COHESITY | Hewlett Packard Enterprise

Technical White Paper

Integrated Data Protection with Cohesity and HPE Servers

Table of Contents

Executive Summary

Modern enterprises have become boundless data generating machines. They need systems that can store, process, and archive data for various business purposes including regulations and business continuity. These systems may impact business outcomes by making data available for development, analyzing it for insights, or providing a service to the end customer. However, current architectures with siloed approach across storage tiers are creating gross inefficiencies and putting severe strain in business operations.

This white paper describes the storage architecture that Hewlett Packard Enterprise (HPE) and Cohesity have created to eliminate silos and put all of an enterprise's data to use. Whether data resides on primary or secondary storage HPE and Cohesity can put all data to good use, which maximizes the return on any storage investment. The architecture is based on HPE's industry leading HPE 3PAR all-flash arrays and Cohesity's Hyperconverged Secondary Storage platform running on qualified HPE servers. The paper talks about how the infinitely scalable nature of this solution can address the needs of even the largest and most complex data-driven enterprises.

Introduction

Business agility and operational speeds are not only products of innovation, they may also be required for the survival of most modern companies. The key to achieving these requirements is to create a simplified and easy-to-manage IT infrastructure. This is commonly attempted through the use of a myriad of tools and "bleeding-edge" architectures. However, this goal has been greatly stymied by a lack of innovation in secondary storage, where more than 80% of an enterprise's data resides. This is especially galling considering how data and hence storage-dependent enterprises are fast becoming.

The Solutions Architects and IT experts have largely focused on implementing "modern" primary storage architectures. They are employing agile methods of development to extract more value for their businesses. This has led to a mad rush of storage companies "innovating" faster speeds and feeds for primary storage systems using the latest hardware and flash technologies. While more and more data is being generated on primary storage environments, methods of data protection and data management have not seen any major innovation in the last decade to handle this magnitude of data.
This lack of focus is evident when looking at the data protection systems available in the industry today. Data protection systems are complex, cumbersome, and treated only as an insurance policy, which provides minimal return on investment (ROI). Being the bellwether in the technology industry, HPE recognized this gap and felt compelled to bring a next generation solution to its customers. HPE's customers have trusted and always looked at HPE to innovate and showcase best of breed technologies. HPE's partnership with Cohesity intends to bridge this gap by unifying primary and secondary storage, all flash and hybrid tiers, as well as traditional and hyperconverged architectures into a single data fabric.
Cohesity has created a hyperconverged secondary storage platform that offers a radically efficient solution to fragmented data siloes by unifying all secondary use cases on a single, infinitely scalable platform. These use cases span tiers 2 through 4 of the generally acknowledged storage tiers that include Backup, Disaster Recovery, Test/Dev, File Shares, Archiving, and Analytics.

Cohesity is a true Software Defined Storage System (SDSS) that enables policy-based provisioning and management of data storage independent of the underlying hardware. This approach aims to transform physical storage hardware without disrupting application access and to provide users a choice in the hardware used to build data center storage services at lower costs. Cohesity's overarching goal is to consolidate secondary storage into a single scale-out platform to radically simplify data protection and put secondary data to work. This is a big paradigm shift in how data protection is traditionally thought of as just an expensive insurance policy. This new way of protecting data allows the same platform to manage data using copy data management. Now, data protection systems may extend to other secondary storage use cases such as test and development, analytics, and file services. This allows enterprises to do more with less and simplify their storage landscape by eliminating solutions which only address point problems and add management overhead on an already stretched IT organization.

DevOps allows Engineering to innovate faster and enables it to divide responsibilities with Corporate IT. DevOps has become an important workflow for agile development. Cohesity enable organizations to protect their data and allows Engineering to leverage virtual clones of the production dataset to upgrade, develop and build newer and better solutions using the patented Snaptree technology.

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 secondary storage software.

The goals of this white paper are twofold:

1. Showcase the HPE hardware + Cohesity software-defined secondary storage integration. The Cohesity SDSS solution has been qualified on select HPE servers. This integration enables HPE to adopt Cohesity as the next generation data protection and data management platform and thereby bring the benefits of simplicity in managing the data sprawl in secondary storage.
2. Show how enterprise customers can face challenges of exponential data growth in their environments with a solution, which reduces CAPEX and OPEX spending, while delivering simplicity, speed, and ease of deployment. We intend to highlight the overall value realized by eliminating the need to manage and maintain multiple copies of data by consolidating secondary storage workflows onto a single platform.

Audience

The target audience for this document includes backup, storage, and virtualization administrators, consulting data center architects, field engineers, and desktop specialists who want to implement a next generation data center. A working knowledge of VMware vSphere®, SQL Server, HPE 3PAR StoreServ Storage, backups, networking, and data center design is assumed, but not a prerequisite to read and understand this document.

Solution overview

System Requirements

For detailed system requirements and setup, please see HPE SPOCK and HPE 3PAR Operating System Implementation Guides. You can find all of the documents in the link h20272.www2.hpe.com/spock/.

For detailed Cohesity system requirements please visit the link support.cohesity.com.

Topology

The topology that we are using for the primary configuration is FC (Fibre Channel). For more details on how to setup the topology, please see the HPE 3PAR Operating System Implementation Guides at h17007.www1.hpe.com/us/en/storage/info-library/index.aspx.

Cohesity

Cohesity is Software Defined Storage that has the mission to simplify secondary storage by eliminating fragmented silos in enterprise data centers. Cohesity collapses secondary storage workflows like backup, DevOps, and analytics onto a unified, intelligent data platform. The Cohesity DataPlatform enables organizations to eliminate redundant copies of data while also giving a way to tap into the untouched “dark data” that previously existed within the data center.

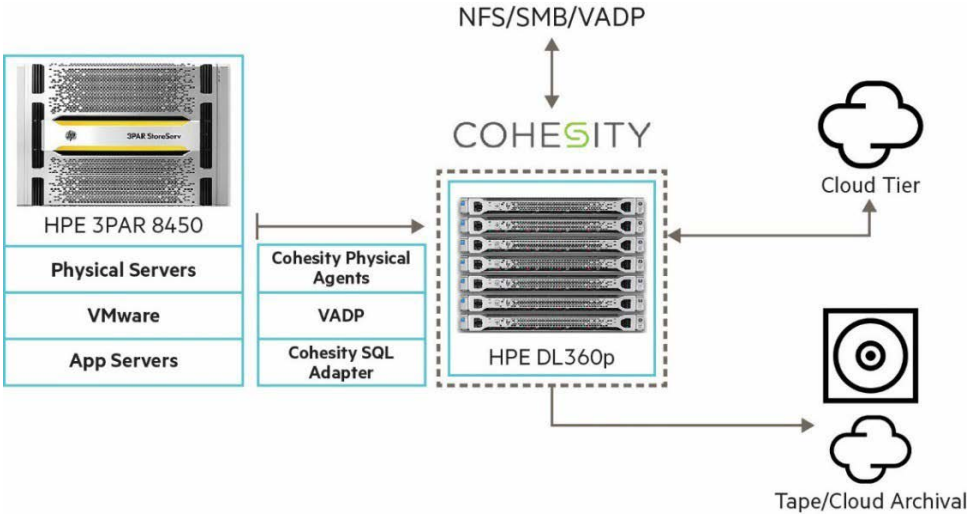

Figure 1. Overview of the Cohesity Software Defined Secondary Storage Solution

Cohesity DataPlatform, integrated with HPE 3PAR StoreServ provides the simplicity and value that comes from secondary storage while providing enterprise-class flash-based storage for performance and reliability. The HPE 3PAR StoreServ is designed to deliver performance and efficiency to I/O-intensive infrastructure workloads. It is the platform of choice for enterprise customers who run databases, virtualization, and bare-metal applications requiring the highest performance.

Cohesity complements HPE 3PAR StoreServ with a distributed platform architecture for data protection and data management designed for high availability. The Cohesity nodes deploy a shared-nothing topology while having no single point of failure or inherent bottlenecks.
Consequently, both performance and capacity can scale linearly as more physical nodes are added to the cluster. The distributed file system spans across all nodes in the cluster and natively provides global deduplication, compression, and encryption.
In addition to the Copy Data Management that is built into Cohesity's Data Platform, Cohesity allows for zero-space instant writeable clones from the backups. Therefore, the same Data Platform can be used to provision and orchestrate in test and development environments. The entire platform is exposed through a single pane of glass that manages all secondary storage workflows that span multiple clusters, geographical sites, and clouds. The platform is fully programmable using REST APIs, which can be leveraged by developers for both DevOps and test and development workflows.
By consolidating all the tier 2–4 storage workflows into a single unified scale-out, web-scale platform, Cohesity reduces overall TCO and improves business agility for enterprise customers. The solution eliminates fragmentation and significantly improves storage efficiency.

Qualified HPE Servers for Cohesity Software

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 secondary storage software. Refer to the Cohesity+HPE solution page for supported HPE servers with Cohesity software. For additional information on specific HPE servers, please refer to HPE server site.

Note: The latest generation of qualified HPE servers for Cohesity software should always be considered when designing a new solution.

HPE 3PAR StoreServ

IT has never been more important to doing business, which means that IT infrastructure must be simpler, smarter, faster, more flexible, and more business-aligned than ever. In the idea economy, business success is defined by how quickly you can turn ideas into value. Is your infrastructure ready?

When it comes to tier-1 storage, HPE 3PAR StoreServ Storage has you covered. The foundation of the HPE Storage portfolio, HPE 3PAR StoreServ Storage offers a range of models that give you effortless, tier-1 flash with midrange affordability and help you affordably consolidate all of your applications onto enterprise flash.

HPE 3PAR StoreServ Storage allows you to break down the silos that stand between you and the efficiency and agility required to succeed in the idea economy. It's the last primary storage architecture you need—regardless of whether you are a midsize enterprise experiencing rapid growth in your virtualized environment, a large enterprise looking to support IT as a Service (ITaaS), or a global service provider building a hybrid or managed private cloud.

HPE 3PAR StoreServ was built to meet the extreme requirements of massively consolidated cloud service providers. Its remarkable speed—3M+ IOPS—and proven system architecture has been extended to transform mainstream midrange and enterprise deployments, with solutions from a few terabytes up to more than 20 PB scale.

The HPE 3PAR StoreServ family of flash-optimized data storage systems modernizes your data center to instantly handle unpredictable workloads with 99.9999% data availability. Get rapid and automated provisioning, multi-tenant design, hardware-accelerated deduplication and compression, and sub-1 ms latency—all within a single tier-1 storage architecture that starts small and scales big.

HPE 3PAR StoreServ is the industry's #1 midrange Fibre Channel storage array and a leading all-flash array. It has helped thousands of customers eliminate complex storage silos and modernize their storage infrastructure.

From midrange to enterprise to all-flash across unified file and block workloads, you can start with a few terabytes in a single system and scale to over eighty petabytes in a four-system federation with a common OS, feature set, and management.

Configuration overview

While Cohesity is not currently integrated directly with the HPE 3PAR array, it can still protect, back up, and restore data stored on it. We accomplish this by protecting the host's volumes that are presented from the HPE 3PAR array.
HPE 3PAR StoreServ Storage primary configuration setup:

The primary configuration consists of three physical hosts (CentOS, VMware® ESXi™, and Microsoft® Hyper-V) along with 12 virtual machines running various operating systems. Each physical host is connected to a 16 GB Fibre Channel switch, which is zoned to the HPE 3PAR StoreServ. Virtual volumes are than presented from the StoreServ array to each physical host.
Cohesity secondary configuration setup:

The secondary configuration consists of a cluster of three HPE servers running Cohesity software. Each server is connected to a 10 Gb Ethernet switch. The Cohesity software is then configured to protect all data written to the primary storage (HPE 3PAR StoreServ).

