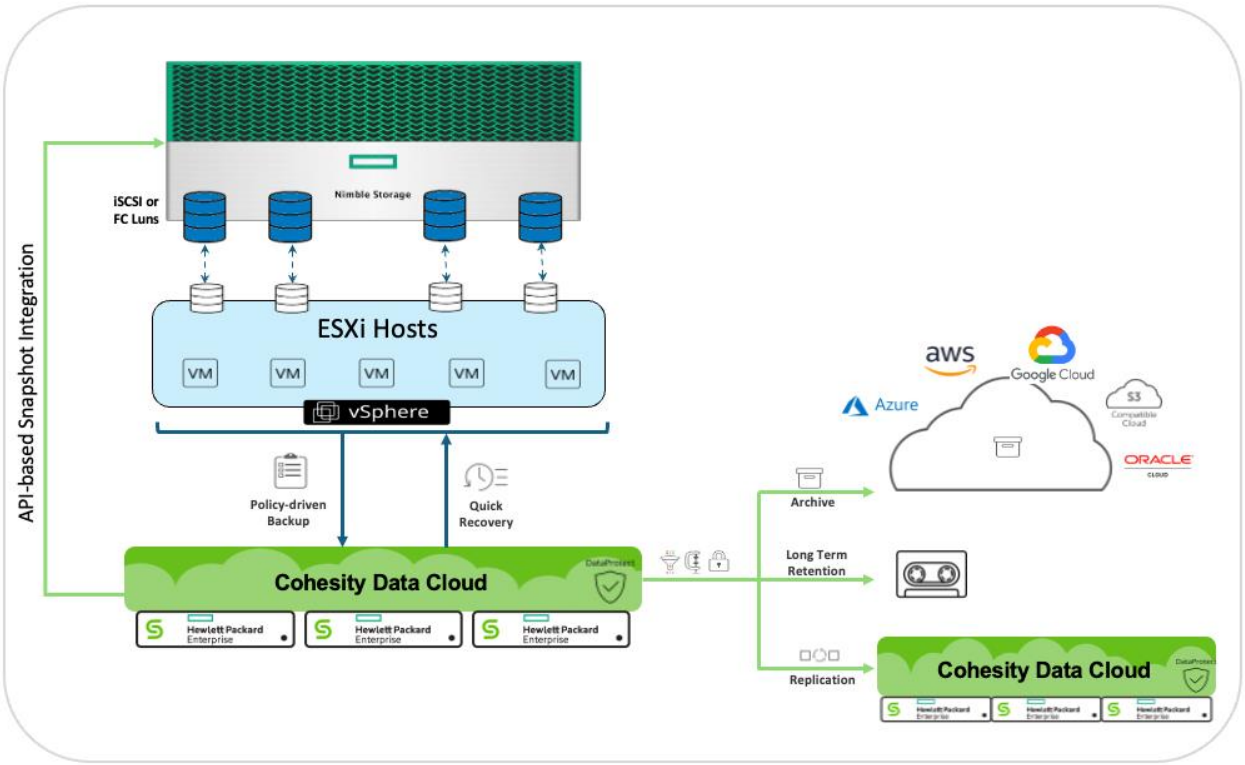

Cohesity Overview

Cohesity consolidates secondary data and applications, including backups, files, objects, test/dev, and analytics, on a single, software-defined platform. Built on web-scale architecture, Cohesity is a scale-out solution based on a unique distributed file system, SpanFS™.

While most organizations initially use Cohesity to overcome mass data fragmentation by simplifying data protection, its flexible architecture allows easy expansion to additional use cases such as NAS and object storage workloads, further increasing operational simplicity and improving TCO (Total Cost of Ownership). Cohesity works on-premises (on qualified Cisco, HPE, Dell, or Cohesity C-Series platforms), in the public cloud, and in remote and branch offices on hypervisors of your choice, such as VMware and Hyper-V™.

HPE Servers for Cohesity

Cohesity's software can be installed on select HPE servers. The combination of a small footprint, energy efficiency, and flexible options makes HPE servers an optimal platform for Cohesity's software-defined platform. For more details, see HPE Solutions with Cohesity.

HPE Nimble Storage All Flash Array

HPE Nimble Storage arrays are engineered for high performance (using flash) and lowered costs (using dense, capacity-optimized disks). HPE Nimble Storage All Flash Arrays combine a flash-efficient architecture with HPE InfoSight predictive analytics to achieve fast, reliable access to data and 99.9999% guaranteed availability. Radically simple to deploy and use, the arrays are cloud-ready, providing data mobility to the cloud through HPE Cloud Volumes.

HPE Nimble Storage arrays provide:

- Accelerated performance for higher throughput (I/Os per second) and sub-millisecond latencies.
- Higher storage efficiency to reduce the storage footprint by 30-75%.
- Non-disruptive scaling to fit changing application needs through increased performance, capacity, or both.
- Maximized data and storage availability with integrated data protection and disaster recovery.
- Simplified storage management and reduced day-to-day operational overhead.

Cohesity and HPE Nimble Storage Integration — Features & Benefits

Cohesity's integration with HPE Nimble Storage can protect applications running on HPE Nimble Storage using array-based snapshots. This results in more efficient backups as well as enhanced application performance.

Table 1: Features and Benefits of Cohesity Integration with HPE Nimble Storage Arrays

Features	Details
Reduced VM Stun	The VM stun that takes place during VM snapshot creation, deletion, and delta commit is completely avoided in Cohesity's integration with HPE Nimble Storage.
Reduced VM backup time with Cohesity MegaFile	To avoid bottlenecks when a VM is assigned to a single node, the MegaFile feature splits each virtual disk in the VM into parts, 'MegaFile chunks,' and distributes them to all nodes in the cluster. As a result, as the number of nodes in a Cohesity cluster grows, the time to ingest decreases dramatically.
Reduced VM snapshot lifespan with HPE Nimble Storage Integration	The VM snapshot hold time is reduced, and the snapshot is offloaded to an HPE Nimble Storage array, where the host resources are less utilized.
Simple data protection	Simplify backup environments by eliminating the need for media servers and master servers.

Features	Details
Application-consistent protection	Cohesity can perform application-consistent backups to ensure the database can recover faster and avoid data corruption. The system provides 'always-on' availability.
Distributed platform	You can scale out capacity and linear performance simply by adding nodes to the cluster. Eliminate the need for forklift upgrades and data migrations.
Native cloud integration	Integrate with AWS, Microsoft Azure, and/or Google Cloud Platform (GCP) for long-term data retention, data tiering, and disaster recovery.
Copy data management	Ability to spin up clone database copies from backups for test and dev environments. Cohesity acts as an NFS target for Oracle data files, control files, and redo logs.
Lower TCO (Total Cost of Ownership)	Our solution consolidates backup software licenses, media and catalog servers, and backup targets. What's more, global deduplication, compression, and snapshots dramatically reduce physical storage usage. Pay-as-you-grow expandability that reduces the need to over-provision.

Reduced VM Backup Time with Cohesity MegaFile

When backing up VMs, the primary challenges are VM stun times and application I/O timeouts, which are caused by long VM snapshot lifespans.

The longer your backup operations (i.e., VM snapshot lifespans) are, the greater your VM stun times, especially during the VM snapshot commit (delete) phase, which can cause application I/O timeouts. Given that the backup copy operation is a function of the total data that is read and then copied, the VM snapshot lifespan grows with the size of the data being backed up, as illustrated in Figure 2 below.

Figure 2: VM Snapshot Lifespan when Reading from VM Snapshot

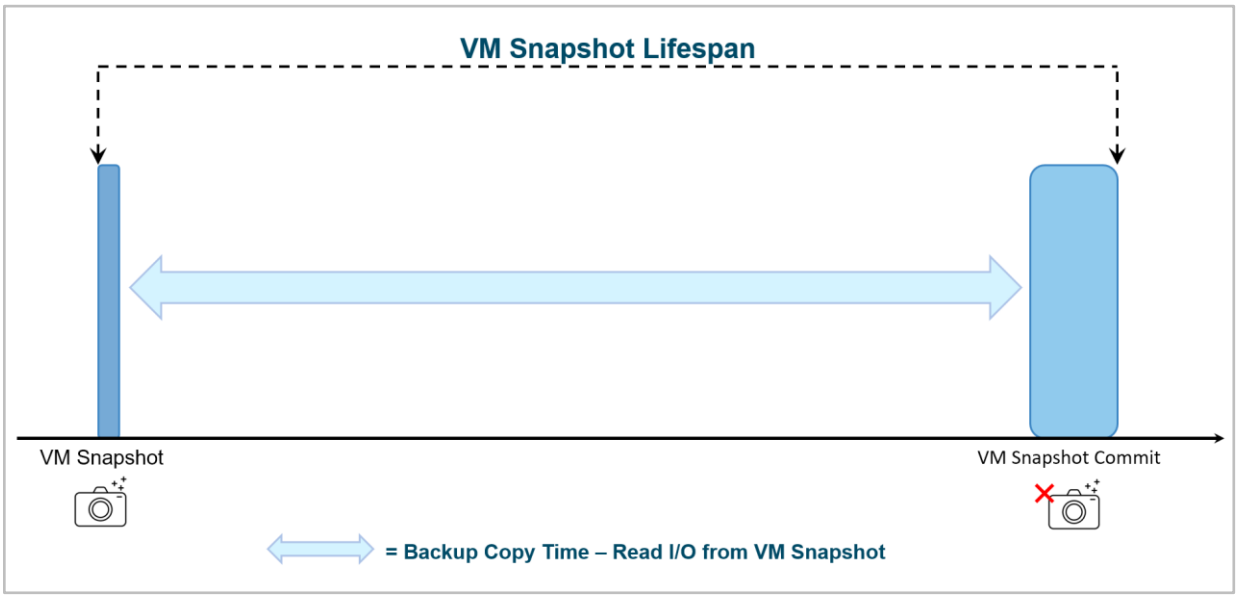

Cohesity natively alleviates the VM snapshot lifespan with the MegaFile feature. The MegaFile feature reduces the time to ingest a single large VMDK by splitting it into several smaller parts (MegaFile chunks) and distributing them to all nodes of the cluster, where they are ingested in parallel.

For example, if you have a 4-node cluster and a 2 TB VMDK, it will be split into four MegaFile chunks, each of which is 512 GB. In this example, we reduce the backup time by 4x, as illustrated in Figure 3 below. If you have an 8-node cluster and a 2 TB VMDK, then it is split into eight MegaFile chunks, each 256 GB, and the ingest time is reduced by a factor of 8. As a result, as the number of nodes in a Cohesity cluster grows, the time to ingest decreases dramatically, which reduces the VM snapshot lifespan.