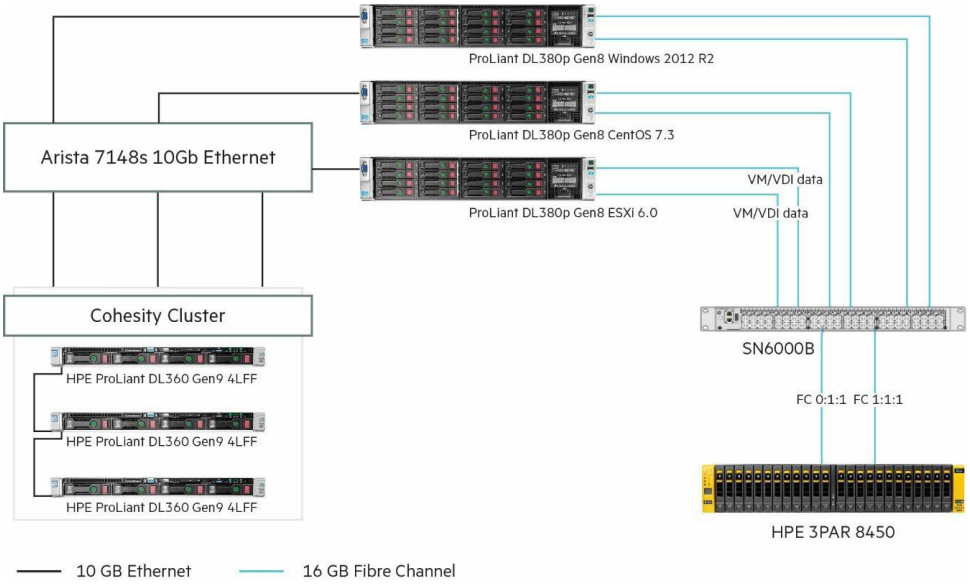

Figure 2. Topology diagram

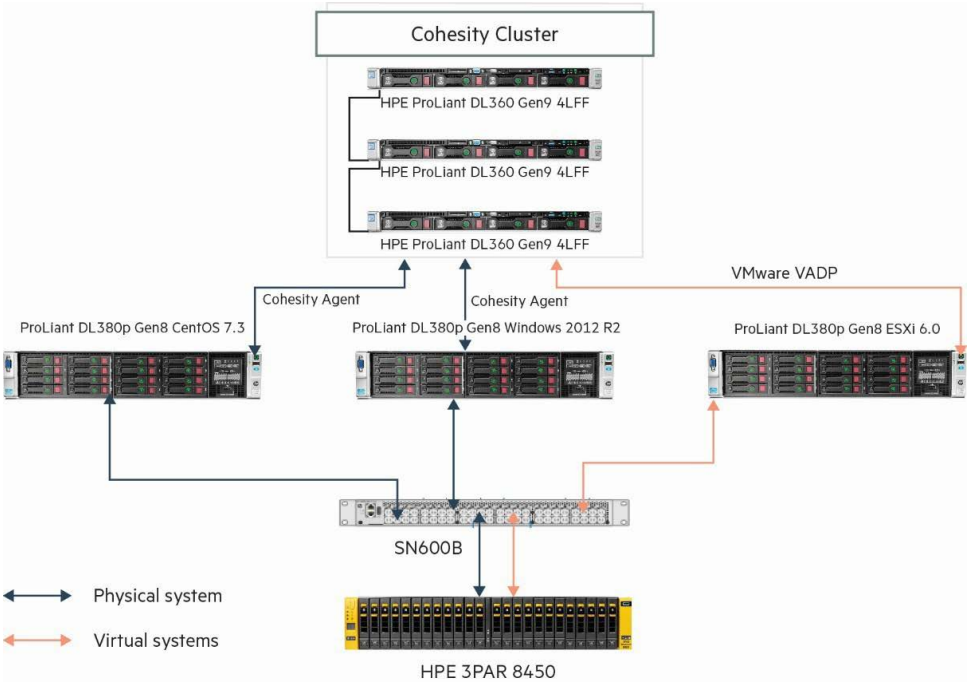

Figure 3. Logical diagram of data flow

Primary storage configuration

HPE StoreServ Storage

Export LUNs to the host(s):

```
# showvlun -host dl380pg8-80 Active VLUNs
Lun VVName     HostName        -Host_WWN/iSCSI_Name-                Port Type Status ID 20 vv_pg8-80.0 dl380pg8-80
51402EC000F0C6E6     0:2:2 host active          1
21 vv_pg8-80.1 dl380pg8-80 51402EC000F0C6E6        0:2:2 host active          1
20 vv_pg8-80.0 dl380pg8-80 51402EC000F0C6E4        1:2:2 host active 1
21 vv_pg8-80.1 dl380pg8-80 51402EC000F0C6E4        1:2:2 host active 1


VLUN TemplatesLun VVName    HostName        -Host_WWN/iSCSI_Name- Port Type 20 vv_pg8-80.0 dl380pg8-80 ----------
------     --- host
21 vv_pg8-80.1 dl380pg8-80 ----------------        --- host
```

Linux® Hosts

Validate that the LUNs are seen on the SCSI and multipath layer (if applicable). LUNs seen in SCSI layer:

```
# lsscsi
[1:0:0:20]       disk      3PARdata VV      3310     /dev/sdb
[1:0:0:21]       disk      3PARdata VV      3310     /dev/sdc
[1:0:0:254]      enclosu 3PARdata SES     3310     -


[4:0:0:20]       disk      3PARdata VV      3310     /dev/sdd
[4:0:0:21]       disk      3PARdata VV      3310     /dev/sde
[4:0:0:254]      enclosu 3PARdata SES     3310     -
```

LUNs seen in multipath layer:

```
# multipath -ll
360002ac0000000000000070007e859 dm-0 3PARdata,VV
size=50G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
`-+- policy='round-robin 0' prio=50 status=active
|- 1:0:0:20       sdb 8:16 active ready running
`- 4:0:0:20       sdd 8:48 active ready running

360002ac0000000000000080007e859 dm-1 3PARdata,VV
size=50G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
`-+- policy='round-robin 0' prio=50 status=active
|- 1:0:0:21       sdc 8:32 active ready running
`- 4:0:0:21       sde 8:64 active ready running
```

Create two logical volumes and filesystems using the HPE 3PAR StoreServ LUNs:

```
# df -h
Filesystem       Size      Used Avail Use% Mounted on
/dev/mapper/cl-root      50G      4.9G      46G      10% /
/dev/mapper/cohesity_test1-lv_cohesity_test1   49G      9.7G      36G      22% /root/mnt_cohesity_test1
/dev/mapper/cohesity_test2-lv_cohesity_test2   49G      9.7G      36G      22% /root/mnt_cohesity_test2
```

Populate data into the filesystem:

```
├┼┼mnt_cohesity_test1
├          ├┼┼1
├          ├┼┼CentOS-7-x86_64-Everything-1611.iso
├          ├┼┼Fedora-Server-dvd-x86_64-25-1.3.iso
├          ├┼┼lost+found
├┼┼mnt_cohesity_test2
├          ├┼┼CentOS-7-x86_64-Everything-1611.iso
├          ├┼┼Fedora-Server-dvd-x86_64-25-1.3.iso
├          ├┼┼lost+found
```

Secondary storage configuration

Cohesity Dashboard

Cohesity Data Platform's policy-based storage management approach makes the system health, capacity, and performance immediately viewable to the user. Upon supplying credentials to the unified management web console, users are presented with an overall health dashboard.

This dashboard reflects the overall health and state of the cluster including the number of jobs that have run, any SLA violations, errors or alerts, as well as a brief data reduction and performance summary. Each of these items can be chosen to show further details. Storage, backup, and virtualization administrators may utilize the dashboard for cluster status reviews, however the platform also provides built-in alerting mechanisms as well as SNMP support for notification of certain conditions.

Complementing the dashboard is the native REST API, where all cluster management tasks can also be driven. Full monitoring and administration capabilities are available via the REST API and associated documentation that resides on the cluster.

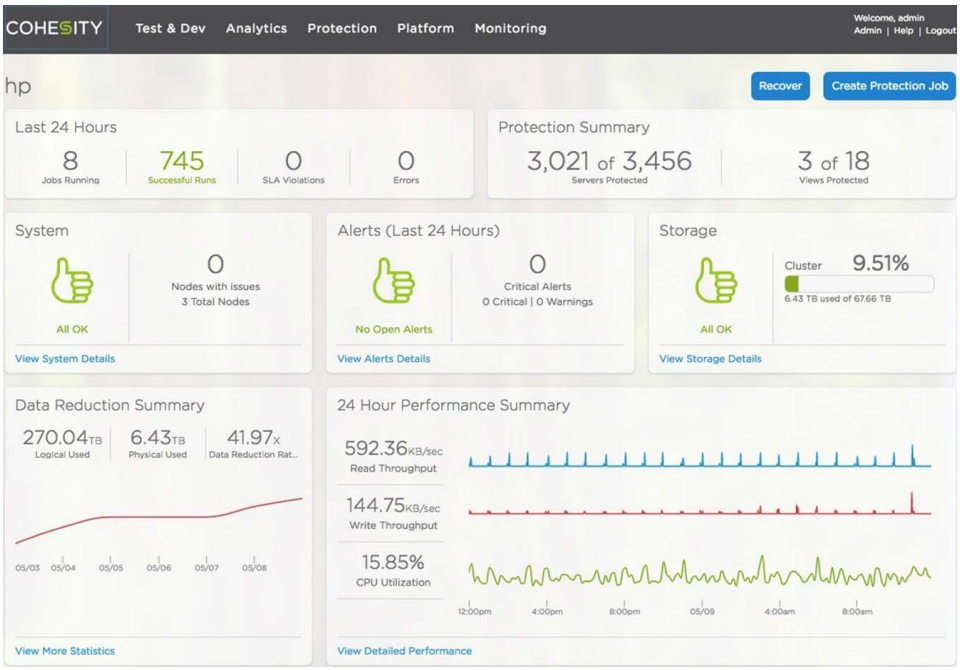

Figure 4. Cohesity Dashboard

Partition

A Cohesity Partition defines the set of physical nodes assigned to a logical volume. A partition uses the nodes assigned to it for physical storage. Each Cohesity Partition can map to a logical partition within your organization such as a department or group. You assign nodes to partitions based on the expected workload. Partitions are configured with a hostname and virtual IPs (VIPs) for external access using well-known protocols such as NFS and SMB. The VIP based design of the partition allows the system to be fully resilient and available even when a network or node level failure occurs.

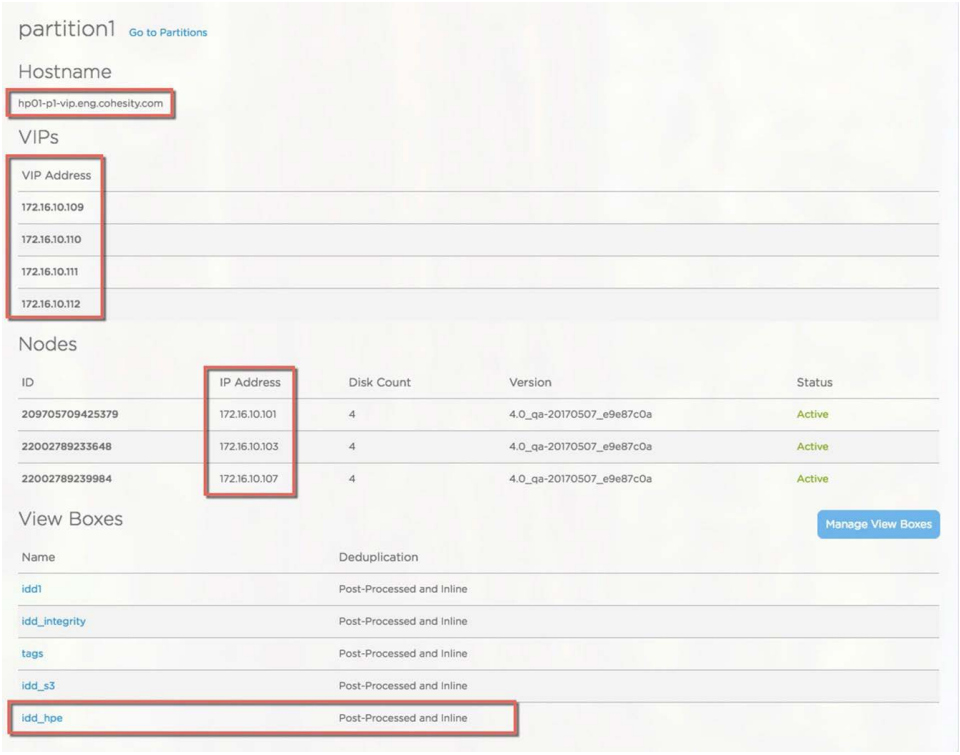

Figure 5. Cohesity Partition configuration

View Boxes

A View Box is a named storage location on a partition. In addition, a View Box defines both the storage policy and when the deduplication process occurs. When a Protection Job is configured, the user must specify a View Box from which the Cohesity Cluster can store Snapshots when executed.

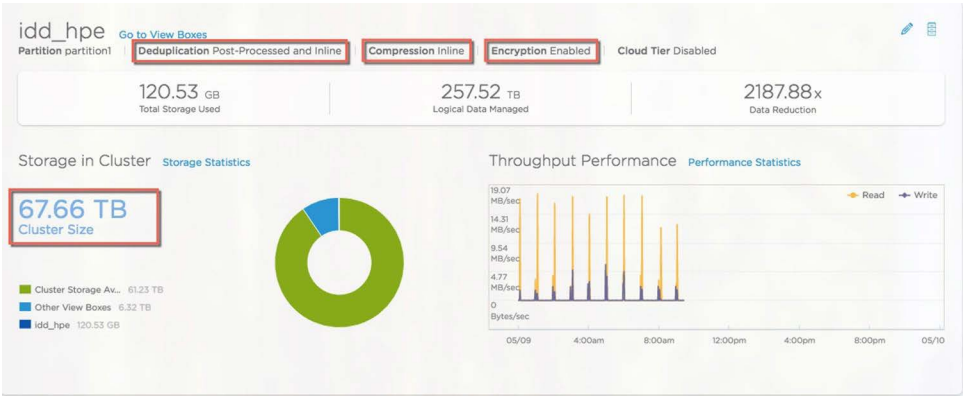

Figure 6. Cohesity view Box configuration

Views

A View represents the mount points into a specific View Box. Views provide NFS or SMB/CIFS protocol access for files, snapshots, and clones. QoS can be set for each view that can tune performance for the target workload.

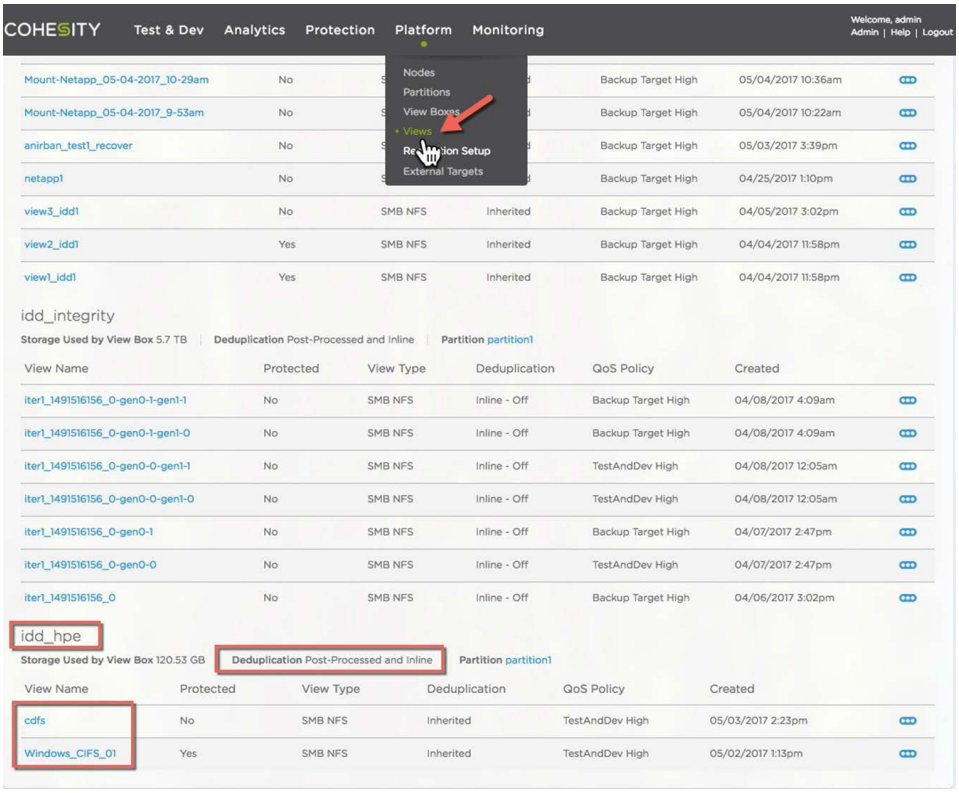

Figure 7. Cohesity view configuration

Use cases

Use case 1: Recovering files and VMs in VMware® ESX®

To begin the use case, VMs are registered and fully protected.

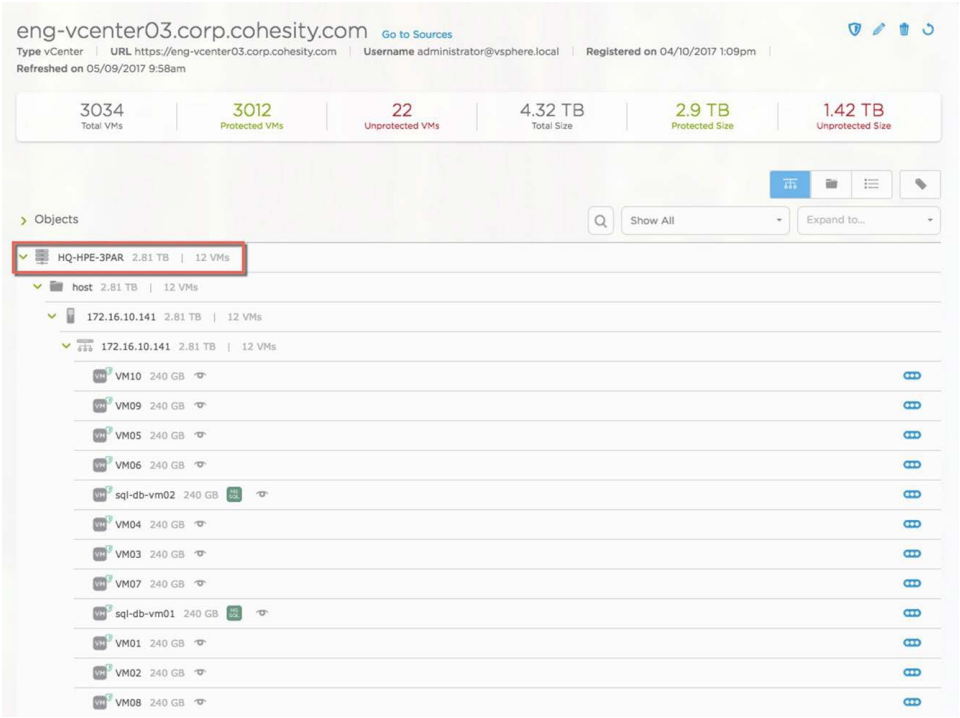

Figure 8. VM inventory used

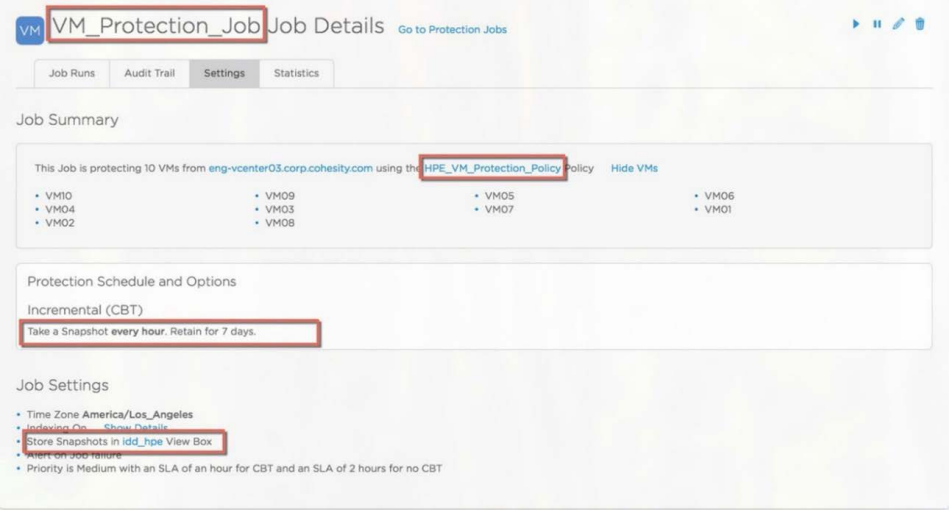

Figure 9. VM data protection policy

VM Recovery

In the following example we'll demonstrate the simplicity of restoring a VM from Cohesity. We will "accidentally" delete VM04 from VMware® vCenter™ and follow a simple recovery procedure to recover the VM instantaneously.