Figure 3: Reduced VM Snapshot Lifespan with Cohesity MegaFile

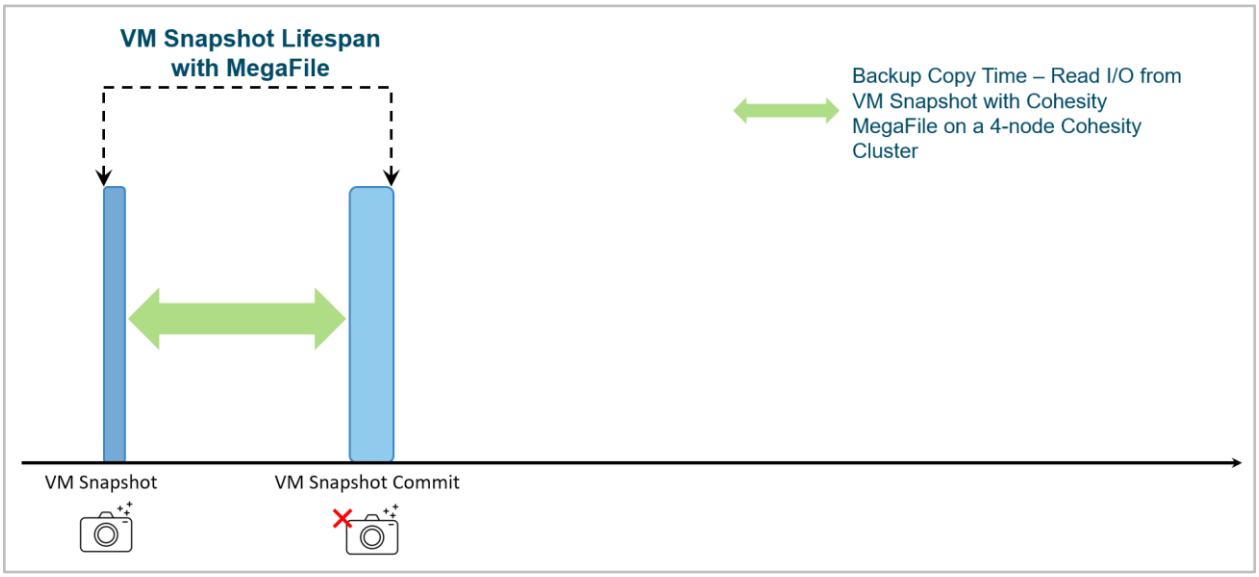

Reduced VM Snapshot Lifespan with HPE Nimble Storage Integration

While the time savings with Cohesity MegaFile are dramatic, they are still a function of the total size of the data and the number of nodes in your Cohesity cluster. To nearly eliminate the VM stun times, our integrated solution offloads the backup copy operation from a VM snapshot to a storage array snapshot. Instead of reading the data from the VM snapshot and keeping it open, with our integration, we create the VM snapshot and then create a corresponding storage array snapshot of the underlying datastore on which the VM is hosted. We then delete the VM snapshot as soon as the storage array snapshot is successfully created and copy the VM data from the storage array snapshot using the Cohesity MegaFile feature, thereby reducing the backup time still further.

Figure 4: VM Snapshot Lifespan when Reading from HPE Nimble Storage Array Snapshot

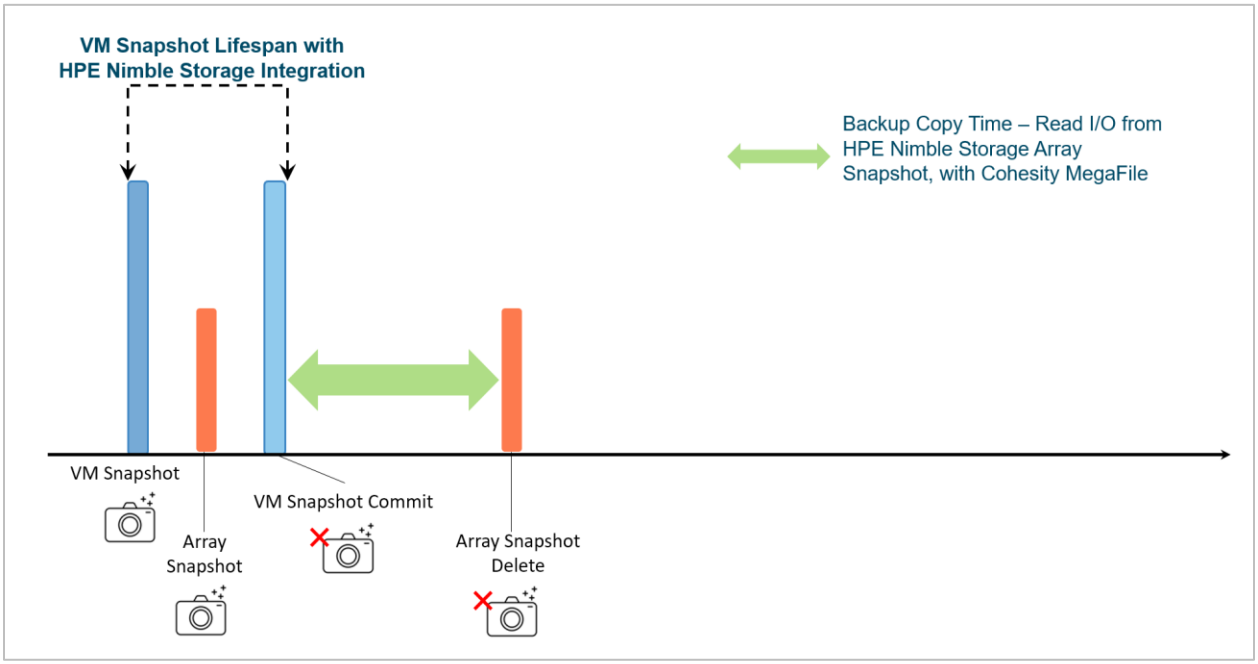

When we use the storage array snapshot instead of the VM snapshot, the VM snapshot lifespan is no longer a function of the size of the data. Therefore, our integration with the storage array snapshot results in much shorter, consistent VM snapshot lifespans, regardless of the size of the VM or the size of the cluster. For example, the VM snapshot lifespan for a VM with 1 TB and another with 10 TB is the same when using the Cohesity-HPE Nimble Storage integration.

With our integration, we have almost completely eliminated VM stun time issues and also reduced the time necessary for backup with the MegaFile feature, by ingesting data in parallel.

Test Results Illustrate Multiplied Reduction of VM-based Snapshot Lifespan

To demonstrate this effect, we tested with 40 VMs of varying sizes on an HPE Nimble Storage datastore and repeated the backup operation with and without storage array snapshot integration. As you can see below, the blue bars representing the VM snapshot lifespan *with* integration are constant, regardless of VM size, unlike the red bars which represent the VM snapshot lifespan *without* integration.

Figure 5: VM Snapshot Lifespan Comparison - With and Without HPE Nimble Storage Array Snapshot Integration

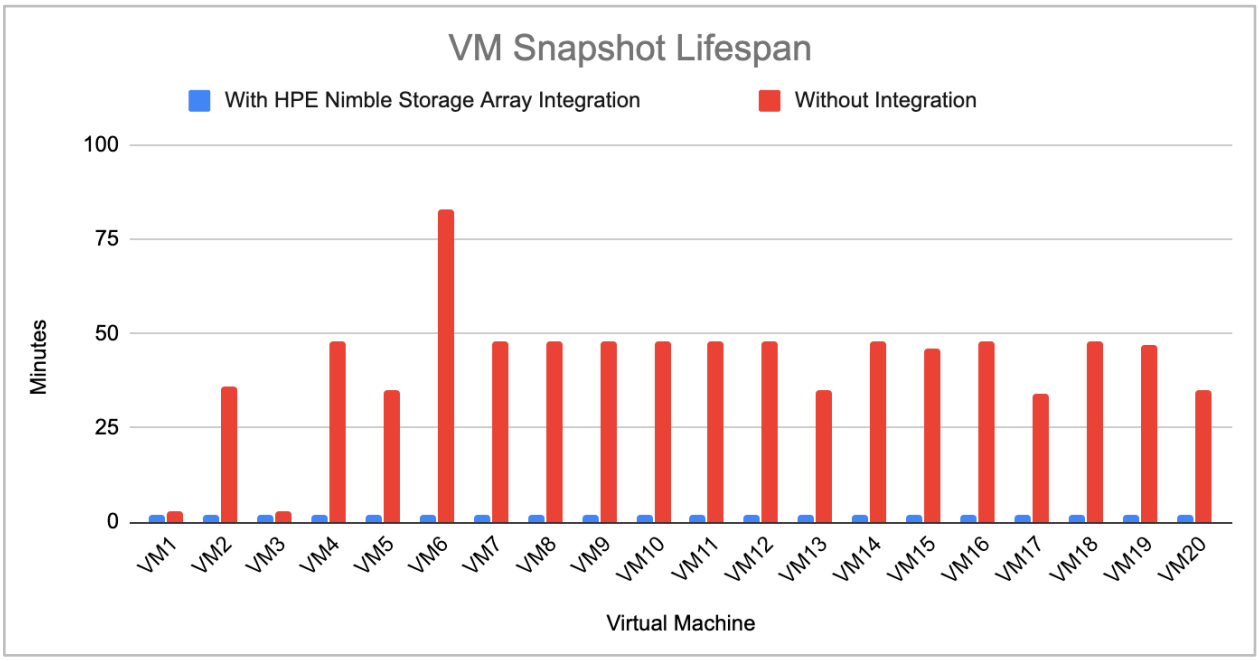

Use Cohesity to Protect VMs on HPE Nimble Storage

Cohesity can protect the virtual machines on HPE Nimble Storage volumes by integrating HPE Nimble Storage snapshots. This integrated approach minimizes performance degradation on the production systems, improves storage efficiency, decreases management complexity, and reduces overall costs.

To start protecting your virtual machines:

1. Register Hypervisor on Cohesity.
2. Register the HPE Nimble Storage array as a Source on Cohesity.
3. Choose a Cohesity Protection Policy.
4. Create a Cohesity Protection Group.

Figure 6: Four Steps to Protect VMs on HPE Nimble Storage

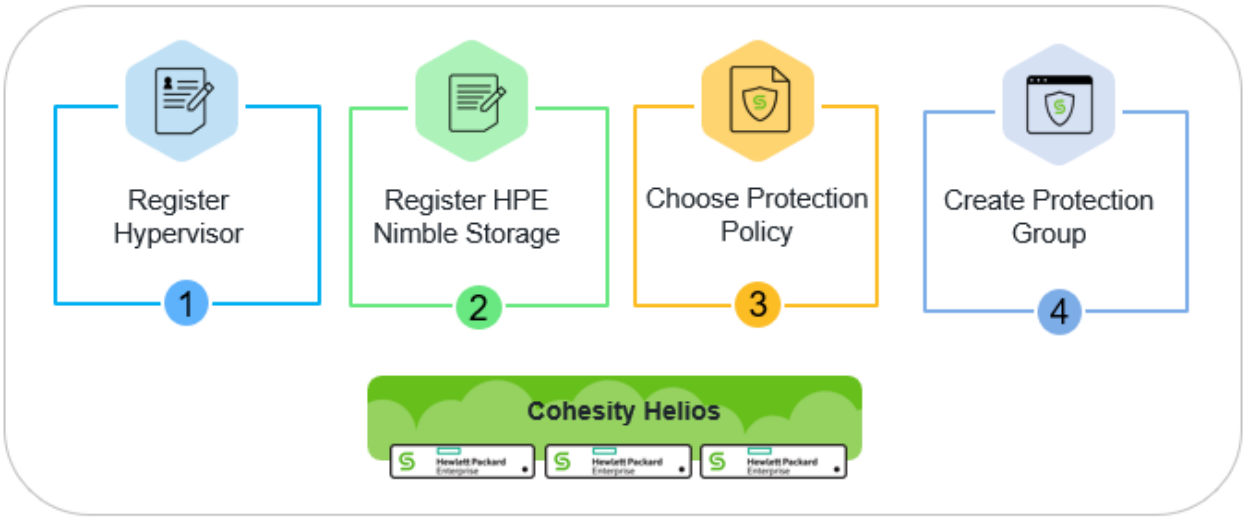

Figure 7 below illustrates the high-level backup process after the integration of Cohesity and HPE Nimble storage.

Figure 7: Cohesity's VM Backup Process with HPE Nimble Storage Integration

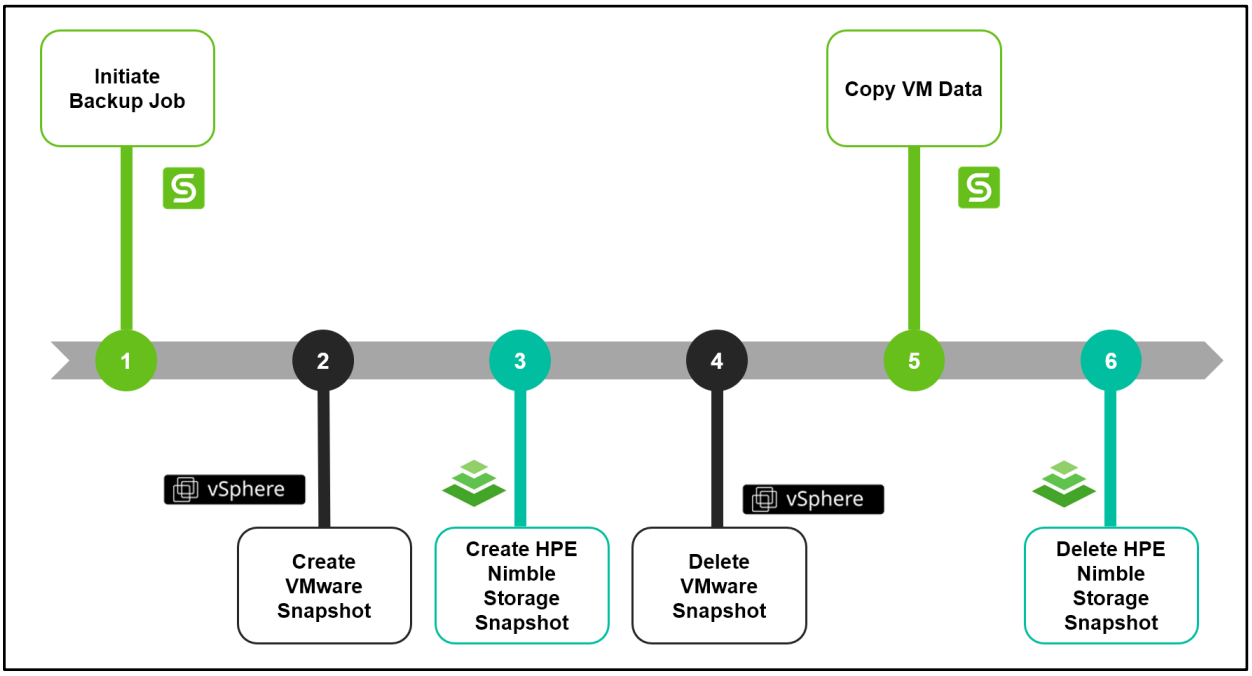

Register Hypervisor on Cohesity

To register VMware vCenter on Cohesity:

1. Log in to Cohesity and navigate to **Data Protection > Sources.**

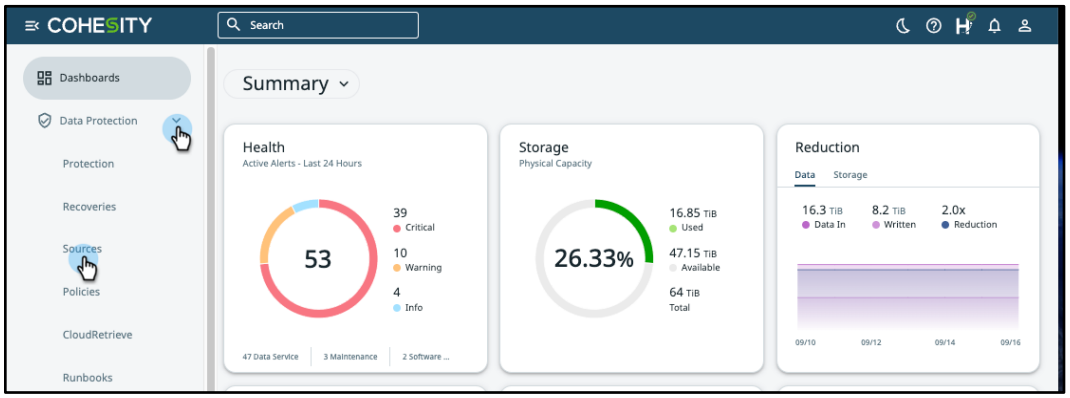

2. On the **Sources** page, click **Register** and select **Virtual Machines.**

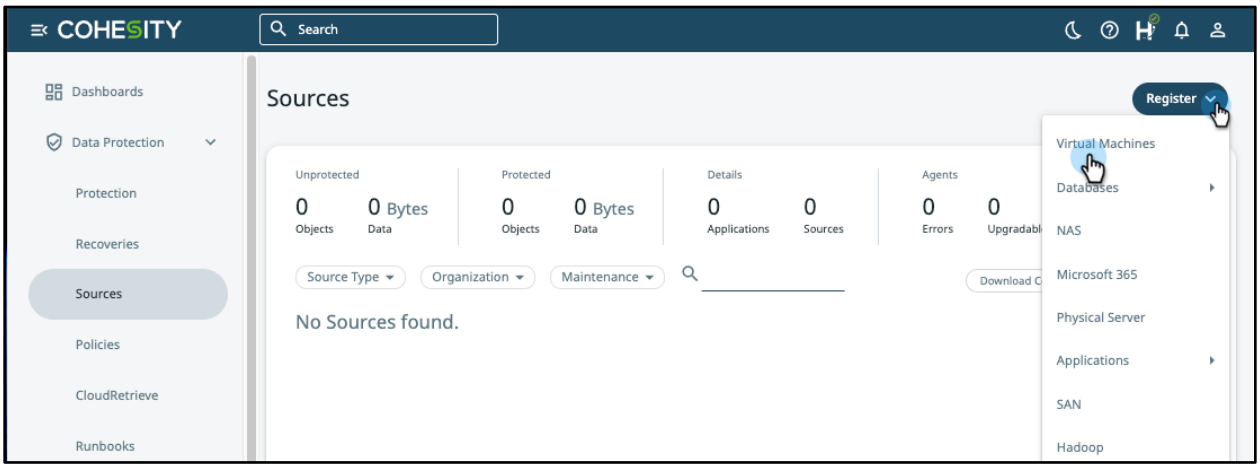

3. On the **Register Virtual Machines** page, enter:
   a. Under **Select Source Type**, select **VMware vCenter**.
   b. **Hostname** or **IP Address**. Enter the management hostname or IP address of your vCenter.
   c. **Username**.
   d. **Password**.
   e. Then click **Register**.

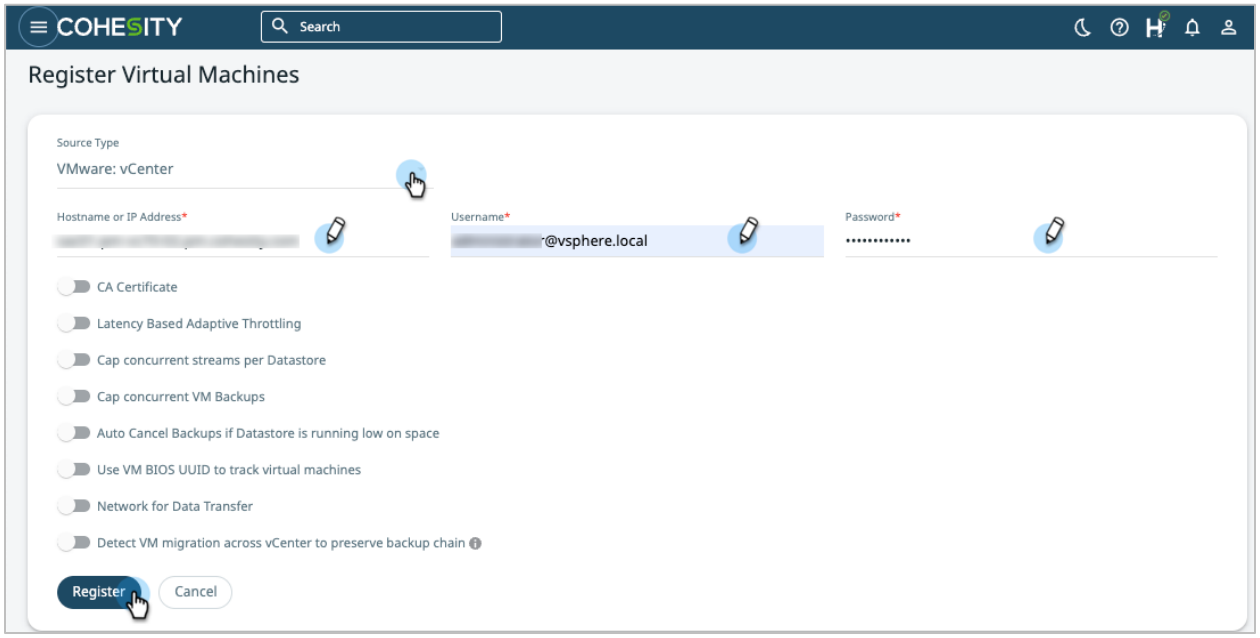

Your VMware vCenter is now registered as a Source with Cohesity and is available to add to a Protection Group after you choose a Protection Policy.