1. Simulate an accidental deletion of a VM.

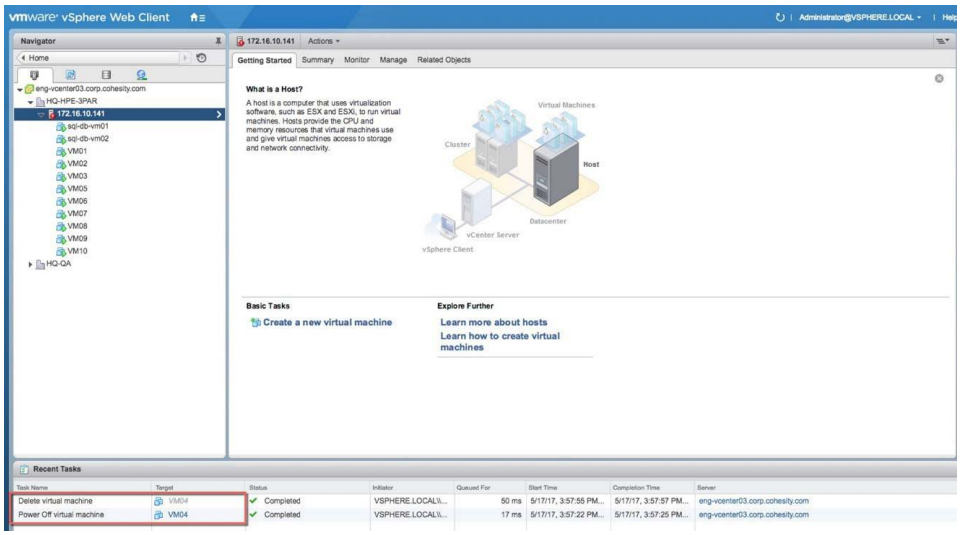

Figure 10. vCenter showing VM being deleted

2. Select "VMs" as the recovery option. This will recover VMs back to the source vCenter or even a different vCenter.

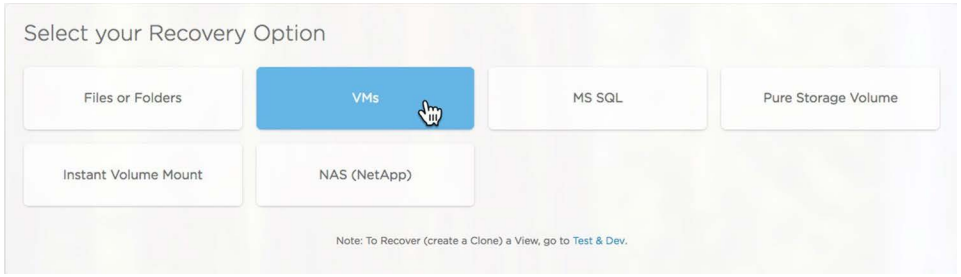

Figure 11. VM Recovery Console

3. Search for the VM to be recovered using regex search function and add it to cart.

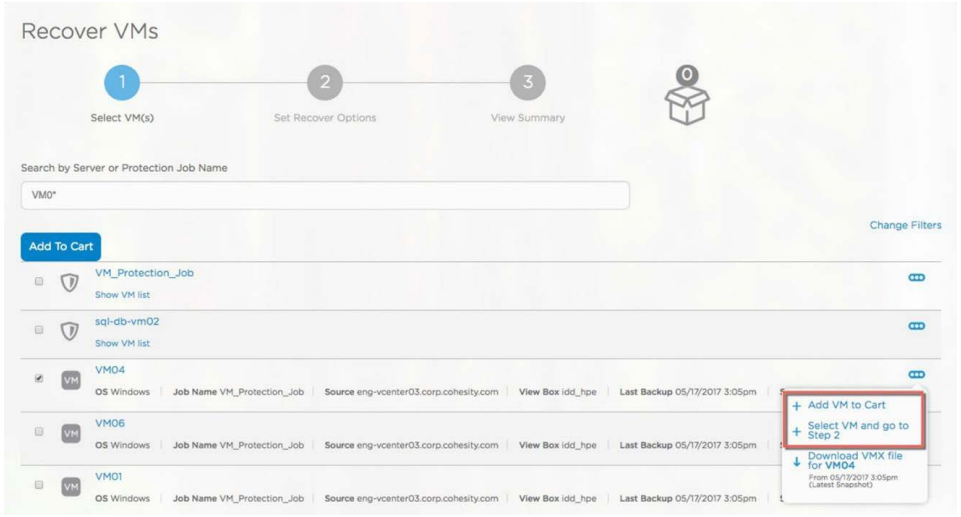

Figure 12. VM Recovery Workflow

4. Click the snapshot instance to modify the RPO.

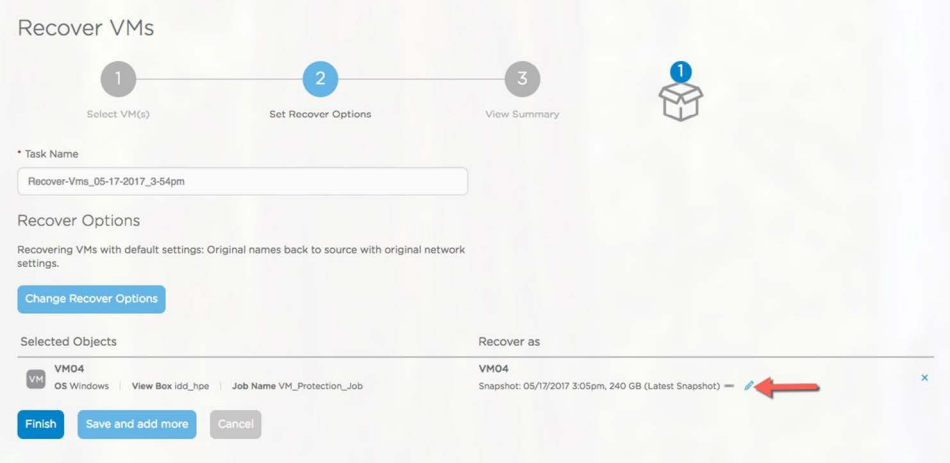

Figure 13. VM Recovery Workflow—Continued

5. Select the RPO from the recover points for the given VM.

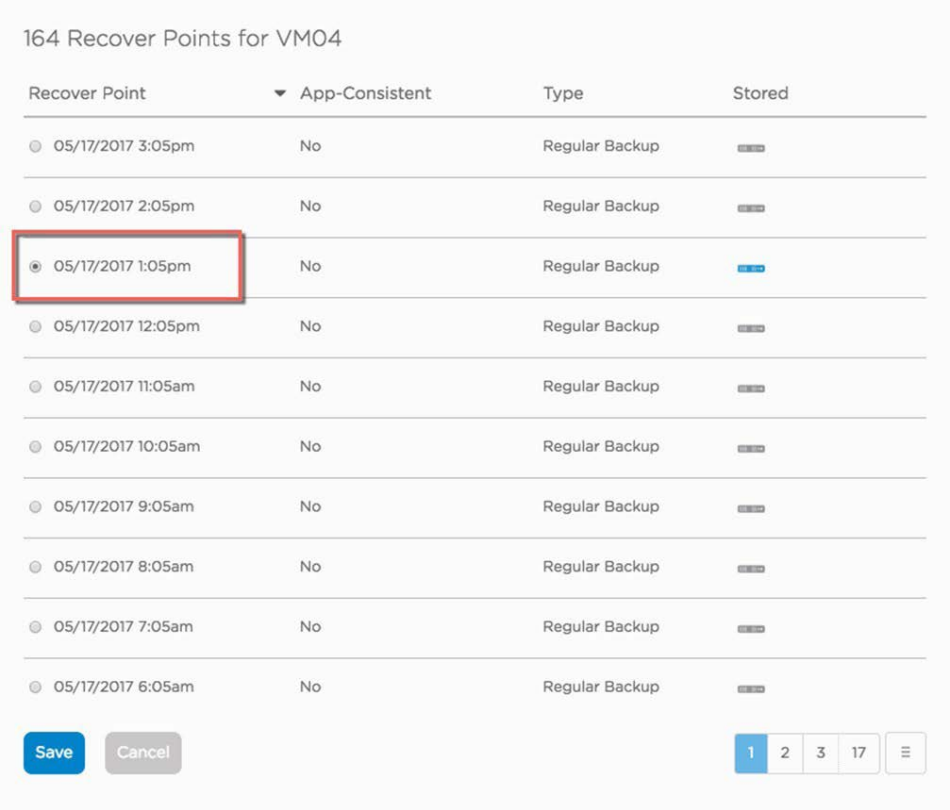

Figure 14. VM Recovery Workflow—Continued

6. Click Finish to start the recovery process.

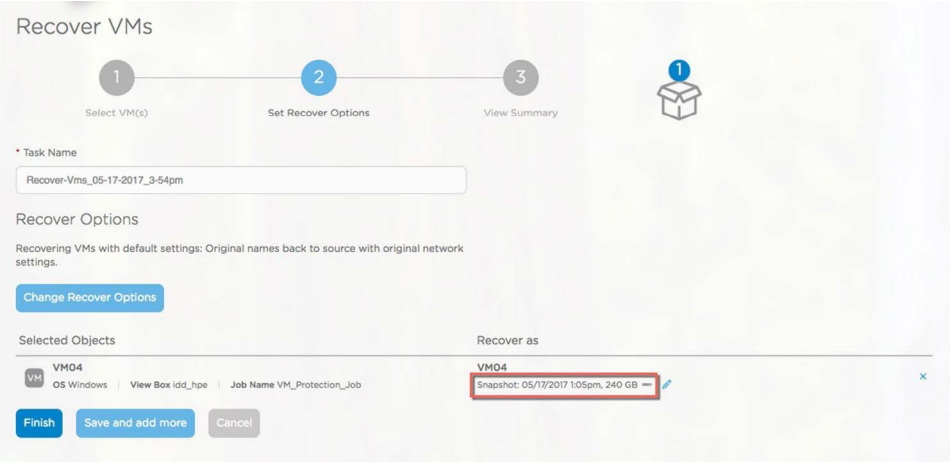

Figure 15. VM Recovery Workflow—Continued

7. Validate the VM has been restored.

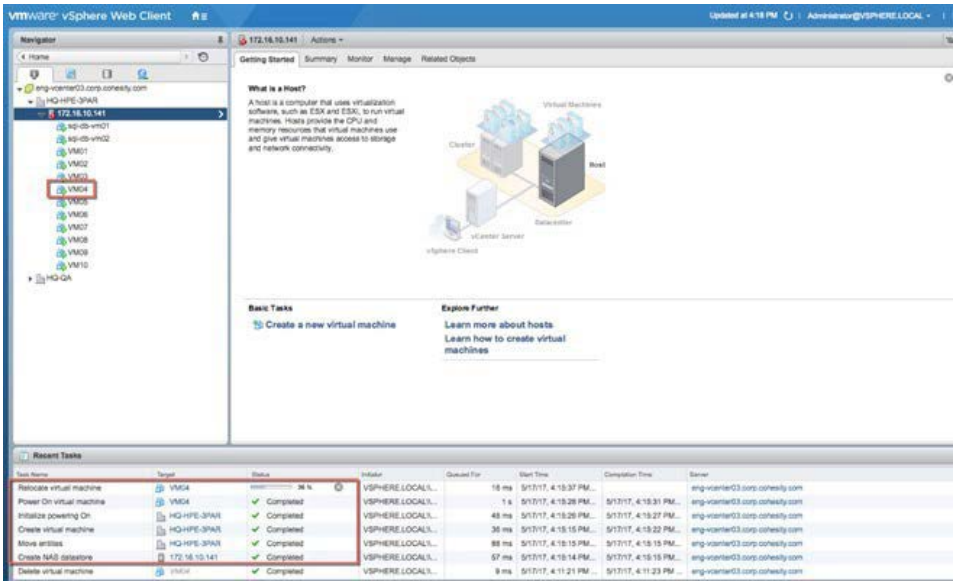

Figure 16. VM Recovery Workflow—Continued

The following tasks are performed for a successful recovery of the VM.

- User manually triggers a Cohesity VM Recovery task with the selected snapshot, target, networking settings, VM name, and target datastore.
- Cohesity contacts VMware® endpoint to validate current inventory and chosen recovery task settings.
- Cohesity creates an internal view and clones the VM snapshot and mounts the view to the target ESXi host(s).
- A new VM object is created using the original VM configuration file and chosen recovery settings. Network configuration changes take place at this step.
- VM is (optionally) powered on (Note that the VM is now available for use).
- Storage vMotion is initiated to move the datastore from the Cohesity cluster to the primary datastore.
- Storage vMotion completes, VMware non-disruptively migrates datastore access from the Cohesity cluster snapshot to the primary datastore.
- Cohesity requests the datastore to unmount.
- ESXi host unmounts datastore.
- Cohesity releases the view.

8. VM is powered on and recovers from Cohesity while it's being relocated back to the HPE 3PAR primary storage in the background. The application comes online as part of VM startup and is available all the while during the relocation of the VM.