Register HPE Nimble Storage Array on Cohesity

To register HPE Nimble Storage on Cohesity:

1. Log in to Cohesity and navigate to **Data Protection > Sources.**
2. On the **Sources** page, click **Register** and select **SAN.**

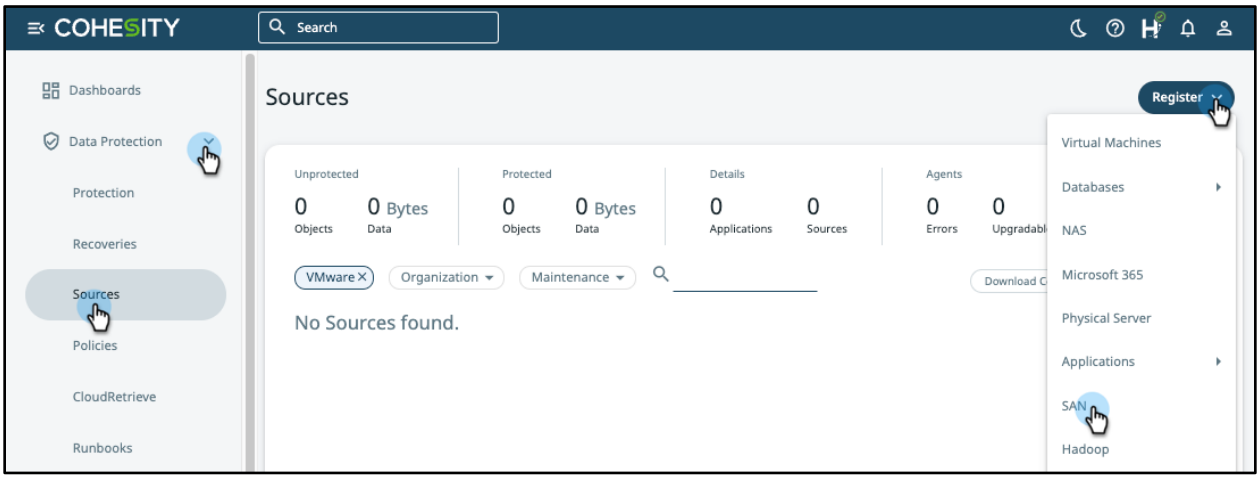

3. On the **Register Storage Array** page, enter:
   a. **Source Type**. Select Nimble.
   b. **Management Host or IP Address**. Enter the management hostname or IP address of your HPE Nimble Storage array.
   c. **Username**.
   d. **Password**.
   e. Then click **Register**.

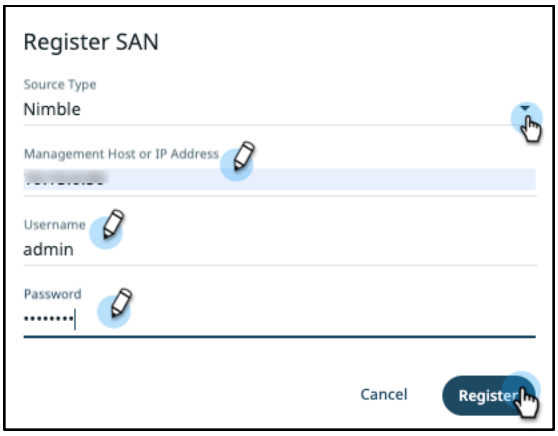

Your HPE Nimble Storage array is now registered as a Source with Cohesity and is available to add to a Protection Group after you choose a Protection Policy.

Choose a Protection Policy

In Cohesity, Protection Groups use Protection Policies. Protection Policies reflect business needs, such as backup and archival frequency and retention requirements, while a Protection Group defines operational requirements, such as the source objects to protect, the Protection Policy to use, the Storage Domain to use, and more. This process of combining business needs (Protection Policy) with the objects to protect and the associated operational requirements (Protection Group) provides rich flexibility to customers.

Cohesity includes three standard policies: **Gold**, **Silver**, and **Bronze**. For their default settings, see Manage Policies in the online Help. If the default settings of the standard policies do not meet your needs, you can create a customized Policy. See Create or Edit a Standard Policy in the online Help.

Create a Protection Group

Protection Groups combine operational requirements — such as the objects to protect, indexing, alerts, exclusions, inclusions, and more. — with the business requirements that are defined in a Protection Policy. Multiple Protection Groups can use the same Protection Policy, but each Protection Group can have only one Policy.

To create a Protection Group:

1. Log in to Cohesity and navigate to **Data Protection > Protection**.

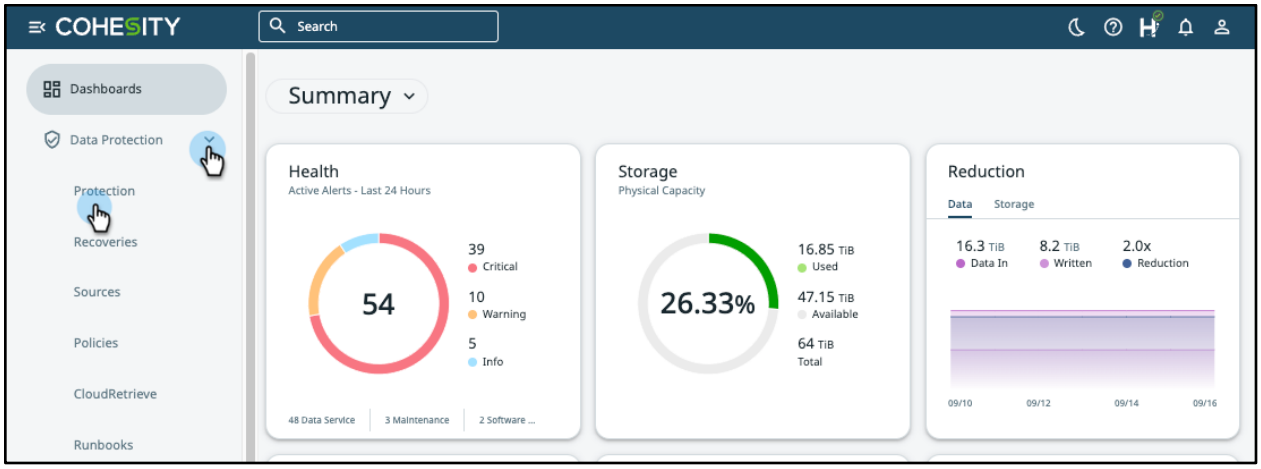

2. Click **Protect** and select **Virtual Machines** to protect the virtual machines that are on the datastore hosted by your HPE Nimble Storage.

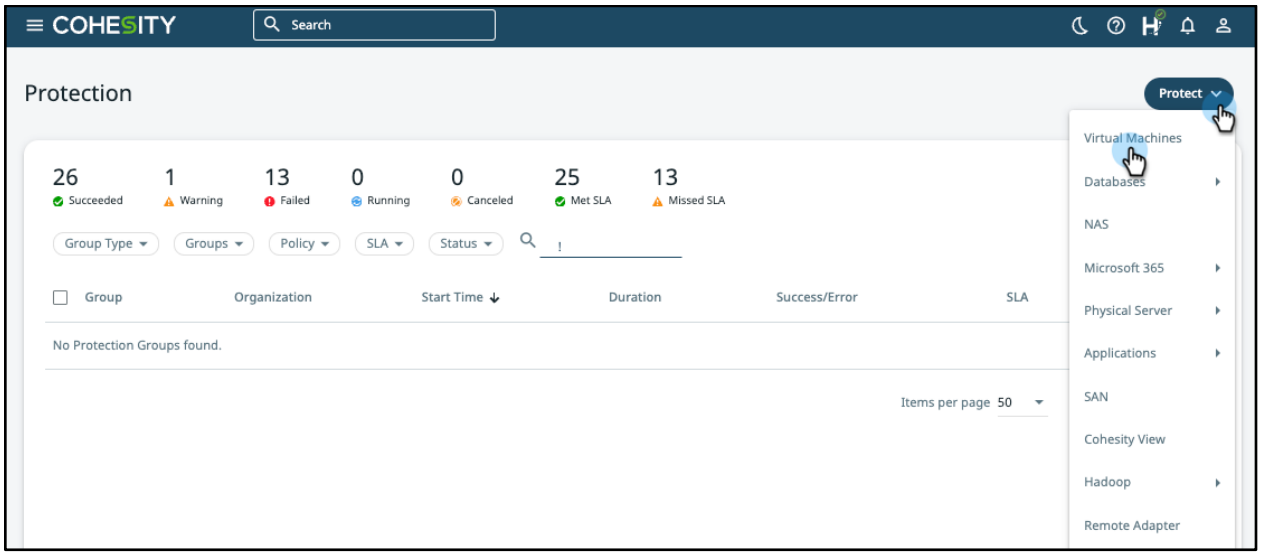

3. Click **Add Objects** and select the virtual server source that you registered earlier. Select the objects to protect and click **Save Selection**.

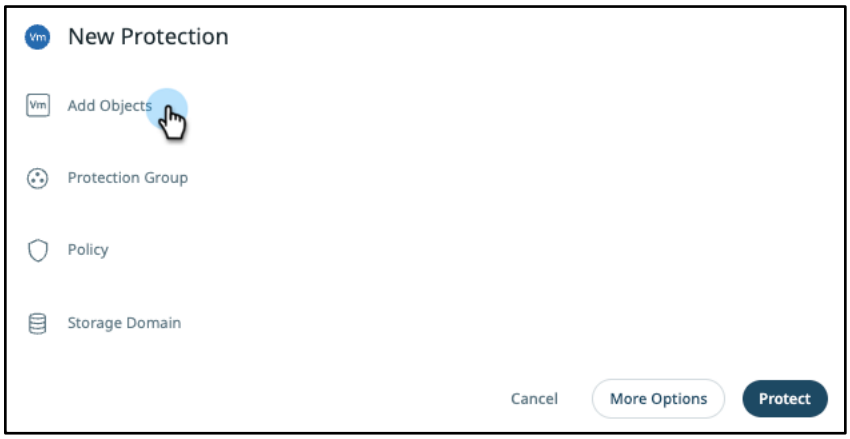

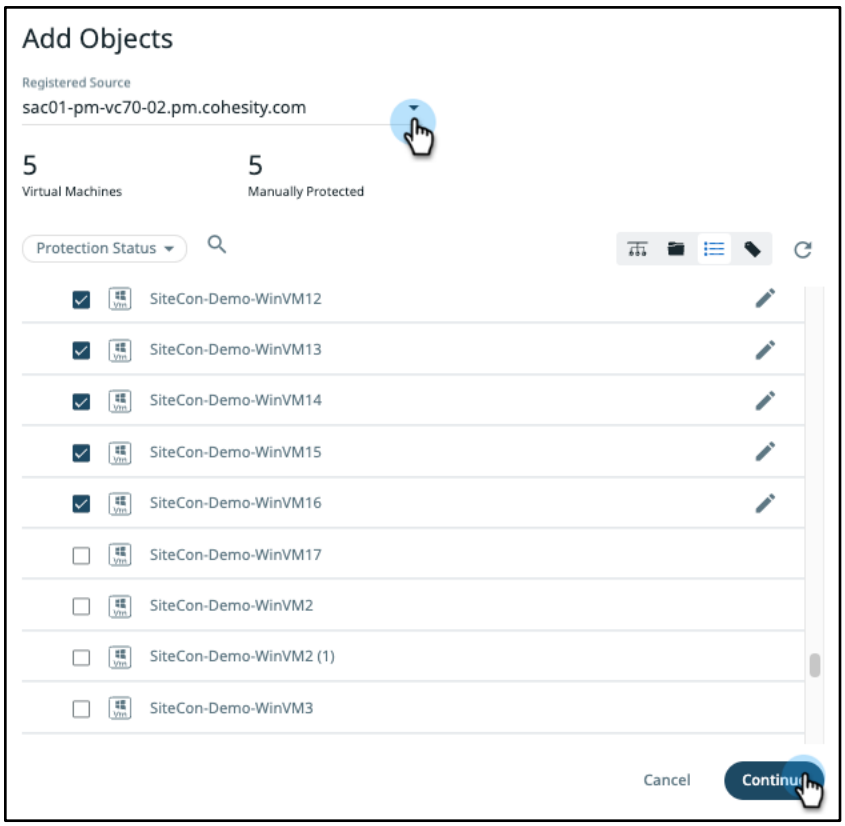
Add Objects
Registered Source
sac01-pm-vc70-02.pm.cohesity.com
5
Virtual Machines
5
Manually Protected
Protection Status
SiteCon-Demo-WinVM12
SiteCon-Demo-WinVM13
SiteCon-Demo-WinVM14
SiteCon-Demo-WinVM15
SiteCon-Demo-WinVM16
SiteCon-Demo-WinVM17
SiteCon-Demo-WinVM2
SiteCon-Demo-WinVM2 (1)
SiteCon-Demo-WinVM3
Cancel
Continu

4. Under **Protection Group**, enter the **Group Name**. Under **Policy**, select an existing policy. (If none of the existing policies meet your needs, scroll to the bottom of the list and select **Create Policy**. See Create or Edit a Standard Policy in the online Help.)

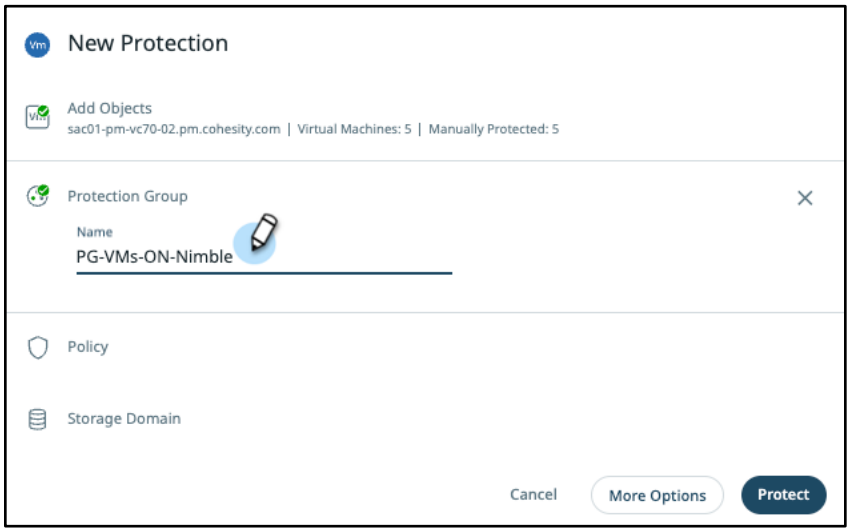

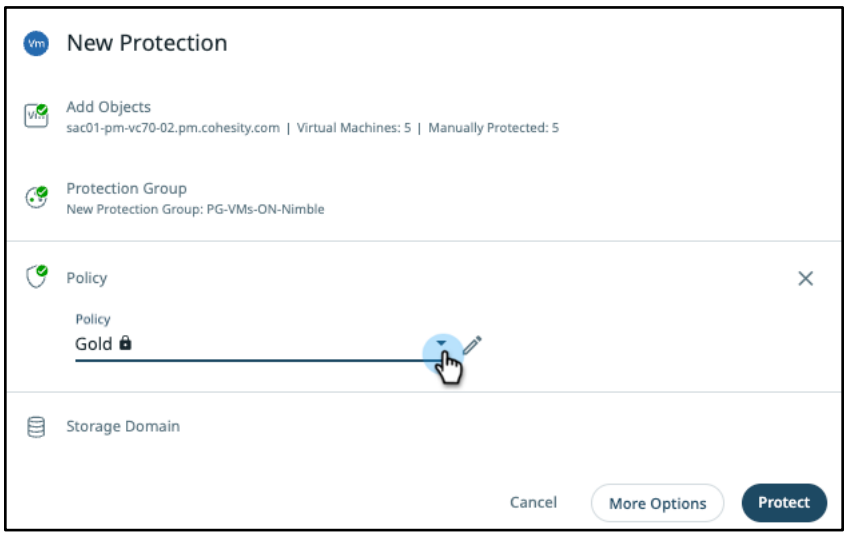

5. Under **Storage Domain**, select a **Storage Domain** and click **More Options.**

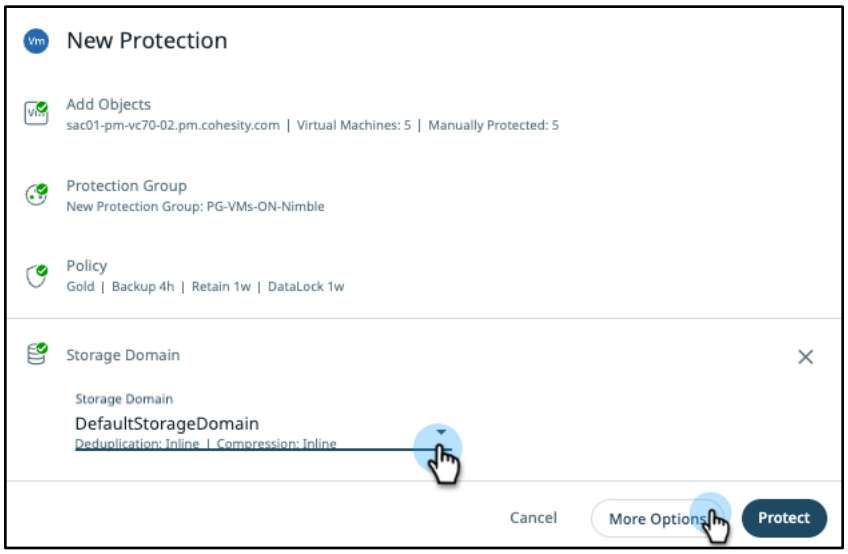

6. Expand **Additional Settings** to enable **Leverage Storage Snapshots for Data Protection** and select **Storage Array/NAS**.

7. Click **Protect**.

   Your new protection group is now active and running.

Use Cohesity to Recover VMs on HPE Nimble Storage

Cohesity's platform with HPE Nimble Storage integration can restore an entire VM almost instantly. In fact, with this solution, you can recover *multiple* VMs instantly. Restoring a VM is as important as backing one up. Equally important questions are:

- Where should the VM be restored?
- When should the restore start?
- How long will it take?
- Should the VM be restored to a specific point in time?

Some of these questions can be answered by the company's Recovery Point Objective (RPO), and their Recovery Time Objective (RTO). Figure 8 explains the recovery procedures of a VM.

Figure 8: Cohesity's VM Recovery Process

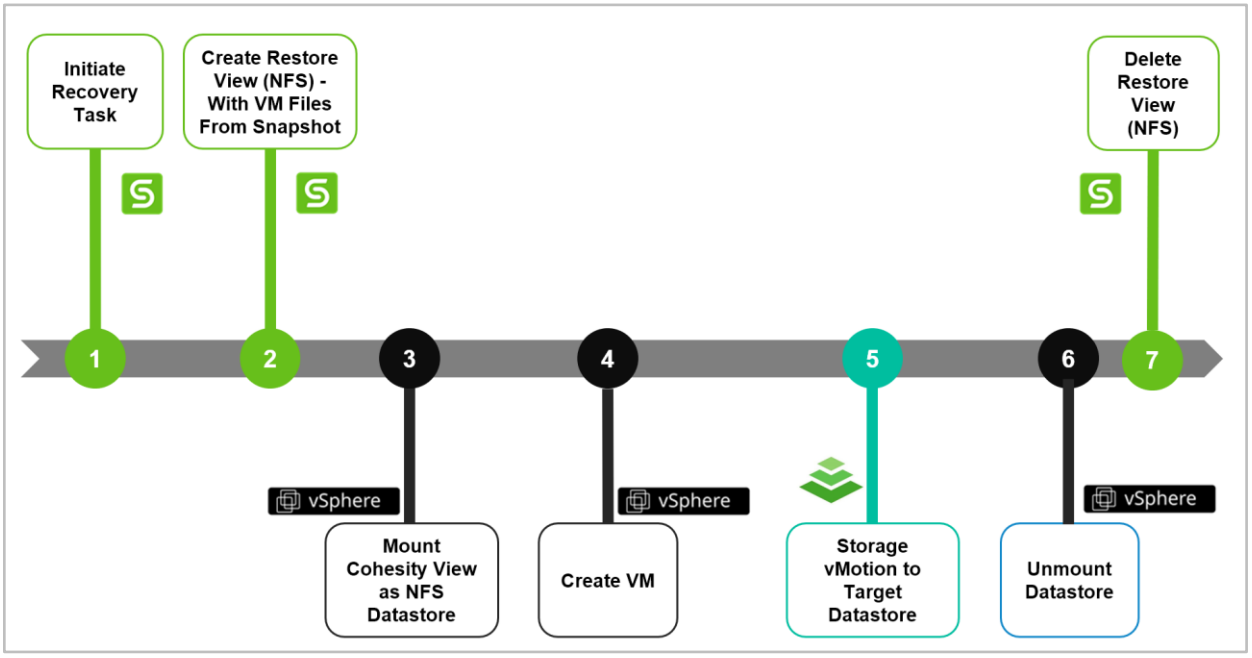

Recover VMs

Cohesity provides the ability to recover Protected Objects (such as VMs) from a snapshot created earlier by a Protection Group.