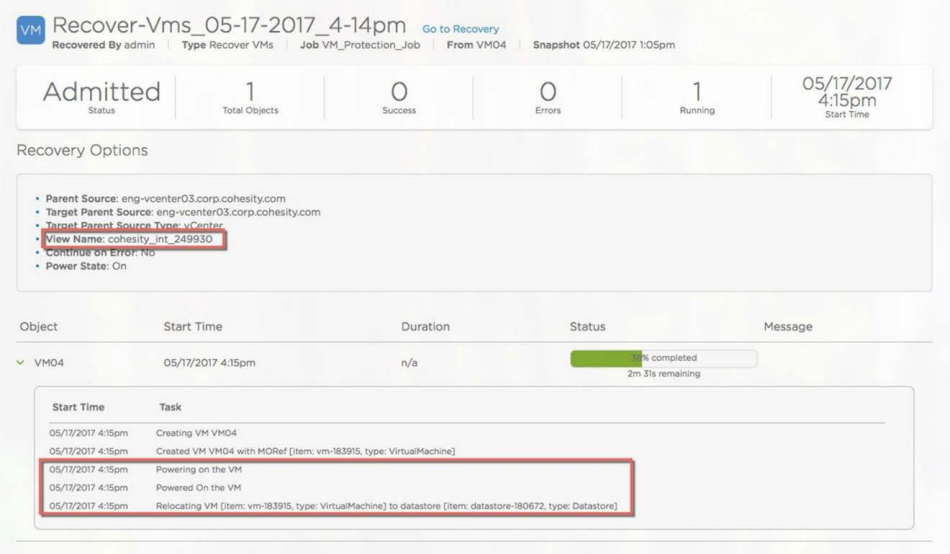

Figure 17. VM Recovery Workflow—continued

Single file level recovery for VMs

In the following example we'll see the simplicity of recovering a file using Cohesity.

1. Choose "Files or Folders" as the Recovery Option.

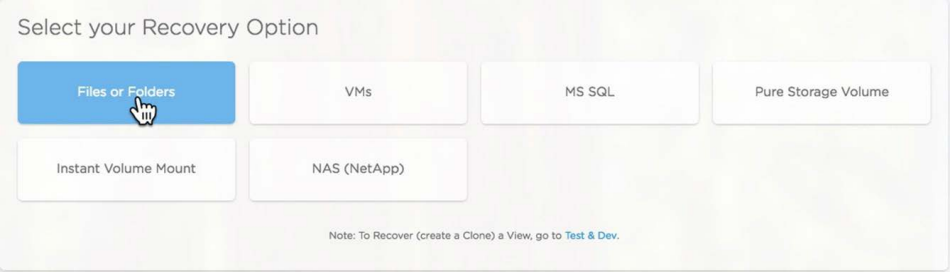

Figure 18. File/Folder Recovery Console

2. Perform a Global Search across the entire data set. Next filter down to the VM or physical server to choose the file

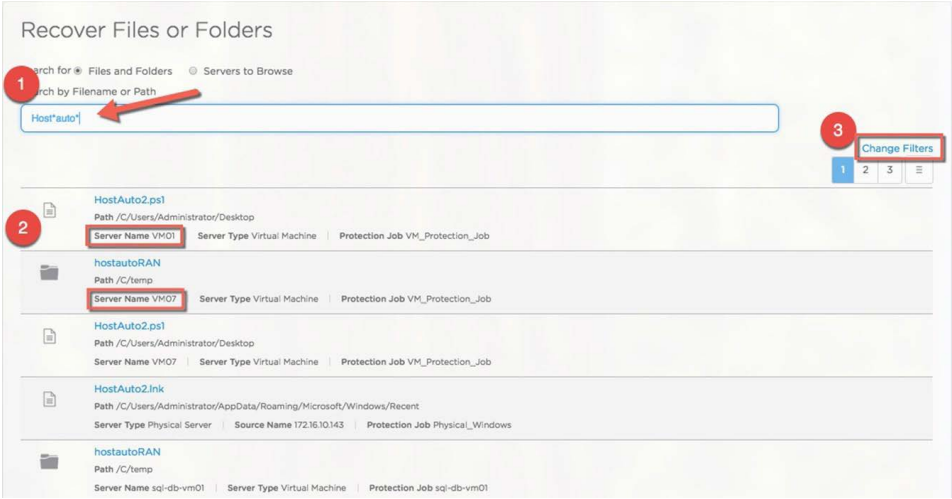

Figure 19. File/Folder Recovery workflow

3. Choose the VM where the file resides.
4. Choose the desired RPO and click "Recover to Server."

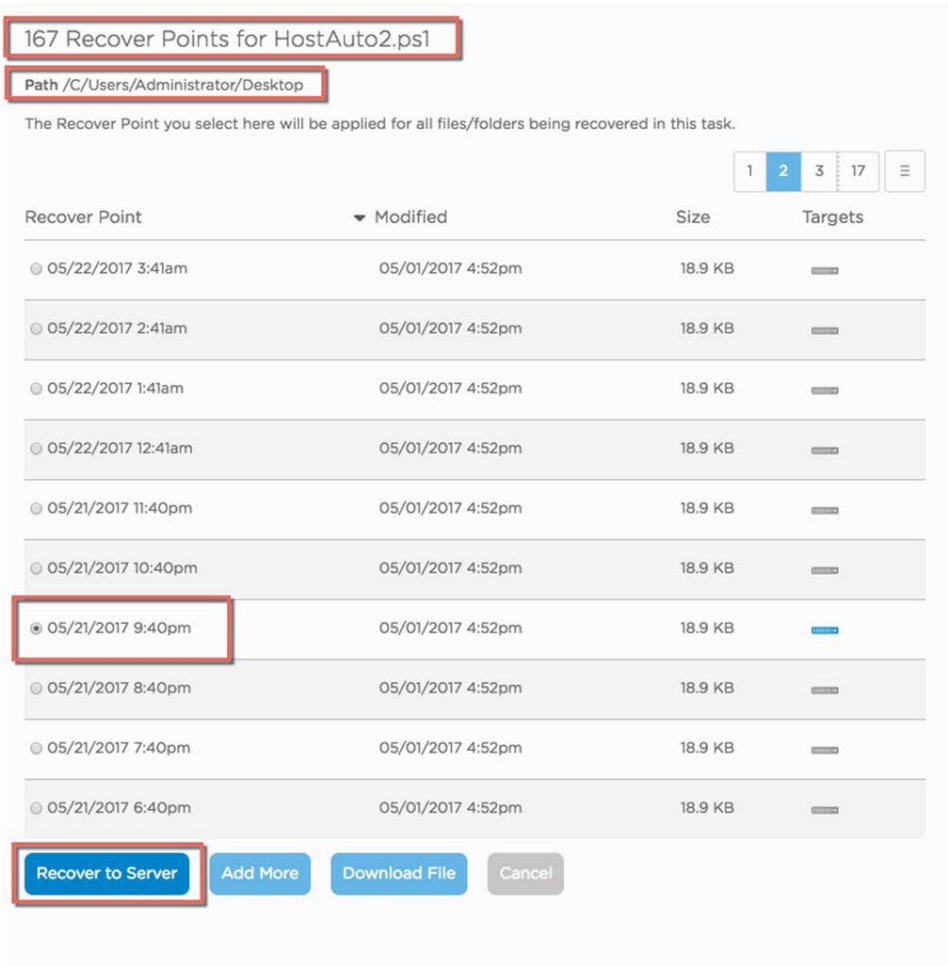

Figure 20. File/Folder Recovery workflow—continued

5. Determine whether to recover to the original location and/or overwrite any existing files.

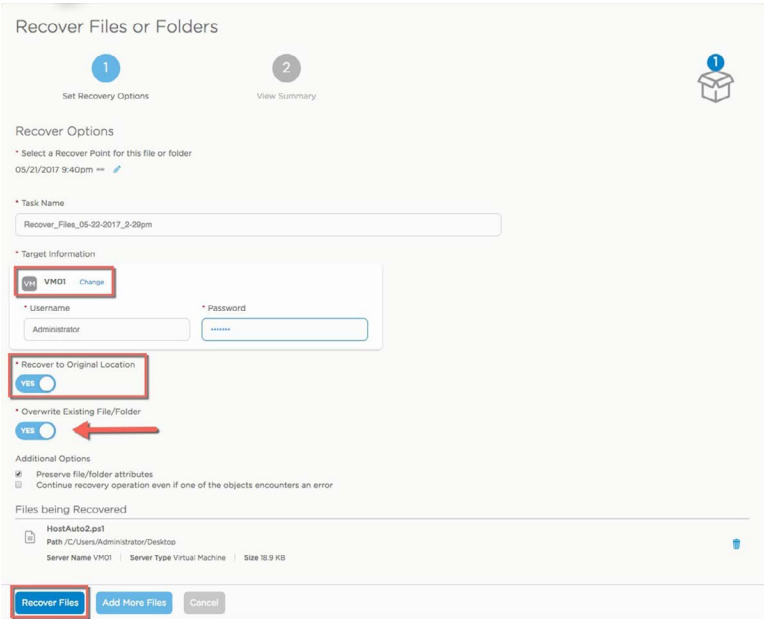

Figure 21. File/Folder Recovery workflow—continued

6. Click “Recover Files” and the following steps are performed by Cohesity during the file/folder recovery:
   - User manually triggers a file/folder recovery task either by searching for files through the elastic search database or via browsing the VMs and their volumes.
   - Cohesity creates an internal view, clones the VM snapshot, and mounts the view to the target ESXi host(s).
   - Cohesity attaches the cloned VMDK files to the target VM from which the files are being recovered.
   - Cohesity deploys a helper utility onto the VM and triggers the restore process.
   - The restore helper utility performs file copy from the attached disks (originally from the backup) onto the recovery location. The utility additionally preserves the file attributes and other properties based on the user preferences.
   - Once the file/folder copy completes, the disks are detached from the VM.
   - Cohesity requests the datastore to unmount.
   - ESXi host unmounts the datastore.
   - Cohesity releases the view.

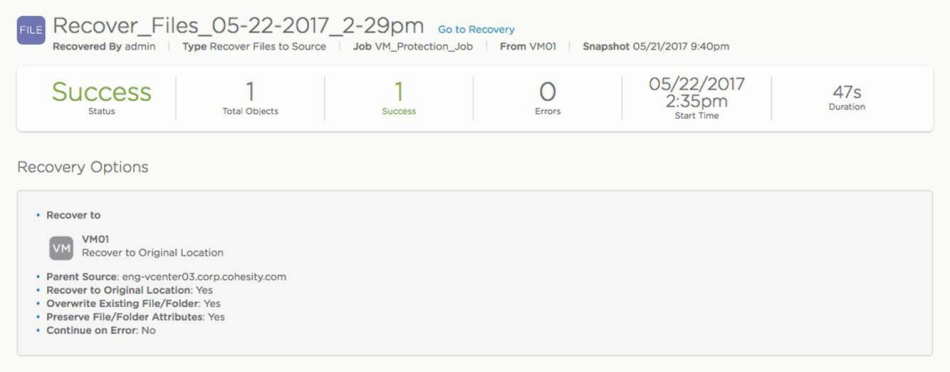

Figure 22. Recovery successful

Use case 2: Microsoft SQL Integration, Protection, and Recovery in VMware ESX

MS SQL Protection

To begin the use case, the SQL Server must be registered. We will then use the policy engine to protect it.