To recover a VM:

1. Log in to Cohesity and navigate to **Data Protection** > **Recoveries**.

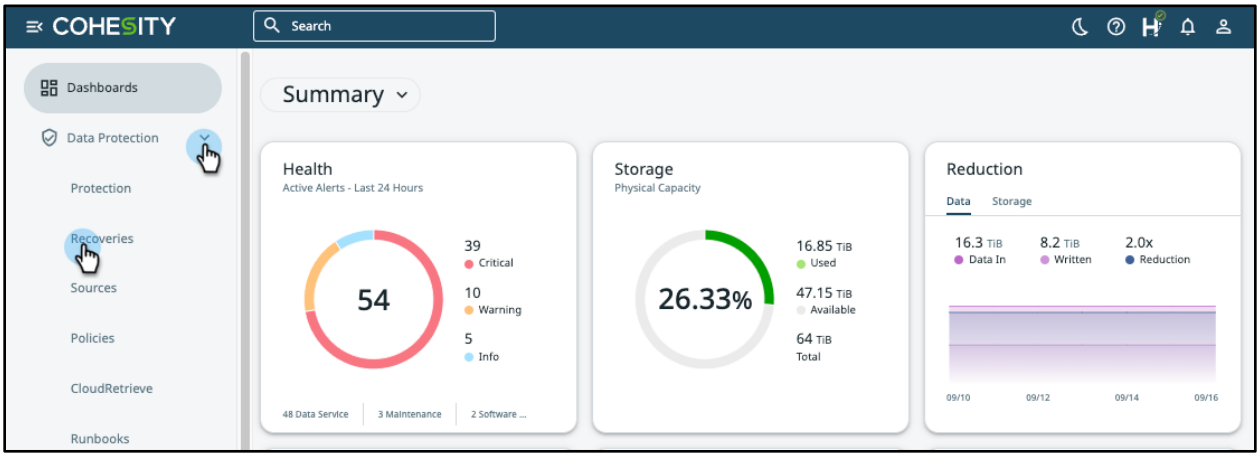

2. On the **Recovery** page, click **Recover** and select **Virtual Machines** and **VMs.**

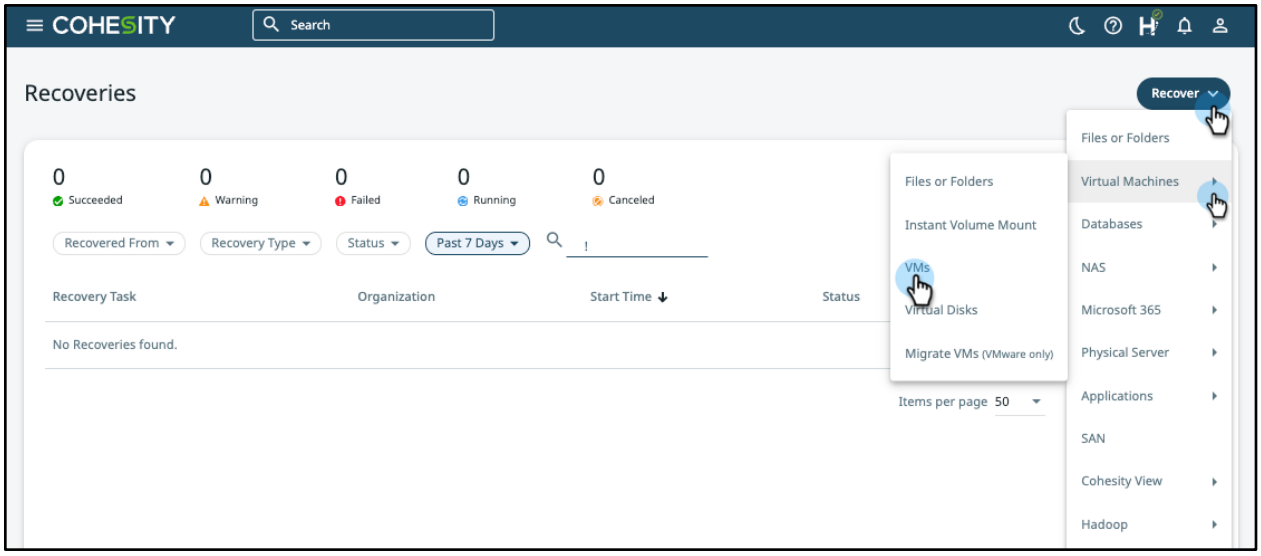

3. Enter the search query for the Objects you want to recover. You can search by **Server** or **Protection Group Name**.

NOTE: Cohesity supports wildcard characters and search filters.

4. Select the **objects** and their Recover Point to recover (By default, the most recent snapshot of the VM is selected.). Click **Recovery Options**.

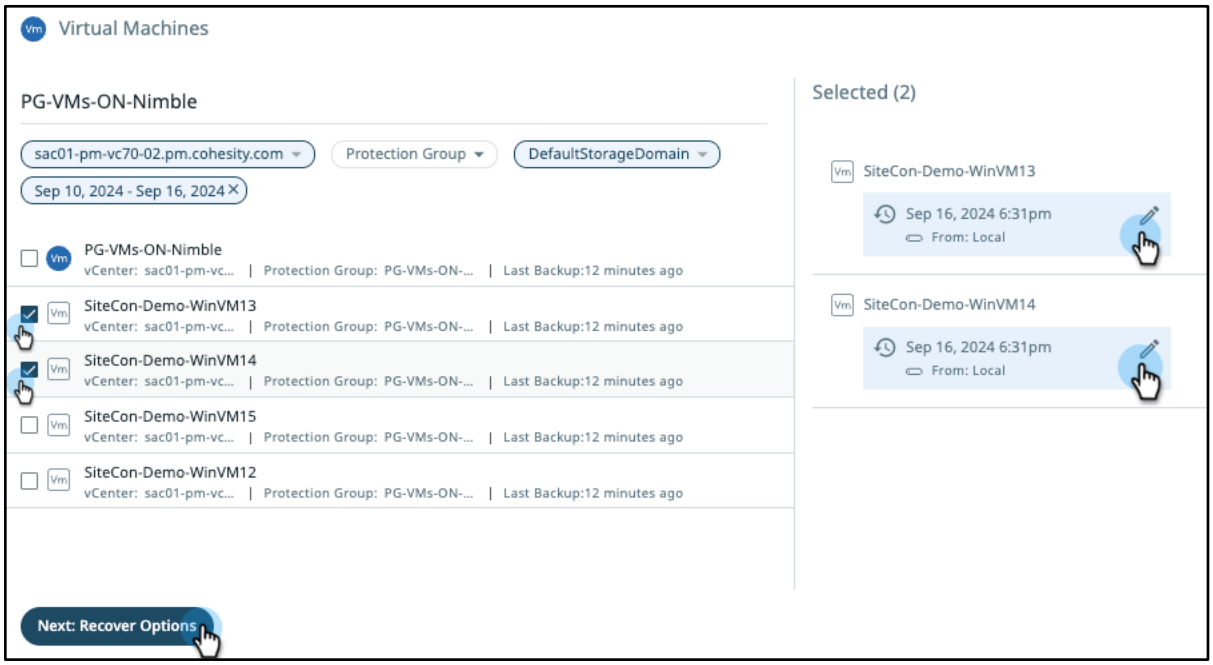

5. Set Recovery options.
   a. Select the **Recovery Location**. (Original Location or New Location)
   b. Select the **Recovery Method** (Instant Recovery or Copy Recovery)
   c. Select the **Existing VM Handling** (Overwrite Existing VM, Keep Existing VM or None)
   d. Enable **Network** (Attach, Start Connected)
   e. Select **Rename** (add a Prefix and/or Suffix to VM name)

f. Enable **Power State** (Power on)

g. Enable **Leverage SAN Transport for Copy Recovery** (*Optional* - Only available with Copy Recovery)

h. Enable **Continue on Error** (*Optional* - Continue recovery even if errors occur when recovering VMs)

i. Select **Cluster Interface** (Auto Select or Custom)

j. Edit **Task Name** (*Optional* - Change the default task name)

k. Finally, click **Recover** to trigger the recovery.

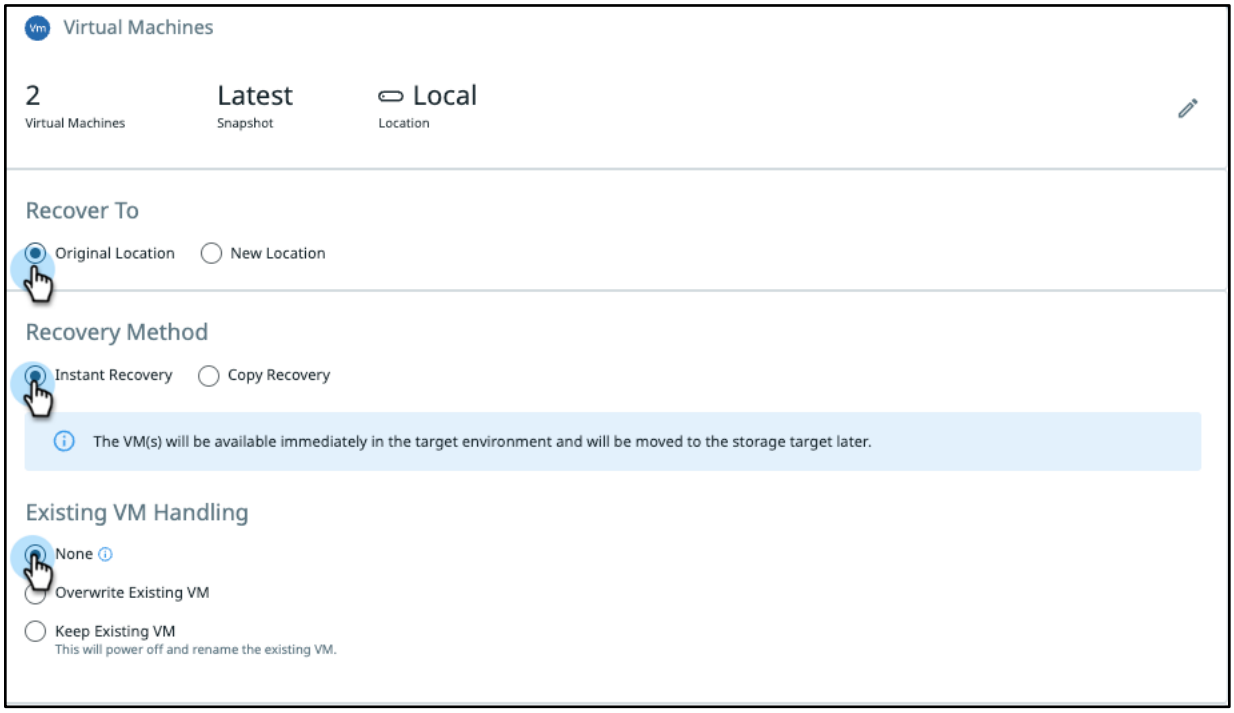

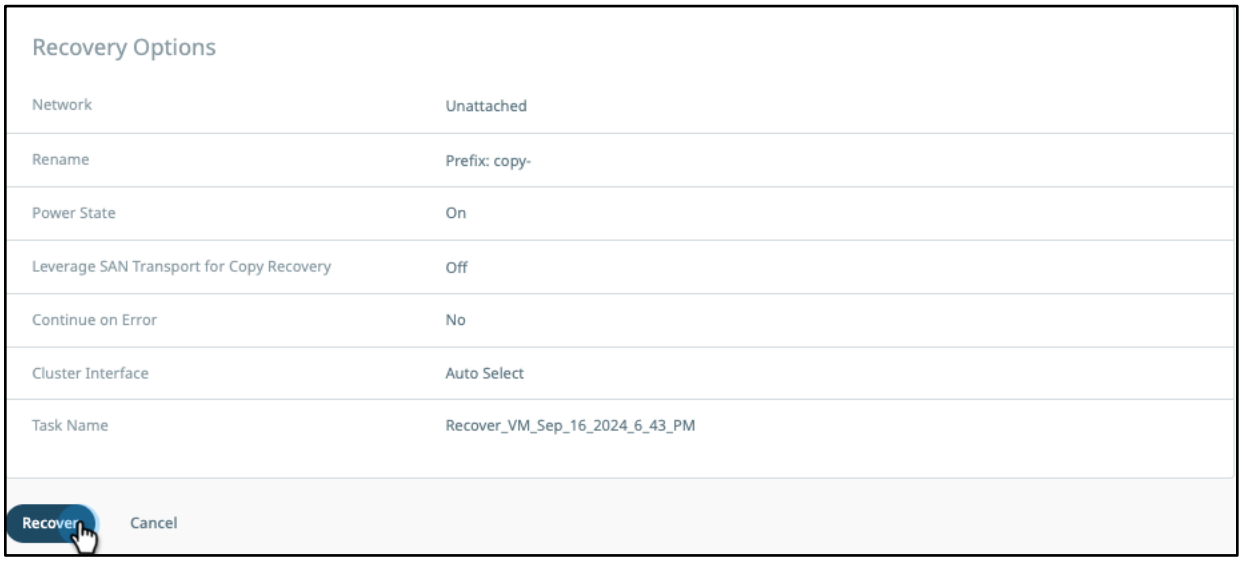

Use CloudArchive for Long-Term Retention

The exponential growth of data volumes and the resulting IT management demands have prompted businesses to seek more cost-effective, reliable data storage and protection solutions. What's more, most organizations require periodic off-site storage of full backup sets (for example, one per month for 12 months, and one per year for seven years). This was historically accomplished with tapes. Now, Cohesity provides a policy-based method to archive to public clouds (AWS, Azure, GCP), to any S3-compatible storage, tape, and/or to any NFS mount point. Cohesity CloudArchive offers a complete, self-contained copy of your backup, containing backup data, backup metadata, indexing data, and deduplication fingerprints.

Backup administrators can take advantage of Cohesity CloudArchive to address long-term data retention requirements. The archived data is efficiently transferred and stored by sending only deduplicated, compressed, incremental backups, thereby reducing network and storage utilization. For added security, you can also enable Archive Object lock to lock archives on external targets and prevent data from being modified, deleted, or overwritten. Refer to Archive Object Lock for more details.

Figure 9: Leverage the Public Cloud for Long-term Retention

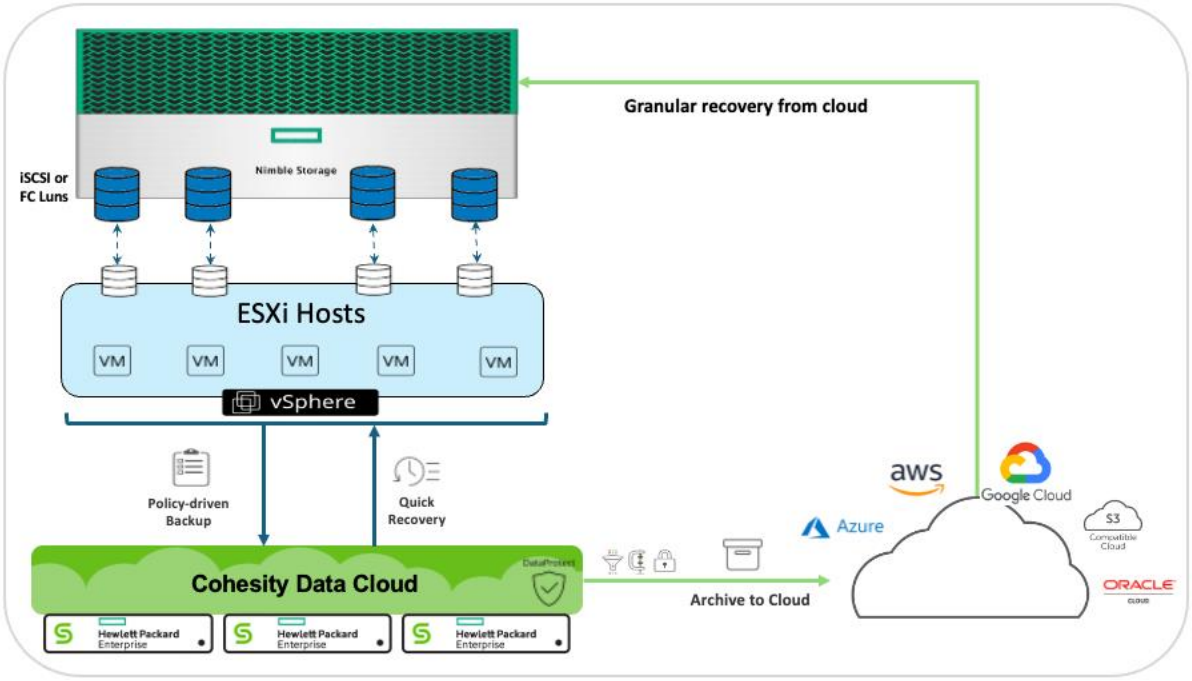

CloudArchive is very flexible. You can use it with AWS, Azure, GCP, NAS, and S3-Compatible cloud object storage.

Maintain Business Continuity with Disaster Recovery

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

- **Replication**. A simple way to store and then get your data back in the event of major business disruptions such as natural disasters and IT failures.
- **CloudArchive**. Archives your data in the cloud for later use with Cloud Recover (to the original cluster) and CloudRetrieve (to a new cluster).

Replicate Backups to Other Cohesity Clusters

A Cohesity cluster can replicate your backup snapshots to another Cohesity cluster in the same data center, in a different data center, or in the public cloud.

Figure 10: Replicate Backups to Other Cohesity Clusters

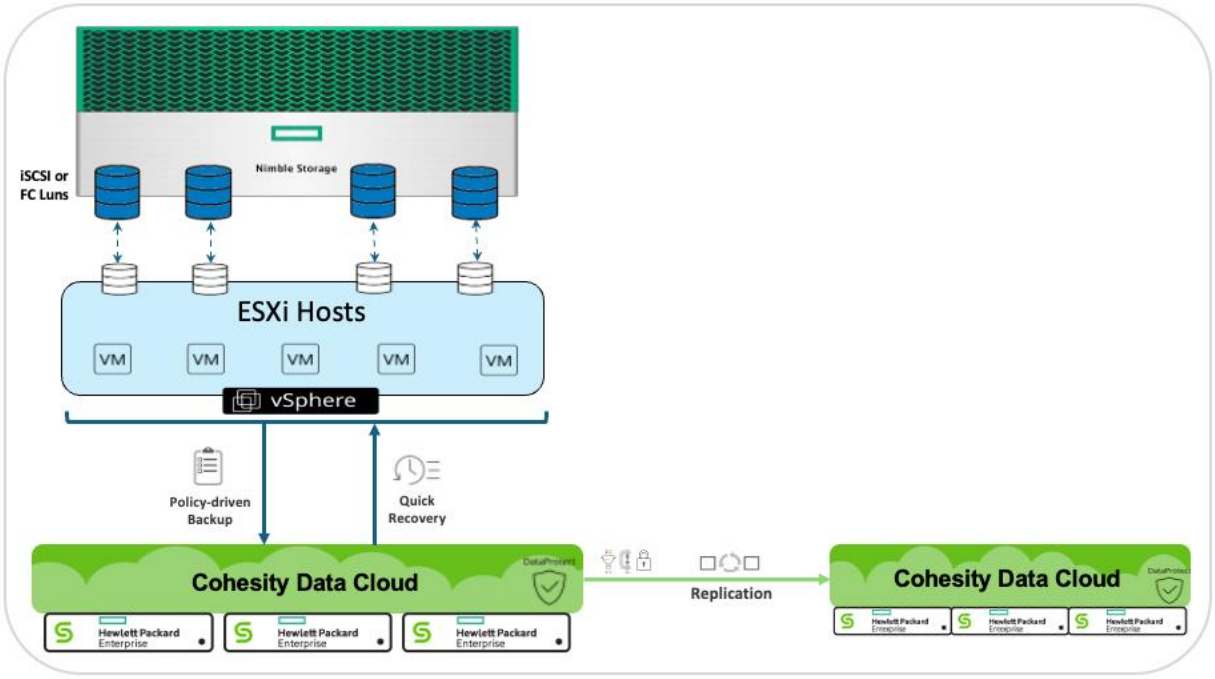

To enable replication, within a Protection Policy, you can configure the settings for replicating protected objects to remote clusters. You can schedule recurring replication, which further improves your data resiliency. You can select a replication target and replication frequency and define the retention time. You can also add new remote replication targets and sources, allowing you to complete entire workflows without leaving the page you are on.