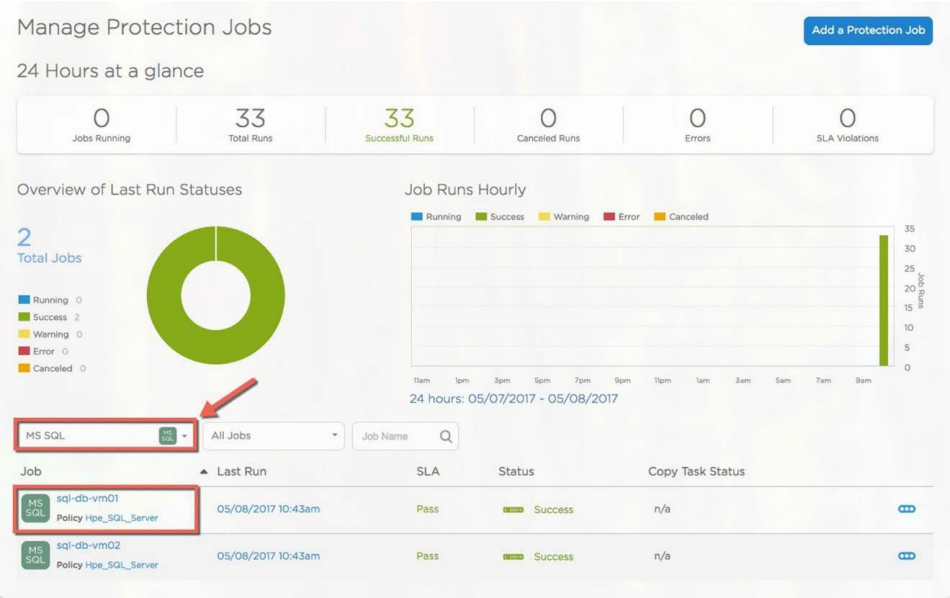

Figure 23. Filter to see MS SQL servers being protected

The following Policy is used to protect the 2 MS SQL database VMs in this example. The SQL Policy dictates the database servers to be backed up every hour with log capture every 15 mins.

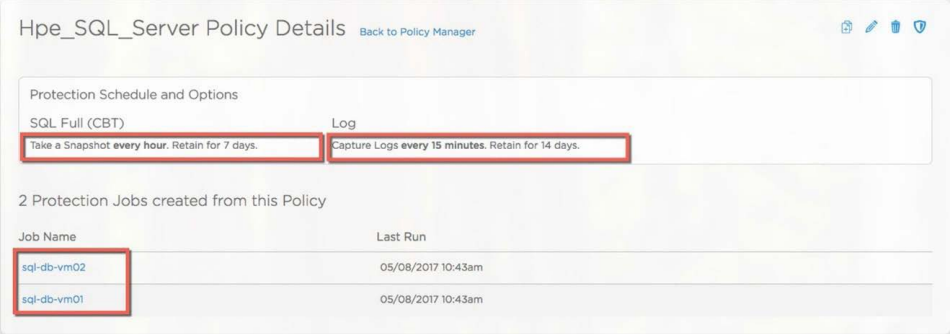

Figure 24. MS SQL Data Protection Policy

MS SQL Recovery

In the following steps we will walk through the database recovery process.

1. Search for the database to be restored.
2. Validate the database name and associated job name.

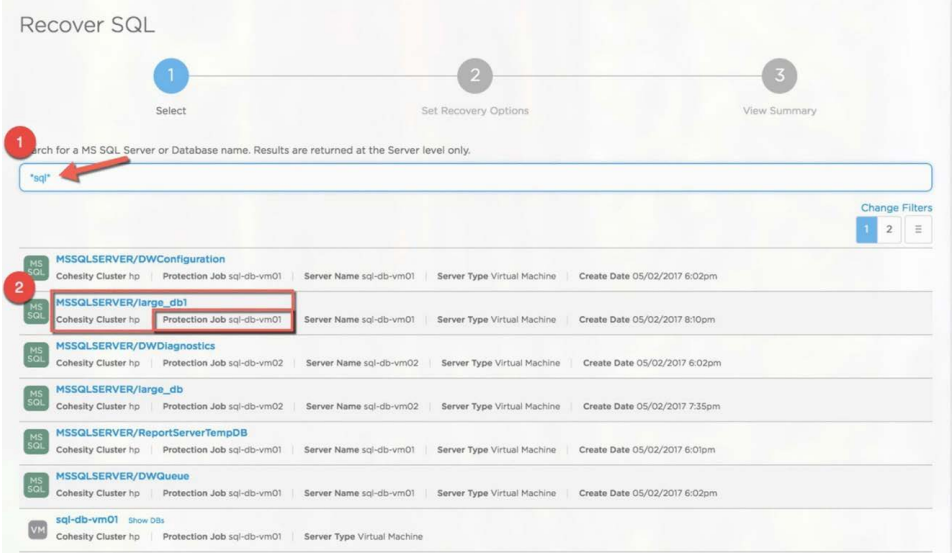

Figure 25. MS SQL Database recovery

3. Edit the restore point from available backups of the database.

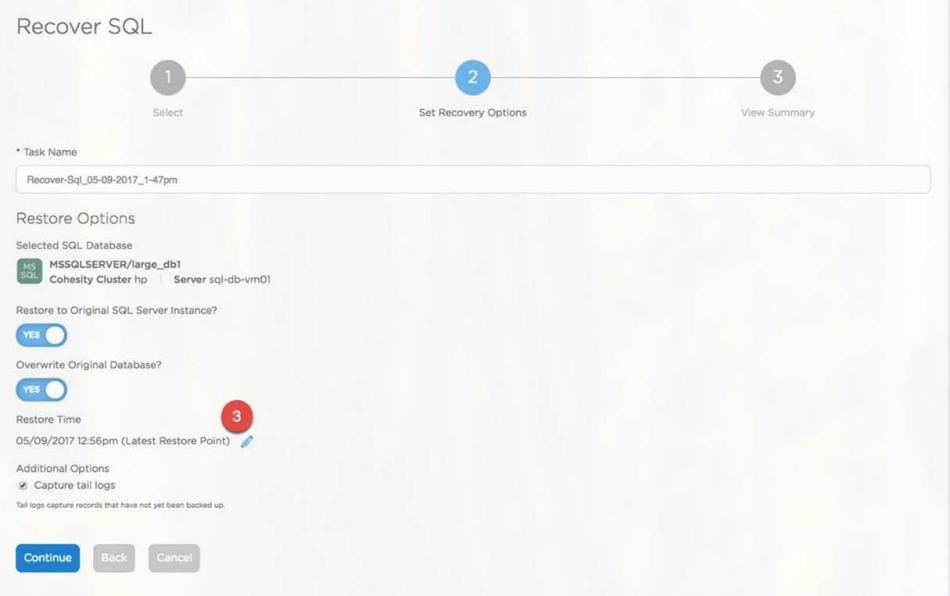

Figure 26. MS SQL Database recovery—continued

4. Choose the required RPO.
5. Select to replay the logs back to the required time.

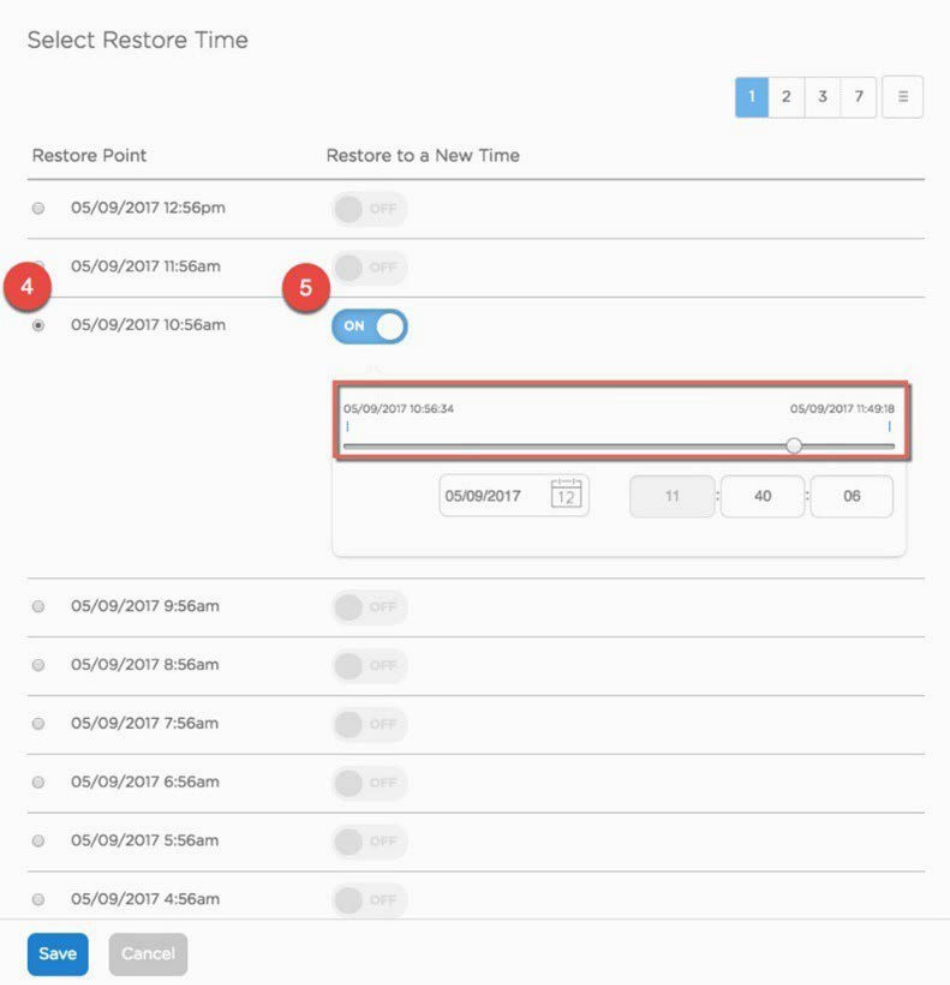

Figure 27. MS SQL Database recovery—continued

While restoring an individual MS SQL Server database, you can restore the database to its original VM and MS SQL Server instance. You can also restore to a non-original location such as a different VM or a different MS SQL Server instance on the original VM.

The example above assumes restoring to the database's original location and that the database failure has already occurred. The database's MS SQL Server VM is registered as an MS SQL Server and was backed up by a Cohesity Protection Job that includes transaction logs.

Restoring an individual MS SQL Server database to its original location and overwriting the existing database consists of the following actions: The Cohesity user must select the following task options:

- Restore to the original SQL Server VM and instance and overwrite the database
- VMware VM Snapshot and individual MS SQL Server database
- Point in time (optional)
- Capture tail logs (optional)
- The Cohesity Cluster copies the database from the specified Snapshot to the original location
- If applicable, transaction logs are applied to restore the database to the specified point in time
- The database is now restored to its specified location

Use case 3: Protecting data on Microsoft Windows and Linux Physical Servers

In the following example Windows® Physical servers are protected using Cohesity physical agents. Installing agents and registering physical servers is covered in the product documentation. Here we are filtering down to the physical servers, and validating the data protection policy used for these servers.