To apply replication to your Protection Policy:

1. Navigate to **Data Protection** > **Policies**.
2. Click **Create Policy** button to create a new Policy, or edit an existing Policy.
3. Click the **Add Replication** icon to open your replication options.

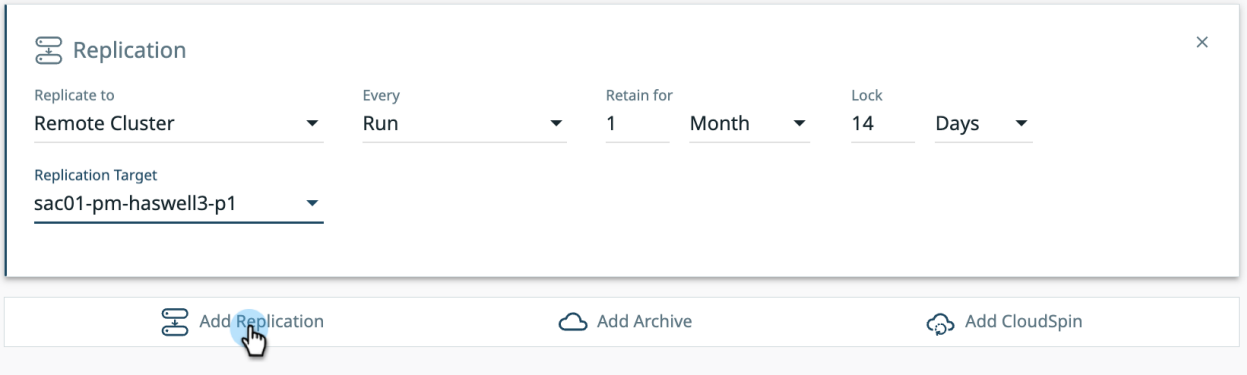

For more, see Replication in the online Help.

Access Your Cloud-stored Data

Once the data is archived using Cohesity's CloudArchive feature, backup administrators can also take advantage of the Cloud Recover and CloudRetrieve features:

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

Figure 11: Cloud Recover to Original Cluster & CloudRetrieve to New Cluster

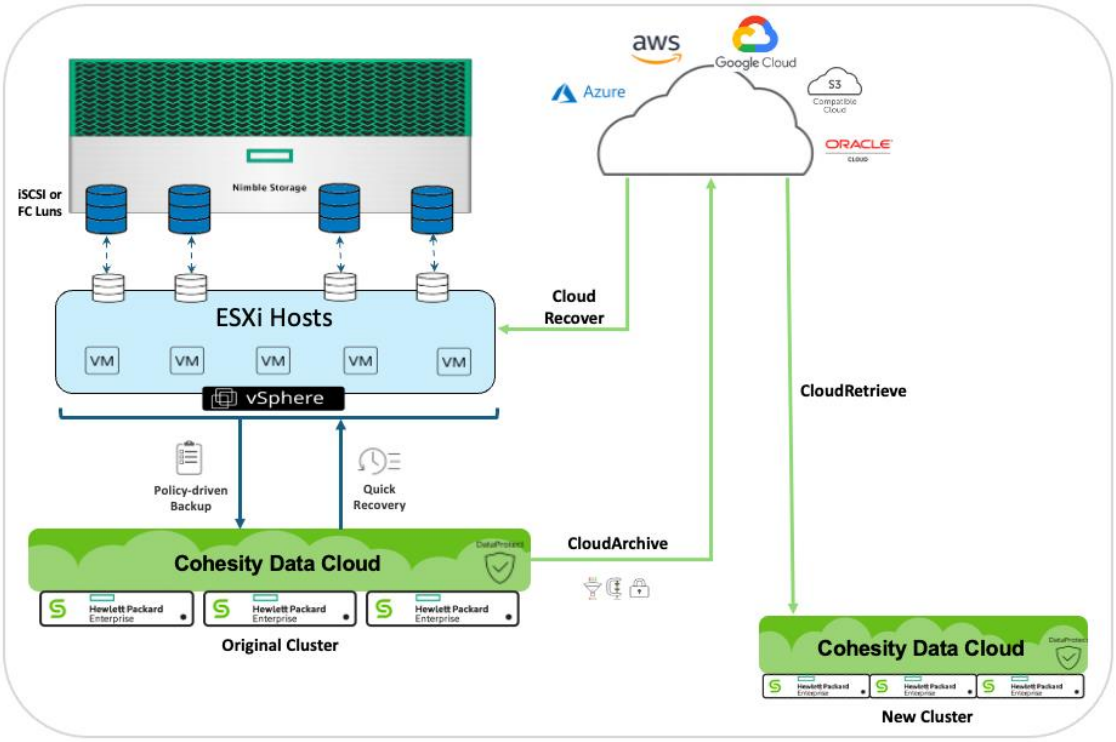

To learn more about CloudArchive, see the *CloudArchive & CloudRetrieve Deployment & Recovery Guide* for AWS, Azure, GCP, NAS, and S3-Compatible cloud object storage.

Reduce TCO with Policy-based Cloud Tier

The performance, availability, and cost requirements of storing and accessing your data can change based on your business needs. Cohesity Cloud Tier allows you to move data to lower-cost storage for infrequently accessed data, reducing operating expenses and helping you meet compliance and access frequency requirements. Cohesity can automatically move data between different tiers.

Data can be down-tiered from your Cohesity cluster to External Targets such as public cloud infrastructure providers (AWS, Azure, or Google Cloud Platform) or any S3-compatible External Target, using a threshold that you define in the Protection Policy. Similarly, hot data in External Targets can be up-tiered back to the Cohesity cluster.

Figure 12: Cohesity Cloud Tier Manages Data Tiering with a Policy-defined Threshold

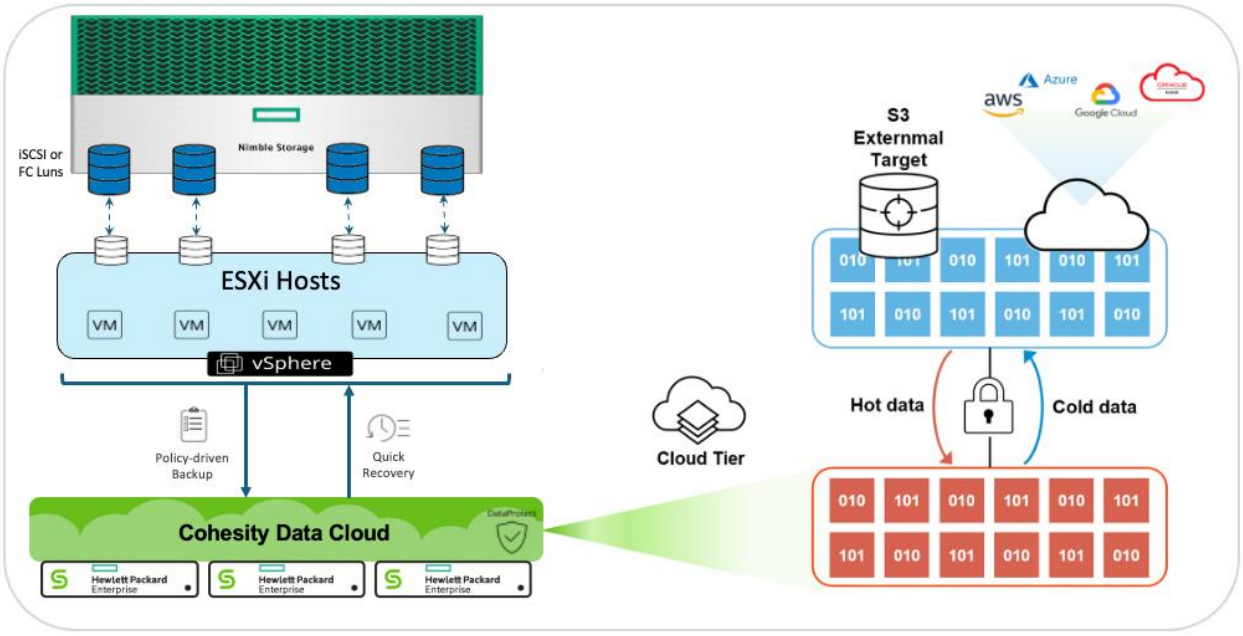

The tiering of cold data to the External Target is based on a policy with two factors: tiering threshold and data policy (i.e., the age of the data, or length of time since it was last accessed). Tiering happens only if utilization exceeds the tiering threshold and some data blocks meet the data policy.

- **Tiering threshold**: The percentage of space utilization set to trigger the tiering of cold data.
- **Data policy**: Specifies the duration of time that the data must be inactive for it to be eligible for tiering.

The tiering threshold and data policy can be set on the Cohesity cluster, on an individual Storage Domain, or both. When set on the cluster, all Storage Domains inherit the setting. When set on an individual Storage Domain, it is not applied to any other Storage Domains.

When the threshold is set on both the cluster and a Storage Domain, the Storage Domain settings take precedence. This allows you to set a global threshold for your cluster and a different threshold for specific Storage Domains, giving you greater flexibility.

When these configured thresholds are breached, data is tiered to the cloud. When tiered data becomes hot data, data is seamlessly tiered from the cloud back to the Cohesity cluster without user intervention.

Following the paradigm upheld throughout by Cohesity, all tiered data is compressed, deduplicated, and encrypted.

For more, see the *Cohesity Cloud Tier Architecture Reference*.

Your Feedback

Was this document helpful? Send us your feedback!

About the Authors

Sadik Sayed is a Technical Solution Engineer at Cohesity. In his role, he focused on VMware data protection, NAS data protection, and SiteContinuity.

Other essential contributors included:

- Karthick Radhakrishnan, Director – Technical Solution Engineering.
- Adaikkappan Arumugam, Sr Director – Product Solution.

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.3	Oct 2024	Minor Updates
1.2	July 2024	Republishing
1.1	Aug 2021	Cohesity rebranding updates
1.0	Oct 2020	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 42 of the Fortune 100.

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