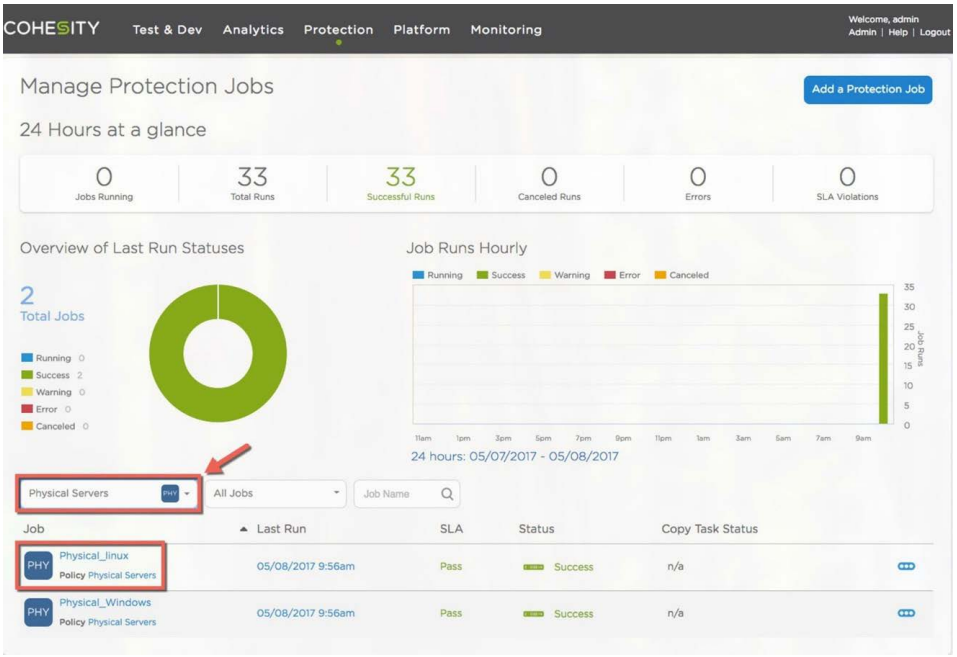

Figure 28. Filter to find Physical Servers being protected

The following policy shows physical Linux and Windows Server® being protected every hour and data being retained for 30 days.

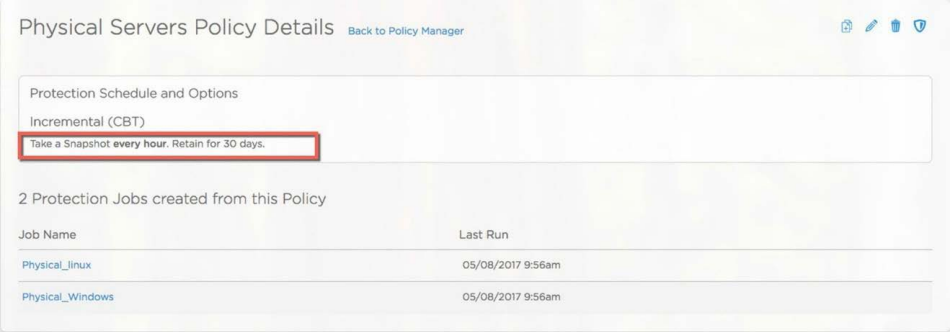

Figure 29. Physical server data protection policy

Now let's walk through the process of recovering files/folders from Cohesity using the Instant Volume mount option. Some benefits of this restore mechanism include granular recovery of Microsoft Exchange, SQL, and SharePoint data. Cohesity's simple workflows can instantly mount backup volumes to the Kroll Ontrack® PowerControls™ application, enabling application owners to perform granular recovery.

Note: Files and folders can also be restored back to physical servers using the workflow of searching globally for files and restoring back to physical servers directly. This workflow is the same as file/folder recovery showcased for VMs above.

1. Choose Instant Volume Mount as the recovery option.

Figure 30. Physical server File/Folder recovery using Instant Volume Mount

2. Search for Physical server protected by the job name or server name.

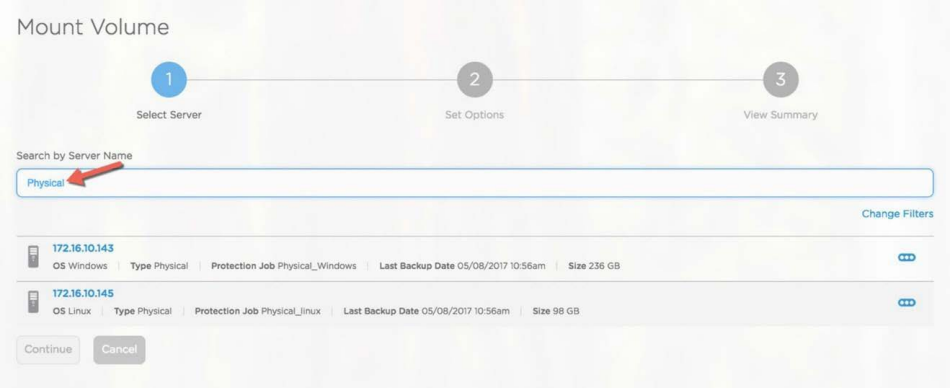

Figure 31. Physical server File/Folder recovery

3. Select the volumes and the restore point for the recovery operation as shown below.

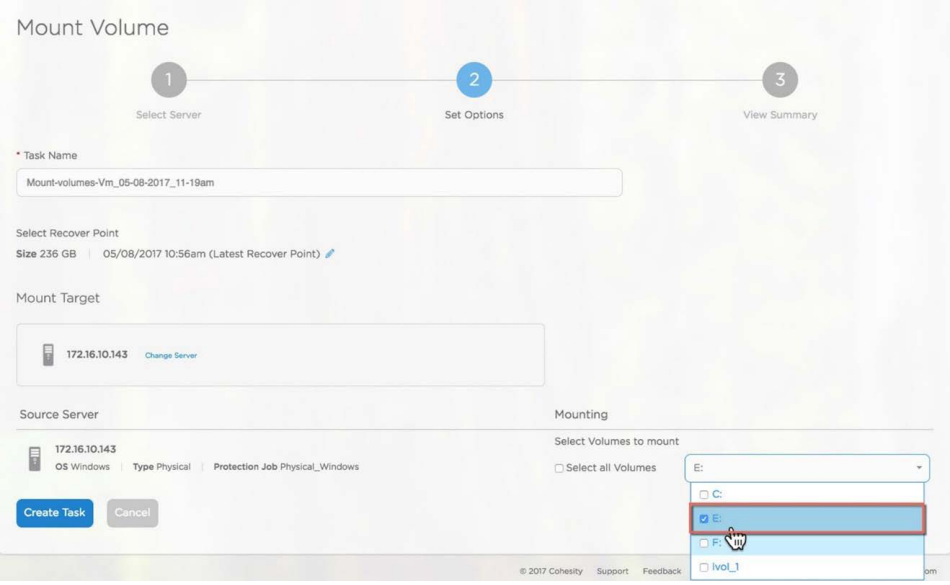

Figure 32. Physical server File/Folder recovery—continued

4. As the server was protected every hour according to the policy, we have the option to select from any recovery point.

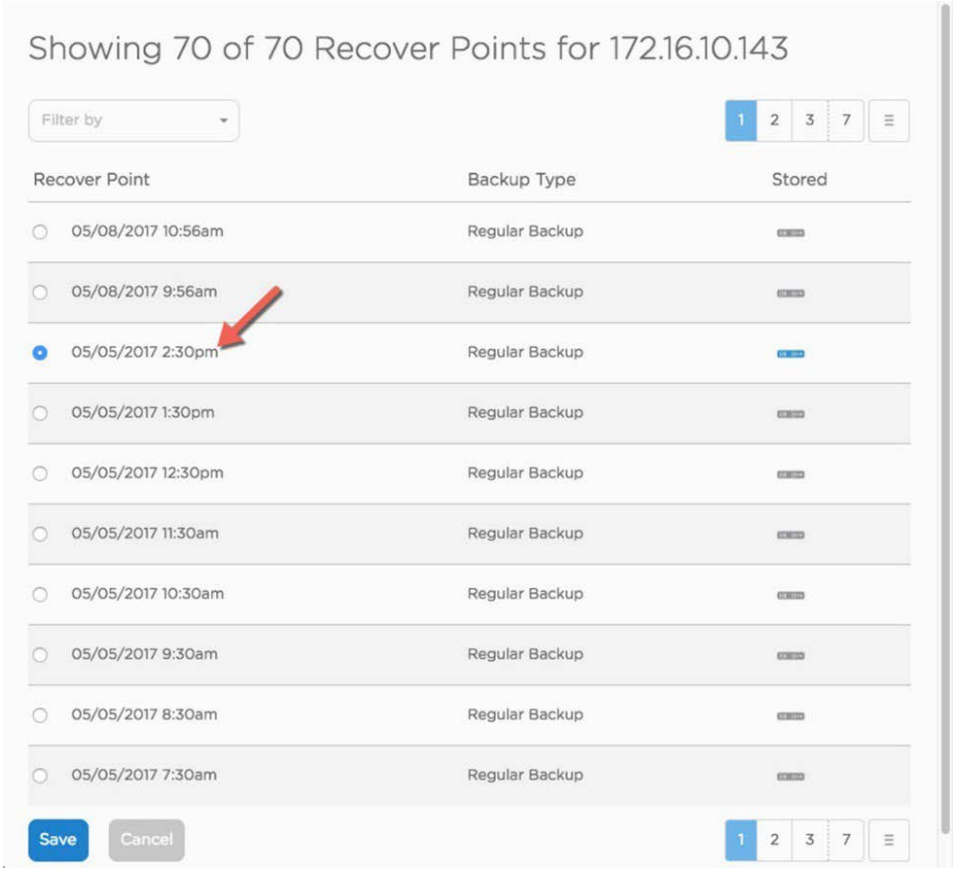

Figure 33. Physical server File/Folder recovery—continued

5. Once the target server has been selected, the task can be executed.

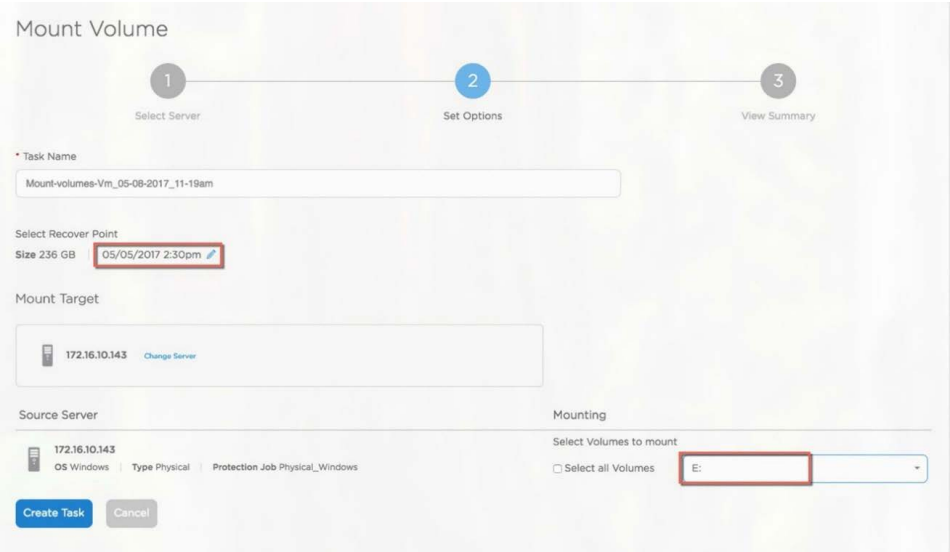

Figure 34. Physical server File/Folder recovery—continued

6. As seen below, the mount point is presented as a directory structure to the physical server and the administrator now has all the backed up files on the server to perform granular recovery.

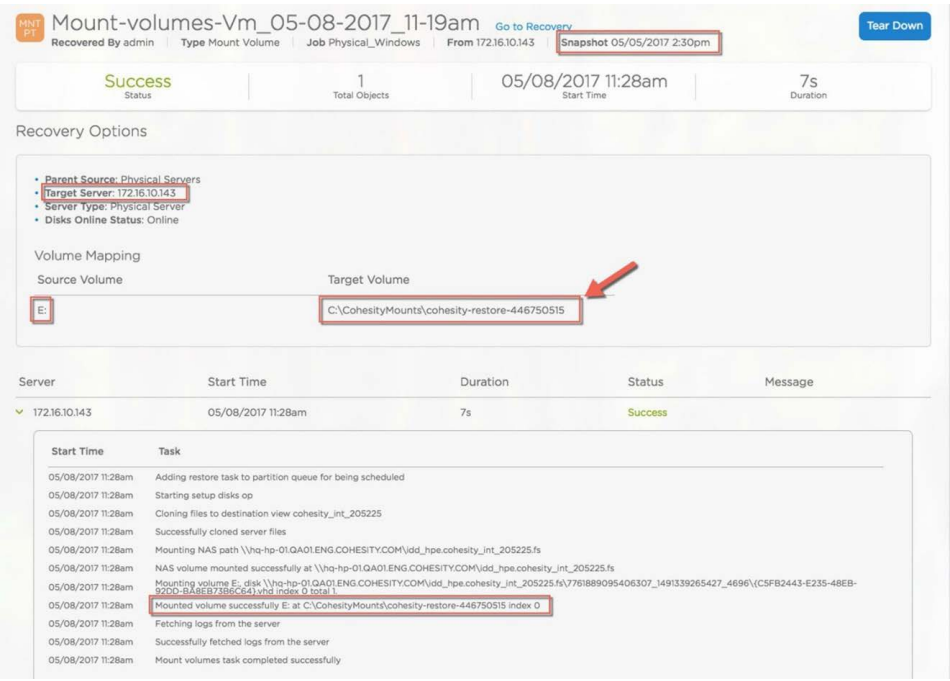

Figure 35. Physical server File/Folder recovery—continued

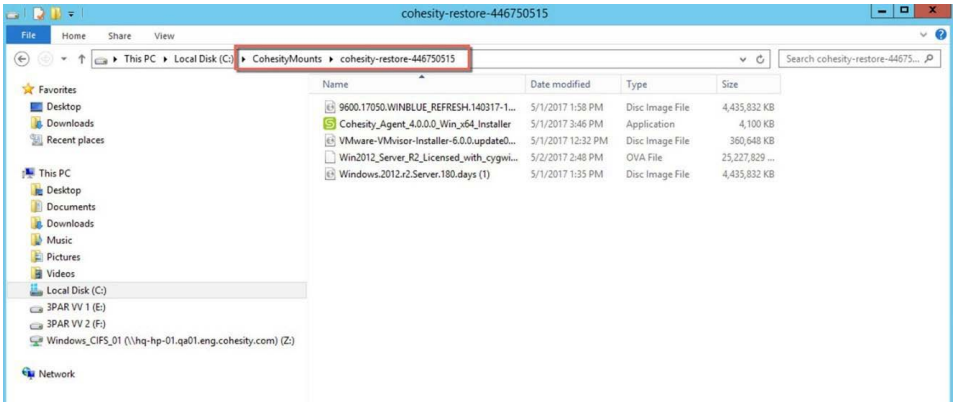

Figure 36. Physical server File/Folder recovery—continued

7. Once the restore operation has been completed, the mount point which was presented can be removed from Cohesity. This would gracefully unmount the volume and remove the recovery directory structure.

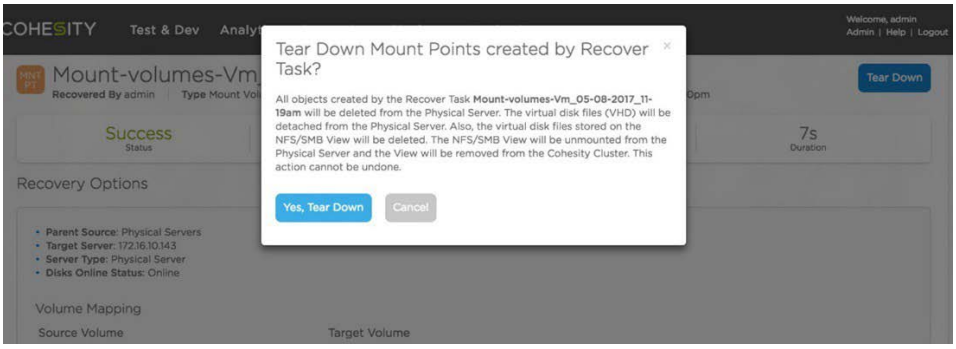

Figure 37. Physical server File/Folder recovery—continued

Use case 4: Creating a CIFS Share on secondary storage

Map the View (mentioned above in the View section) Windows_CIFS_01 as a mapped CIFS drive for light file services use cases to Microsoft Windows systems.

1. Click on the Windows_CIFS_01 view and get the path from the Cohesity UI.

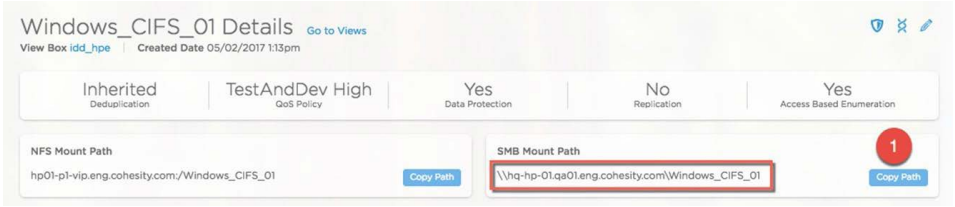

Figure 38. CIFS View mount to be used as CIFS mapped drive

2. Choose “Map network drive.”

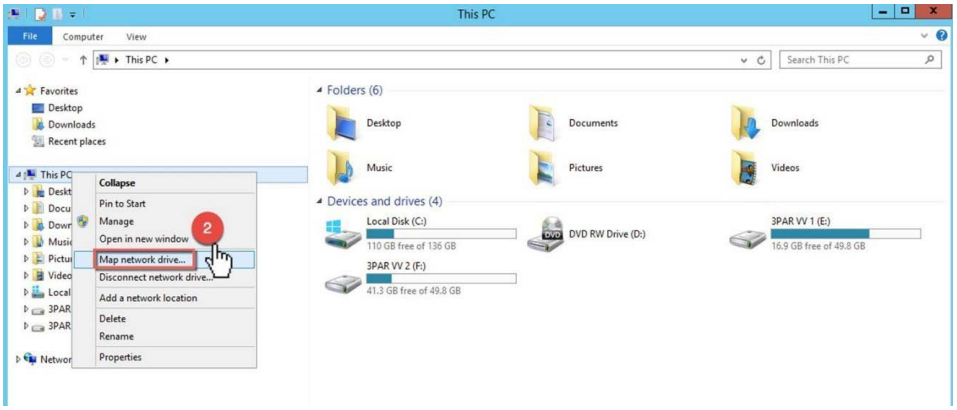

Figure 39. Map network drive

3. Specify a drive letter and folder.

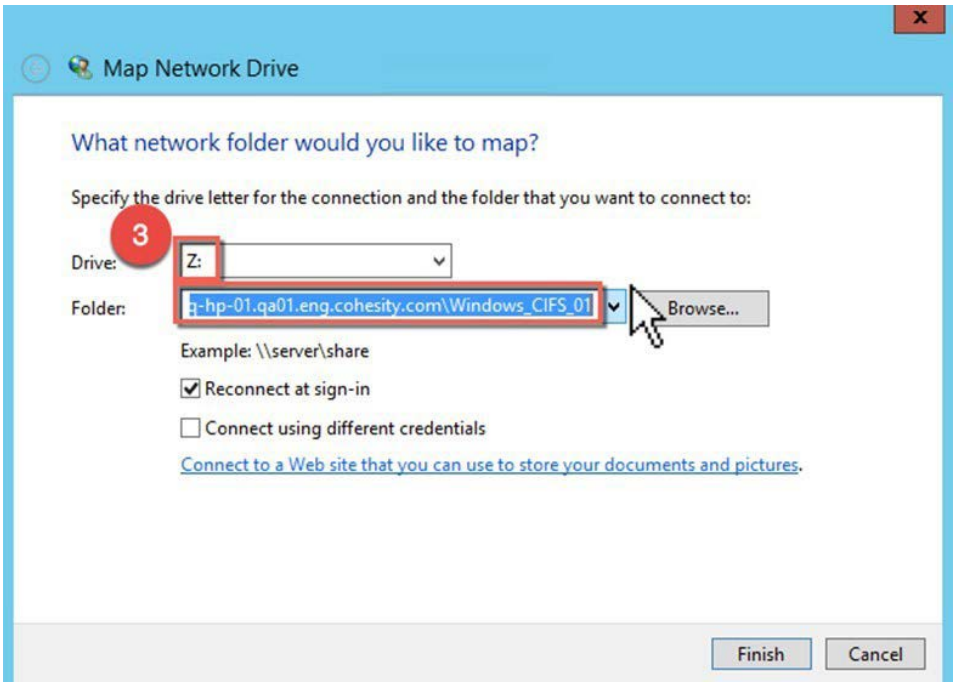

Figure 40. Provide the drive information and the network folder information

4. Validate that the drive is mapped on the host.

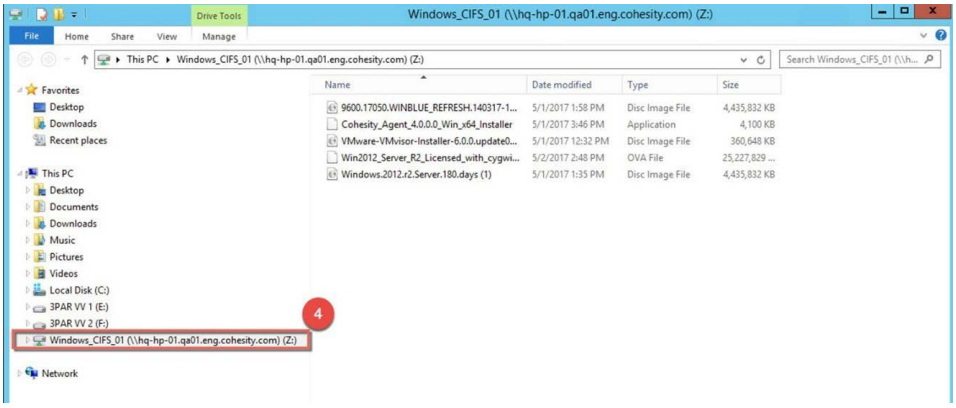

Figure 41. Network drive mapped in the Windows host

This mapped drive can now be used for light file services use cases, and the View's data can be protected with Cohesity snapshots. This View based data protection is configured by using the data protection policy-based constructs.

Conclusion

We have now demonstrated the seamless integration of Cohesity with HPE 3PAR StoreServ primary storage. This paper showcased Cohesity's software defined secondary storage solution and the ease in which today's modern data center engineer can deploy, manage, recover, and expand from tier-1 into tier-4 workloads.
The following backup and recovery procedures were covered:

- Backup and Restore single or multiple VM(s)
- Backup and Restore single or multiple file(s)
- Backup and Restore single or multiple physical Windows and Linux Server(s)
- Backup and Restore a SQL database

With this joint integration, HPE and Cohesity deliver the simplicity to manage data at scale without the fragmentation and inefficiencies of silos and data lakes.

Learn more at [hpe.com/us/en/storage/hpe-complete.html.](hpe.com/us/en/storage/hpe-complete.html)

Cohesity, Inc.
Address 300 Park Ave., Suite 300, San Jose, CA 95110
Email contact@cohesity.com www.cohesity.com
@cohesity